

000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 INFIMED: LOW-RESOURCE MEDICAL MLLMs WITH ADVANCING UNDERSTANDING AND REASONING

Anonymous authors

Paper under double-blind review

ABSTRACT

Multimodal Large Language Models (MLLMs) have achieved remarkable progress in domains such as visual understanding and mathematical reasoning. However, their application in the medical domain is constrained by two key challenges: (1) multimodal medical datasets are scarce and often contain sparse information, limiting reasoning depth; and (2) Reinforcement Learning with Verifiable Rewards (RLVR), though effective in general domains, cannot reliably improve model performance in the medical domain. To overcome these challenges, during the supervised fine-tuning (SFT) stage, we incorporate high-quality textual reasoning data and general multimodal data alongside multimodal medical data to efficiently enhance foundational medical capabilities and restore the base model’s reasoning ability. Moreover, considering that there are some multimodal medical datasets with sparse information, we further synthesize reflective-pattern-injected chain-of-thought (CoT) in addition to general CoT samples, equipping the model with initial reflective reasoning capabilities that provide a structured foundation for subsequent RLVR training. Finally, we introduce our InfiMed-Series models, InfiMed-SFT-3B and InfiMed-RL-3B, both of which deliver state-of-the-art performance across seven multimodal medical benchmarks. Notably, InfiMed-RL-3B achieves an average accuracy of 59.2%, outperforming even larger models like InternVL3-8B, which achieves 57.3%. Specifically, during the SFT phase, we utilized 188K samples, while the RLVR phase incorporated 36K samples, demonstrating the efficacy of both training strategies in achieving superior performance. We also conducted a series of extensive experiments, which provide valuable insights that contribute to advancing the performance of MLLMs in medical scenarios.

1 INTRODUCTION

The rapid development of multimodal large language models (MLLMs) in recent years has marked a transformative phase in artificial intelligence, driving substantial progress across diverse domains. Notably, MLLMs have achieved significant breakthroughs in areas such as object recognition (Yin et al., 2025; Liu et al., 2025d), mathematical reasoning (Zhuang et al., 2025; Peng et al., 2024; Liu et al., 2025b), and graphical user interface (GUI) interaction (Liu et al., 2025a; Luo et al., 2025; Qin et al., 2025), largely attributable to the availability of abundant high-quality multimodal datasets. In contrast, the medical domain remains particularly challenging due to the scarcity of high-quality multimodal data, which severely limits the performance of MLLMs in medical scenarios.

To enhance the medical reasoning capabilities of MLLMs, prior work has primarily relied on large-scale, domain-specific supervised fine-tuning (SFT). For instance, LLaVA-Med (Li et al., 2023) directly utilizes the PMC-15M (Zhang et al., 2023b) dataset for medical concept alignment and instruction following. However, its performance is constrained by the inherent noise of the dataset and the limited amount of reasoning information it provides. Recent studies, such as MedGemma (Sellergren et al., 2025), collect larger and higher-quality medical datasets that cover both textual and multimodal modalities, aiming to further enhance the general medical capabilities of MLLMs. While SFT can be effective, it is highly data-intensive and mainly focuses on memorizing training data (Chu et al., 2025). Building on the success of DeepSeek-R1 (Guo et al., 2025), Reinforcement Learning with Verifiable Rewards (RLVR) has shown significant improvements in exploration and generalization for multimodal tasks (Zhang et al., 2025; Liu et al., 2025a;c). RLVR, which often includes a “cold-start” phase in (MLLMs)(Huang et al., 2025; Peng et al., 2025; Liu et al., 2025a), is also

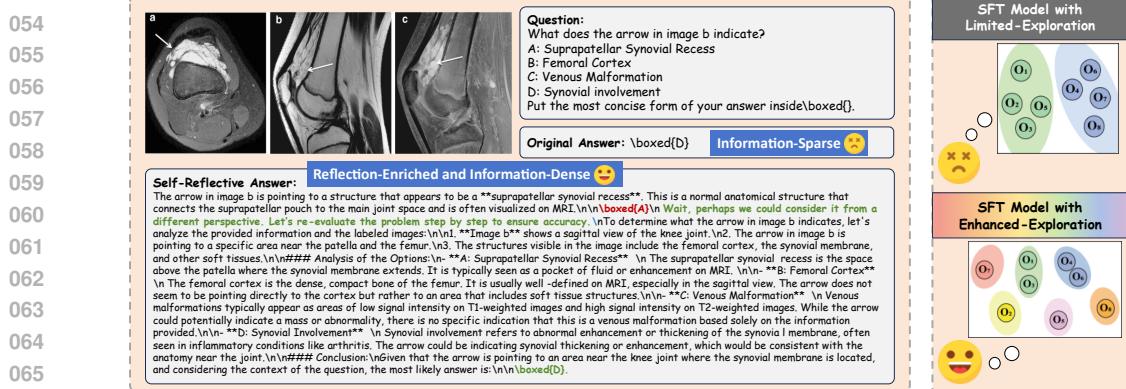


Figure 1: **Left:** Comparison of information-sparse and reflection-enriched, information-dense outputs. **Right:** A model with Enhanced Exploration (bottom) generates a broader, more effective search space, while Limited Exploration (top) results in a narrower, less efficient search.

beginning to find notable applications in medical scenarios (Su et al., 2025; Xu et al., 2025; Pan et al., 2025).

Despite these ongoing efforts, existing approaches still exhibit notable limitations, which can be summarized into two key challenges. First, the scarcity of high-quality multimodal medical datasets remains a bottleneck: most existing datasets suffer from sparse information and contain limited explanatory information, which hinders effective model training and results in poor reasoning performance, as shown in Figure 1. Second, although RLVR has been shown to substantially enhance model performance in other domains, its application in medical scenarios remains underexplored. Existing work either lacks extensive exploration across broad benchmarks (Pan et al., 2025; Su et al., 2025) or fails to effectively improve model performance (Xu et al., 2025).

To address the challenges mentioned above, during the SFT stage, we leverage not only multimodal medical data but also general multimodal data to preserve the model’s visual perception capabilities, while integrating medical textual data to enhance its domain-specific knowledge. Additionally, we introduce a novel synthesis of reflective-pattern-injected chain-of-thought (CoT) data, effectively addressing the information sparsity present in certain multimodal medical datasets. This approach could also provide a more robust exploratory foundation for subsequent RLVR, enabling a cold-start method with limited resources. Building upon this, we train our InfiMed-SFT-3B model on 188K samples, equipping it with both fundamental reasoning and reflective patterns. We then apply RLVR on top of InfiMed-SFT-3B using 36K samples to obtain InfiMed-RL-3B, further enhancing its exploration capabilities and generalization performance. Extensive experiments show that our InfiMed-series models set new SOTA performance across multiple multimodal medical benchmarks, outperforming similarly-sized models like MedGemma-4B-IT and larger models such as InternVL3-8B, demonstrating the effectiveness of our reflective SFT and RLVR approach. We also investigated the impact of data composition and reasoning strategies through a series of exploratory experiments, yielding valuable insights for the advancement of medical MLLM applications.

In summary, the key contributions of our work are as follows: **(1)** We synthesize reflective-pattern-injected CoT data, equipping the model with initial reflective capabilities and a stronger cold-start foundation for subsequent RLVR. **(2)** We employ a low-resource SFT with 188K samples, enabling the model to develop robust reasoning, comprehension, and reflective patterns. Subsequently, RLVR is applied with 36K samples, effectively boosting the model’s exploration capabilities and performance. **(3)** We introduce the InfiMed-series models, **InfiMed-SFT-3B** and **InfiMed-RL-3B**, which achieve SOTA performance among 3B-level MLLMs, with InfiMed-RL-3B outperforming models like MedGemma-4B-IT by 7.64%, and remain competitive even against 7B-level models.

2 RELATED WORK

2.1 MEDICAL MULTIMODAL LARGE LANGUAGE MODELS

In recent years, MLLMs have evolved rapidly and achieved remarkable progress across a wide range of domains, attracting increasing interest in their potential applications within the medical

108 field (AlSaad et al., 2024). Extensive research efforts have been devoted to enhancing MLLMs’
 109 ability to integrate heterogeneous medical data to support critical dimensions in healthcare. Inspired
 110 by the success of medical LLMs like HuatuoGPT (Zhang et al., 2023a), Apollo (Wang et al., 2024),
 111 and Med-PaLM series (Singhal et al., 2023; 2025), recent efforts have increasingly focused on
 112 extending LLM capabilities to multimodal medical. LLaVA-Med (Li et al., 2023) introduces a
 113 biomedical-specialized large language-and-vision model trained on a curated figure-caption dataset
 114 with self-instructed instruction-following data. The model highlights the potential of cost-efficient
 115 training strategies for domain-specific MLLMs. MedGemma (Sellergren et al., 2025) has shown
 116 strong generalization across medical vision-language and text-only tasks, demonstrating advanced
 117 medical understanding and reasoning on multimodal data. Lingshu (Xu et al., 2025) proposed
 118 a domain-specialized multimodal foundation model for medical, supported by a curated dataset
 119 enriched with medical VQA, CoT reasoning, and report annotations. While prior work has made
 120 notable progress in adapting MLLMs to the medical domain, many approaches depend on large
 121 model sizes and substantial computational resources, which limit their accessibility and scalability.
 122

2.2 REASONING IN MEDICAL LARGE LANGUAGE MODELS

124 Interpretable reasoning remains a central desideratum in medical AI, with recent efforts exploring
 125 general CoT prompting (Wei et al., 2022) and program-based logic (Chen et al., 2022) modeling.
 126 Although these approaches have shown potential, they typically rely on costly expert-curated annotations
 127 (Li et al., 2024b), which limits their scalability in real-world clinical settings. RL offers a
 128 compelling alternative by enabling emergent reasoning capabilities without requiring explicit super-
 129 vision, as demonstrated by recent models such as DeepSeek-R1 (Guo et al., 2025), which achieve
 130 notable improvements in reasoning with rule-based reward. Building on this paradigm, RLVR has
 131 been used to improve reasoning reliability, with Group Relative Policy Optimization (Shao et al.,
 132 2024) known for its efficiency and good performance. This method is now increasingly used to
 133 train MLLMs to improve their reasoning ability (Meng et al., 2025; Wang et al., 2025; Tan et al.,
 134 2025). With the success of RLVR, several work leverages it on medical MLLMs. MedVLM-R1 (Pan
 135 et al., 2025) employs RLVR to explicit reasoning in medical VQA, achieving strong performance and
 136 generalization. Its emphasis on reasoning highlights the role of RL in enhancing transparency and
 137 trustworthiness in clinical AI systems. GMAI-VL-R1 (Su et al., 2025) explores RLVR to enhance
 138 reasoning and reflection in multimodal medical models. By introducing a multi-agent reasoning data
 139 synthesis framework, the model outperforms prior models on some complex tasks. Lingshu (Xu et al.,
 140 2025) also leverages an RLVR paradigm, achieving strong performance across medical VQA, report
 141 generation, and text-only QA. Despite these promising advances, prior work has been limited in its
 142 exploration of the RLVR stage.

2.3 OUR DISTINCTION

144 As mentioned above, general CoT consists of multi-step natural language reasoning traces derived
 145 from instruction data. These traces teach the model structured reasoning patterns during SFT and help
 146 it select more appropriate responses during RLVR after learning task-solving behaviors. However,
 147 unlike open-domain tasks where many reasoning paths may be acceptable, medical reasoning is
 148 highly standardized and tightly constrained by clinical knowledge. This significantly limits the
 149 diversity of viable intermediate steps, reducing the exploration space available to RLVR and making
 150 it harder for the model to discover improved reasoning trajectories. To address this limitation, we
 151 introduce reflective-pattern-injected CoT data during SFT. This data provides the model with initial
 152 self-reflection and self-correction capabilities, effectively expanding the reasoning space that RLVR
 153 can explore and enabling more robust improvements on complex medical tasks. Moreover, existing
 154 studies either focus on a narrow set of benchmarks or fail to consistently improve the performance of
 155 the SFT model. In contrast, we not only successfully enhance model performance during the RLVR
 156 stage but also conduct extensive experiments to analyze the features in multimodal medical tasks.
 157

3 METHODOLOGY

158 In this section, we outline our methodology for advancing multimodal medical understanding and
 159 reasoning through RLVR with a self-reflective cold start, which is depicted in Figure 2. Our
 160 approach unfolds in two stages: (1) A cold start phase, in which we uniquely integrate general

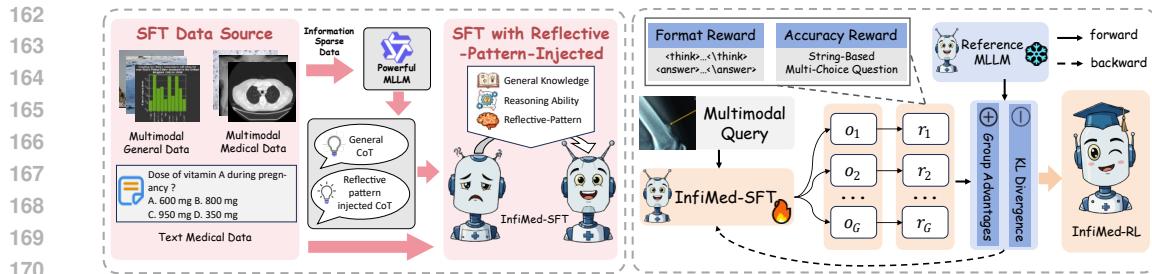


Figure 2: The overall training process of InfiMed-Series models.

multimodal data with medical text reasoning data to simultaneously enhance image understanding and restore fundamental reasoning skills. Crucially, to address the information sparsity of existing medical datasets, we further synthesize both distilled CoT and self-reflective CoT for SFT, thereby establishing a richer and more exploratory reasoning foundation. (2) A RLVR phase, which enables the model to explore a wider spectrum of reasoning trajectories, thereby producing more robust and clinically faithful multimodal reasoning.

3.1 REFLECTIVE-INJECTED SUPERVISED FINE-TUNING

As mentioned above, since SFT constitutes the foundation for subsequent RLVR, we incorporated not only general multimodal data but also text-based medical reasoning data during SFT to strengthen the model’s fundamental multimodal understanding and reasoning capabilities (Sellergren et al., 2025; Xu et al., 2025). However, several existing multimodal medical SFT datasets suffer from insufficient informational richness. For instance, multiple-choice question datasets often only provide the final choice and always lack explicit explanation. To address this limitation, in addition to only generating conventional CoT data to supplement the missing information, we further construct reflective-pattern-injected CoT data, enabling the model to develop more comprehensive and self-corrective reasoning capabilities (Cheng et al., 2024).

The core premise of reflective-pattern-injected is that directly exposing the model to a spectrum of reasoning trajectories, including correct, partially inconsistent, and subtly flawed chains, encourages the development of self-evaluation and error-correction mechanisms.

Formally, given a multimodal medical question q , whose original response consists of insufficient information, consisting of a textual task instruction x and one or more images \mathcal{I} , i.e. $q = \{x, \mathcal{I}\}$. We utilize a powerful MLLM (e.g., Qwen2.5-VL-32B (Bai et al., 2025)) to generate a batch of candidate responses $\{y_i\}$. Subsequently, leveraging rejection sampling, we partition these candidates into two disjoint subsets: $\{y_i^+\}$, corresponding to correct responses, and $\{y_i^-\}$, corresponding to incorrect responses. For the correct responses $\{y_i^+\}$, we further engage a more advanced MLLM (e.g., Qwen2.5-VL-72B (Bai et al., 2025)) to evaluate each response across multiple dimensions, including clinical accuracy, logical reasoning, factual correctness, and completeness.

Finally, we synthesize a reflective-pattern-injected CoT by combining one of the highest-quality responses from $\{y_i^+\}$ with a randomly selected response from $\{y_i^-\}$, thereby creating a novel training instance that emphasizes both reasoning and error-awareness. More details of the reflective-pattern-injected CoT synthesis can be found in the Appendix A.2.

3.2 REINFORCEMENT LEARNING WITH VERIFIABLE REWARDS

3.2.1 OVERALL PROCESS OF RLVR

After the SFT stage with reflective-pattern injection, we use Group Relative Policy Optimization (GRPO) in the RLVR phase to improve stability, building on the method from Deepseek-R1 (Guo et al., 2025). GRPO computes advantages by generating multiple responses for the same query, removing the need for an explicit critic model.

We formally denote the model after the SFT stage with reflective-pattern injection as π_θ , the policy model in RLVR. Given a multimodal medical query q , the policy model $\pi_{\theta_{old}}$ (prior to parameter updates) generates a set of G candidate responses $\{o_i\}_{i=1}^G$. For each response o_i , a rule-based

reward function $R(o, \text{gt})$ is used to evaluate its quality and assign a score r_i , where gt denotes the ground-truth answer. Based on the collection of rewards $\{r_i\}_{i=1}^G$, the group-relative advantages $\{A_i\}_{i=1}^G$, which quantify the relative quality of responses within the batch, can be calculated as:

$$A_i = \frac{r_i - \text{mean}(\{r_1, r_2, \dots, r_G\})}{\text{std}(\{r_1, r_2, \dots, r_G\})}, \quad (1)$$

where $\text{mean}(\cdot)$ indicates the average value, and $\text{std}(\cdot)$ refers to the standard deviation.

Based on the above group-relative advantages, GRPO updates the policy by maximizing the expected advantage-weighted likelihood ratio. The optimization objective can be formulated as:

$$\begin{aligned} \mathcal{J}_{\text{GRPO}}(\theta) = & \mathbb{E}_{[q \sim P(Q), \{o_i\}_{i=1}^G \sim \pi_{\theta_{\text{old}}}(O|q)]} \frac{1}{G} \sum_{i=1}^G \frac{1}{|o_i|} \\ & \sum_{t=1}^{|o_i|} \left\{ \min \left[\frac{\pi_{\theta}(o_i|q)}{\pi_{\theta_{\text{old}}}(o_i|q)} A_i, \text{clip} \left(\frac{\pi_{\theta}(o_i|q)}{\pi_{\theta_{\text{old}}}(o_i|q)}, 1 - \epsilon, 1 + \epsilon \right) A_i \right] - \beta D_{\text{KL}} [\pi_{\theta} \parallel \pi_{\text{ref}}] \right\}, \end{aligned} \quad (2)$$

where the additional Kullback–Leibler term $D_{\text{KL}} [\pi_{\theta} \parallel \pi_{\text{ref}}]$ is applied to penalize divergence from the reference policy model π_{ref} , thereby helping to maintain training stability.

3.2.2 RULE-BASED REWARD CONSTRUCTION

Considering the reward function $R(o, \text{gt})$ aims to guide the policy model to learn a suitable and correct reasoning trajectory, we design our total reward R_{total} , which integrates assessments of both output format correctness and accuracy:

$$R_{\text{total}}(o, \text{gt}) = w_{\text{format}} \cdot R_{\text{format}}(o) + w_{\text{acc}} \cdot R_{\text{accuracy}}(o, \text{gt}), \quad (3)$$

where $R_{\text{format}}(o)$ denotes the reward for the correctness of the output format and $R_{\text{accuracy}}(o, \text{gt})$ denotes the reward for the accuracy of the output o relative to the ground-truth result. The non-negative coefficients w_{format} and w_{acc} serve as hyperparameters weighting the relative contribution of the two components, with $w_{\text{format}} + w_{\text{acc}} = 1$.

The format reward $R_{\text{format}}(o)$ assesses whether the output of the policy model π_{θ} satisfies the predefined format. Notably, $R_{\text{format}}(o) \in \{0, 1\}$, where $R_{\text{format}}(o) = 1$ if all specified format requirements are satisfied; otherwise, $R_{\text{format}}(o) = 0$. Specifically, it verifies two primary aspects:

- **Thinking Progress:** We evaluate whether the model correctly presents its reasoning process according to a predefined format. Specifically, the model may be required to encapsulate its reasoning process and final answer within designated tags.
- **Final Answer Format:** We examine whether the model outputs an explicit final answer, with particular attention to cases where the instructions related to query q require such a response.

The accuracy reward $R_{\text{accuracy}}(o, \text{gt})$ evaluates the correctness of the model output o relative to the ground truth of query q . Importantly, $R_{\text{accuracy}}(o, \text{gt})$ is defined only when the output meets the format constraint, i.e., $R_{\text{format}}(o) = 1$; otherwise, it is zero. This design ensures that the model generates well-structured outputs before being evaluated for correctness. When $R_{\text{format}}(o) = 1$, the computation of $R_{\text{accuracy}}(o, \text{gt})$ depends on the task-specific ground-truth format. The two main tasks' reward functions are as follows; others are presented in the Appendix A.3.

- **String-based Tasks:** For textual answers, $R_{\text{accuracy}}(o, \text{gt})$ is computed by normalizing both the model output and the ground truth (e.g., lowercasing, removing redundant spaces). This function evaluates the extracted answer from the output o , denoted as o_{ans} , by comparing it to the ground truth answer gt . We use the Jaccard function to measure the similarity between o_{ans} and gt . The Jaccard function can be formulated as: $\text{Jaccard}(o_{\text{ans}}, \text{gt}) = \frac{|o_{\text{ans}} \cap \text{gt}|}{|o_{\text{ans}} \cup \text{gt}|}$.
- **Multiple-Choice Questions:** For tasks that require selecting an option from a predefined set, $R_{\text{accuracy}}(o, \text{gt})$ is calculated by directly comparing the model's extracted predicted answer, o_{ans} , with the correct ground truth option, gt . A match results in a reward of 1, while a mismatch yields a reward of 0.

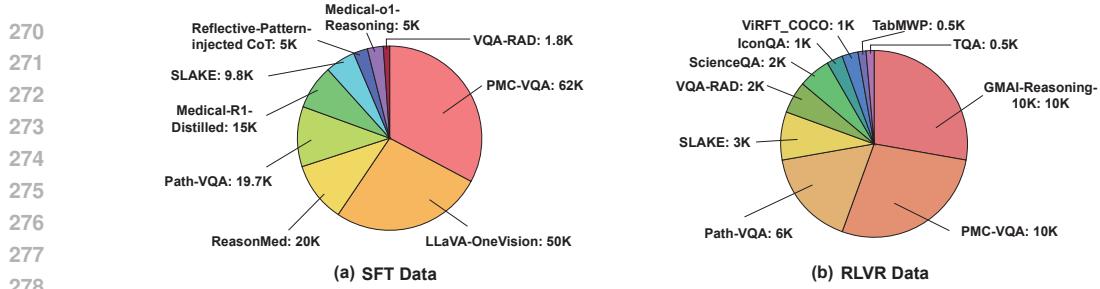


Figure 3: Overview of the training samples for the InfiMed series models in the reflective-injected SFT and RLVR stages.

4 EXPERIMENT

In this section, we present the experimental setup used to train and evaluate our proposed InfiMed-series models, which are built upon Qwen2.5-VL-3B-Instruct (Bai et al., 2025). We detail the implementation process, outline the evaluation benchmarks, and provide a comprehensive comparison with SOTA models. Furthermore, we analyse the impact of our training recipe to better understand its contributions to overall performance. We also aim to address the following research questions.

- **RQ1:** How do the InfiMed-Series models perform compared with other MLLMs across various medical benchmarks?
- **RQ2:** How do different data types and data numbers influence SFT and RLVR performance?
- **RQ3:** Does reasoning really enhance performance in medical tasks?
- **RQ4:** How do models trained with self-reflection via SFT compare to RLVR-optimized models in their quality and reliability of medical responses?
- **RQ5:** Can our SFT data consistently improve model performance across different base MLLM architectures?
- **RQ6:** Does increasing the amount of RLVR training data lead to further improvements in model performance?

4.1 EXPERIMENTAL SETUP

Models. We conduct a comprehensive comparison across a wide range of models. The models include: (1) Proprietary models: GPT-series models (Achiam et al., 2023), Claude Sonnet 4 (Anthropic, 2025), and Gemini-2.5-Flash (Comanici et al., 2025); (2) General open-source models: Qwen2.5-VL series models (Bai et al., 2025), Gemma3 series models (Team et al., 2025) and InternVL series models (Chen et al., 2024c; Zhu et al., 2025); (3) Medical open-source models: MedVLM-R1-2B (Pan et al., 2025), MedGemma-4B-IT (Sellergren et al., 2025), LLaVa-Med-7B (Li et al., 2023), HuatuoGPT-V-7B (Chen et al., 2024b), Lingshu-7B (Xu et al., 2025), BioMediX2-8B (Mullappilly et al., 2024).

Datasets. During the reflective-injected SFT stage, we utilize a total of **188K** samples from three categories: (1) multimodal general data, (2) multimodal medical data, and (3) text-based medical data. In the RLVR stage, we utilize **36K** multimodal medical datasets and multimodal general datasets. An overview of the training datasets is provided in Figure 3. The detailed description of the datasets is provided in the Appendix A.4.

Evaluation Benchmark We adopt seven widely used multimodal medical benchmarks: MMMU-Health&Medicine (MMMU-H&M) (Yue et al., 2024), VQA-RAD (Lau et al., 2018), SLAKE (Liu et al., 2021), PathVQA (He et al., 2020), PMC-VQA (Zhang et al., 2023c), OmniMedVQA (Hu et al., 2024), and MedXpertQA (Zuo et al., 2025). These benchmarks span a wide range of imaging modalities, including X-ray, CT, MRI, PET, ultrasound, and pathology. Collectively, they form a comprehensive evaluation framework for assessing both reasoning ability and proficiency in general medical knowledge. A detailed description of these benchmarks is provided in Appendix A.4.

Additional experimental setup details are provided in Appendix A.4.

Table 1: **Performance comparison of different MLLMs across various medical vision-language benchmarks.** Results of comparison of InfiMed-series models with other MLLMs on medical multi-modal benchmarks. MMMU-H&M, OMVQA, and MedXQA denote MMMU-Health&Medicine, OmniMedVQA, and MedXpertQA-Multimodal, respectively. The best results among models in the 2-4B parameter are **bolded**.

Model	Size	Accuracy (%)							
		MMMU-H&M	VQA-RAD	SLAKE	PathVQA	PMC-VQA	OMVQA	MedXQA	Avg.
<i>Proprietary Models</i>									
GPT-5	-	83.6	67.8	78.1	52.8	60.0	76.4	71.0	70.0
GPT-5-mini	-	80.5	66.3	76.1	52.4	57.6	70.9	60.1	66.3
GPT-5-nano	-	74.1	55.4	69.3	45.4	51.3	66.5	45.1	58.2
GPT-4.1	-	75.2	65.0	72.2	55.5	55.2	75.5	45.2	63.4
Claude Sonnet 4	-	74.6	67.6	70.6	54.2	54.4	65.5	43.3	61.5
Gemini-2.5-Flash	-	76.9	68.5	75.8	55.4	55.4	71.0	52.8	65.1
<i>General Open-source Models</i>									
Qwen2.5VL-3B	3B	51.3	56.8	63.2	37.1	50.6	64.5	20.7	49.2
Gemma3-4B	4B	34.0	49.9	61.1	43.2	47.9	60.9	20.9	45.4
Qwen2.5VL-7B	7B	54.0	65.0	67.6	44.6	51.3	63.5	21.7	52.5
InternVL2.5-8B	8B	53.5	59.4	69.0	42.1	51.3	81.3	21.7	54.0
InternVL3-8B	8B	59.2	65.4	72.8	48.6	53.8	79.1	22.4	57.3
<i>Medical Open-source Models</i>									
MedVLM-R1-2B	2B	35.2	48.6	56.0	32.5	47.6	77.7	20.4	45.4
MedGemma-4B-IT	4B	43.7	49.9	76.4	48.8	49.9	69.8	22.3	51.5
LLaVA-Med-7B	7B	29.3	53.7	48.0	38.8	30.5	44.3	20.3	37.8
HuatuGPT-V-7B	7B	47.3	67.0	67.8	48.0	53.3	74.2	21.6	54.2
Lingshu-7B	7B	54.0	67.9	83.1	61.9	56.3	82.9	26.7	61.8
BioMediX2-8B	8B	39.8	49.2	57.7	37.0	43.5	63.3	21.8	44.6
<i>Ours (InfiniMed-Series)</i>									
InfiniMed-SFT-3B	3B	54.7	58.1	82.0	60.6	53.2	67.0	23.5	57.1
Gemma3-SFT-4B	4B	35.3	59.9	83.3	64.7	53.3	68.7	21.0	55.2
InfiniMed-RL-3B	3B	55.3	60.5	82.4	62.0	58.7	71.7	23.6	59.2

4.2 RESULTS ON VARIOUS MEDICAL BENCHMARKS (RQ1 & RQ5)

Table 1 presents a comprehensive comparison of different MLLMs across seven diverse medical vision-language benchmarks. Among all models, proprietary closed-source models (e.g., GPT-5, Gemini-2.5, Claude) consistently outperform both general-purpose and medical-domain open-source models, achieving the highest average accuracy (e.g., 70.0% for GPT-5). These models set a strong upper bound, particularly excelling on complex benchmarks such as MMMU-H&M and MedXpertQA, indicating their superior reasoning and image understanding capabilities.

Furthermore, comparisons with existing open-source models show that the InfiMed-series models offer significant performance advantages. Both InfiMed-SFT-3B and InfiMed-RL-3B notably outperform other models of similar scale, achieving average accuracies of 57.1% and 59.2%, respectively, across seven multimodal medical benchmarks. We also notice that MedVLM-R1-2B achieves 77.7% on OmniMedVQA, primarily because its training dataset may overlap with a portion of the OmniMedVOA benchmark.

Notably, our 3B models outperform some larger 7B and 8B models, such as HuatuoGPT-V-7B and InternVL2.5-8B, despite their greater scale. Although a gap remains between InfiMed-RL-3B and Lingshu-7B, our model achieves competitive performance with fewer parameters and using a low-resource dataset (188K for SFT and 36K for RLVR), compared to Lingshu-7B's 12M samples, HuatuoGPT-V-7B's 1.3M samples, highlighting the efficiency and effectiveness of our training.

Although models such as MedVLM-R1-2B, GMAI-VL-R1-7B (not open-sourced), and Lingshu-7B provide some evidence that RLVR can be effective after SFT, they either target only a narrow range of benchmarks or fail to achieve consistent overall gains. By contrast, the 2.1% overall improvement of InfiMed-RL-3B over InfiMed-SFT-3B, along with consistent gains across seven medical benchmarks, clearly demonstrates that RLVR not only could enhance model performance in the medical domain but also complements our SFT phase training, thereby substantiating the effectiveness of RLVR.

To further evaluate the model-agnostic robustness of our data, we incorporated a fundamentally different model family, the Gemma3 series, and fine-tuned a Gemma3-4B-IT model using our SFT data. Despite the architectural and training differences between Gemma and previously evaluated models such as Owen, our data still produces substantial performance gains over the original Gemma

378
379
380
381
382
383
Table 2: Ablation study examining data composition during the training stage. $\Delta|\text{Data}|$ denotes
384 the amount of data change applied to the training set. w/o-general, w/o-text, and w/o-refcot refer
385 to training configurations where the general multimodal data, textual medical data, and reflective-
386 pattern-injected CoT data are removed, respectively. gen_mm, text, and general_cot denote the
387 general multimodal data, medical textual data, and general CoT data components included in the
388 training corpus.

384 385 386 387 388 389 390 391 392 393 394 395 396 397 398 399 400 401 402 403 404 405 406 407 408 409 410 411 412 413 414 415 416 417 418 419 420 421 422 423 424 425 426 427 428 429 430 431	$\Delta \text{Data} $	Accuracy (%)							
		384 385 386 387 388 389 390 391 392 393 394 395 396 397 398 399 400 401 402 403 404 405 406 407 408 409 410 411 412 413 414 415 416 417 418 419 420 421 422 423 424 425 426 427 428 429 430 431							
<i>Base Model</i>									
Qwen2.5VL-3B	-	51.3	56.8	63.2	37.1	50.6	64.5	20.7	49.2
<i>Ablation Study in SFT Stage on General Multimodal Data</i>									
InfiMed-SFT-3B	-	54.7	58.1	82.0	60.6	53.2	67.0	23.5	57.1
InfiMed-SFT-3B+gen_mm	+20K	54.0	60.7	81.8	55.8	55.5	67.5	22.3	56.8
InfiMed-SFT-3B+gen_mm	-20K	48	58.1	82.1	58.3	51.4	66.9	22.6	55.4
<i>Ablation Study in SFT Stage on Medical Text Data</i>									
InfiMed-SFT-3B-w/o-general	-50K	50.0	60.5	80.5	60.4	51.6	59.7	22.6	55.1
InfiMed-SFT-3B+text	+20K	50.7	63.4	80.4	57.3	54.4	67.3	21.4	56.4
InfiMed-SFT-3B+text	-20K	50.7	60.3	82.2	58.1	53.8	67.3	22.6	56.4
InfiMed-SFT-3B-w/o-text	-40K	44.0	61.0	81.6	60.4	51.1	64.7	21.9	54.9
<i>Ablation Study in SFT Stage on Reflective CoT Data</i>									
InfiMed-SFT-3B-w/generalcot	-	50.7	60.5	81.9	57.7	52.7	66.4	23.4	56.2
InfiMed-SFT-3B-w/o-refcot	-5K	50.0	60.1	81.3	60.5	53.1	64.7	22.8	56.1
<i>Ablation Study in RLVR Stage</i>									
InfiMed-RL-3B	-	55.3	60.5	82.4	62.0	58.7	71.7	23.6	59.2
InfiMed-RL-3B-w/o-general	-10K	53.3	60.7	81.9	61.6	58.3	70.0	23.6	58.4

baseline and achieves state-of-the-art performance within the 3–4B parameter range (excluding comparisons with InfiMed itself). Although this experiment relies solely on SFT, the strong improvements obtained with a relatively small dataset of 176K samples demonstrate that the effectiveness of our SFT data is not tied to any specific backbone family. These results provide compelling evidence that our data is highly efficient and generalizes well across heterogeneous MLLM architectures.

4.3 ABLATION STUDY ON DATA COMPOSITION (RQ2 & RQ6)

For the SFT and RLVR stage, we assessed the contribution of each component in our dataset by conducting an ablation study. For the SFT, we systematically remove specific types of training data, including general multimodal data, textual medical data, and reflective-pattern-injected CoT data. The overall detailed results are presented in Table 2. Based on these experiments, we draw the following conclusion: ***Unlike general multimodal tasks, medical multimodal problems are inherently comprehensive, requiring the integration of textual, visual, and domain-specific knowledge. As a result, medical training datasets alone are insufficient to ensure robust MLLM performance.*** A detailed analysis is provided below:

During the SFT stage, we observe that each data component serves a distinct function. Removing general multimodal data has a pronounced negative effect on benchmarks like OmniMedVQA, which require nuanced visual understanding. This suggests that general-domain multimodal examples help the model interpret complex visual patterns, align visual and textual features, and handle diverse image information in medical-specific datasets. Excluding textual medical data severely degrades performance on MMMU-H&M, indicating that such data provides critical domain-specific knowledge, including medical terminology, clinical reasoning strategies, and structured question-answering patterns essential for accurate interpretation and reasoning.

In addition to the ablations on each data component, we performed supplementary experiments to examine the effect of moderate changes in data proportions. We increased and decreased the amounts of general multimodal data and textual medical data by 20K samples. As shown in Table 2, these adjustments lead to performance differences across benchmarks, reflecting the complementary roles of the two data types. Our current data composition was determined based on empirical observations from preliminary experiments, which indicated that this setting offers a stable balance between visual understanding and medical-domain reasoning. Although the results suggest that alternative ratios may provide further gains, an exhaustive search for optimal proportions requires a broader investigation.

Moreover, even though reflective-pattern-injected CoT data contains only 5K examples, its removal leads to noticeable declines across most benchmarks, highlighting its role in enhancing multi-step reasoning and self-reflection for complex or ambiguous medical questions. Interestingly, VQA-RAD performance slightly increases after removing certain data, as this older benchmark emphasizes mem-

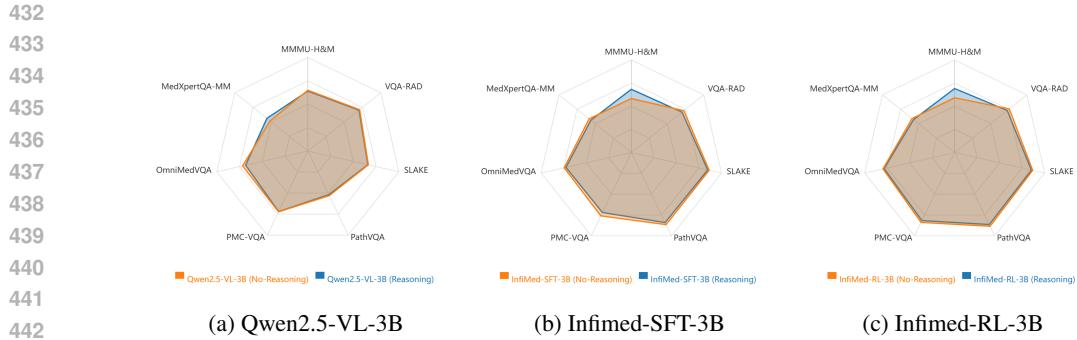


Figure 4: Comparison of direct-answer and reasoning-based prompts on medical benchmarks.

orization; reducing other data effectively increases the relative proportion of VQA-RAD examples, yielding a modest gain.

To further isolate the contribution of reflective CoT, we conducted an additional ablation in which we replaced the reflective CoT data with an equal amount of general CoT samples. The model equipped with reflective CoT consistently outperformed the one trained with general CoT, demonstrating that the reflective formulation itself provides unique benefits by guiding the model to articulate intermediate reasoning, identify potential errors, and refine its final predictions with greater reliability.

The lower half of Table 2 reports the RLVR ablation, focusing on variants excluding general multimodal data. When removed, performance on MMMU-H&M drops below InfiMed-SFT-3B, suggesting that RLVR relying solely on medical multimodal data, which is typically less reasoning-intensive, reduces overall reasoning capability, leading to lower performance. More ablation studies are presented in the Appendix A.5.

4.4 ANALYSIS OF REASONING EFFECTIVENESS IN MEDICAL SCENARIOS (RQ3)

To assess reasoning effectiveness in medical scenarios, we evaluated the model using two prompts: (1) a direct-answer prompt, where the model is asked to output only the final prediction; (2) a reasoning-augmented prompt, where the model is encouraged to generate intermediate reasoning steps before providing the answer. This setup allows us to examine the impact of explicit reasoning. Results are shown in Figure 4.

Our experiments reveal a consistent trend: explicit reasoning prompts tend to reduce performance on most medical benchmarks, with two exceptions. First, both InfiMed-SFT and InfiMed-RL benefit from reasoning on MMMU-H&M. Second, the general-purpose Qwen2.5-VL-3B benefits on both MMMU-H&M and MedXpertQA-MM. This indicates that explicit reasoning alone does not universally improve medical-focused MLLMs, even with additional optimization like RLVR. We attribute the gains on MMMU-H&M and MedXpertQA-MM (for Qwen2.5-VL-3B) to their reasoning-intensive design requiring multi-step logical deduction and cross-modal integration. For Qwen2.5-VL-3B, reasoning prompts structure latent knowledge, reduce uncertainty, and guide coherent intermediate steps. In contrast, InfiMed models have learned efficient, domain-specific strategies for direct medical answering, and enforcing explicit reasoning can disrupt these pathways, lowering performance on MedXpertQA-MM. Meanwhile, many other benchmarks are knowledge-driven, where answers can often be derived directly from visual information or domain expertise. For such tasks, step-by-step reasoning introduces redundant steps, increases hallucination risk, and interferes with the streamlined strategies of medically optimized models.

To further examine whether model size influences this phenomenon, we conducted additional experiments using larger model scales as well as different model families, including the LINGSHU (Xu et al., 2025) and HUATUO (Chen et al., 2024b) series. As shown in Figure 7 in the Appendix, even at the 7B and 32B scales, we continue to observe instances where explicit reasoning leads to performance degradation on several benchmarks. These results align with our primary conclusion that explicit reasoning is not universally advantageous for medical multimodal tasks. The degradation observed in larger models further suggests that introducing explicit reasoning can interfere with model performance on certain tasks.

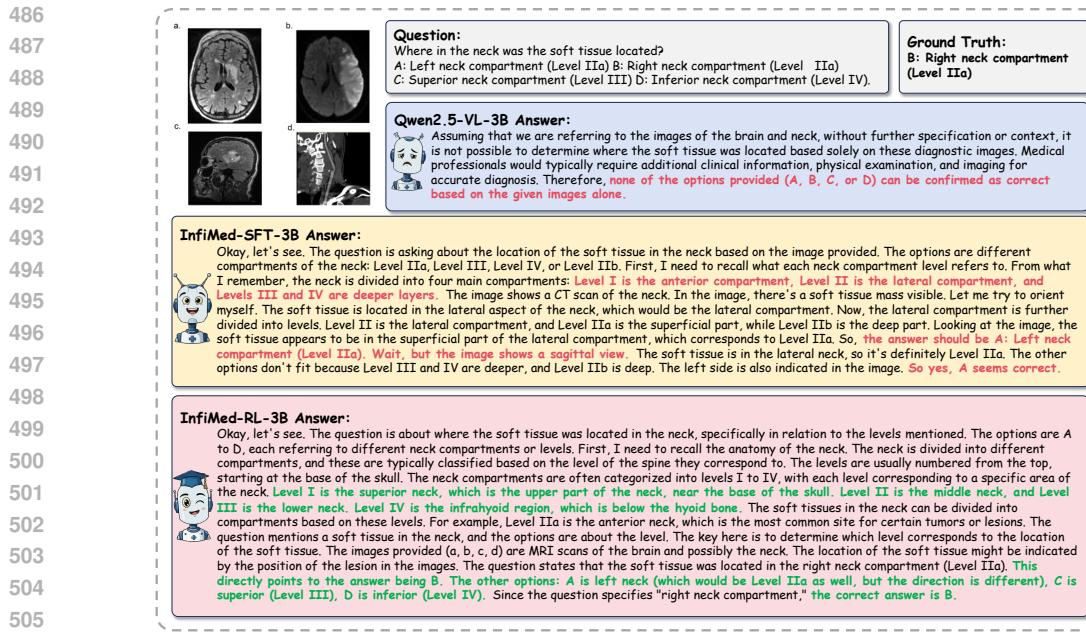


Figure 5: Case study on Qwen2.5-VL-3B, InfiMed-Series models in medical VQA. Red denotes errors or irrelevant content, whereas Green denotes correct or important information.

In summary, these results suggest that ***the effectiveness of explicit reasoning depends on both task type and model***. It may benefit reasoning-intensive tasks and general-purpose models on medical tasks requiring light reasoning, but can hinder performance on recognition-oriented benchmarks where direct factual or visual knowledge suffices.

4.5 CASE STUDY (RQ4)

Beyond benchmark evaluations, we conduct a case study to examine the qualitative differences between the Qwen2.5-VL-3B-Instructw, reflective-pattern-injected InfiMed-SFT-3B, and the RLVR-optimized InfiMed-RL-3B, as illustrated in Figure 5. Our case study highlights distinct response behaviours across models. The general-purpose Qwen2.5-VL-3B adopts a conservative strategy, emphasizing its lack of sufficient medical knowledge and ultimately failing to produce a definitive answer. InfiMed-SFT-3B, by contrast, can generate a reasoning chain and reproduce a reflective pattern. However, despite this reflection, it still converges on the same incorrect answer. This suggests that SFT primarily teaches the model to mimic the form of reflection, yet falls short of enabling genuine understanding or effective application of reflective reasoning. InfiMed-RL-3B, on the other hand, demonstrates a more structured reasoning process. In addition to identifying the correct option, it actively explores and evaluates the other options, reflecting the impact of RLVR in pushing the model beyond memorized patterns toward deliberate and systematic reasoning. More case studies are presented in the Appendix A.6.

5 CONCLUSION

In this work, we introduce the InfiMed-Series models, including InfiMed-SFT-3B and InfiMed-RL-3B, a set of multimodal large language models (MLLMs) specialized for medical tasks. To address the scarcity and sparsity of multimodal medical data, we augmented the training sets with general multimodal and textual medical data and synthesized reflective-pattern-injected chain-of-thought data, enabling the models to acquire initial exploratory capabilities and providing a structured foundation for subsequent Reinforcement Learning with Verifiable Rewards (RLVR) training. Experimental results across diverse medical benchmarks, covering both reasoning-intensive and understanding-oriented tasks, show that the InfiMed-Series models achieve state-of-the-art accuracy among models with similar parameter counts and even surpass some larger models. Beyond performance gains, our analysis provides new insights into the behavior and potential of MLLMs in medical scenarios.

540
541
ETHICS STATEMENT542
543
544
545
This work adheres to the ICLR Code of Ethics. Our study is purely empirical in nature, focused on
advancing the field of medical multimodal large language models. We have exclusively used standard
and publicly available medical datasets and open-source models, which were accessed and applied in
strict accordance with their licenses.546
547
548
549
550
551
552
We acknowledge the significant potential broader impacts and risks associated with the use of MLLMs
in healthcare. This includes concerns related to patient safety, clinical accuracy, and the potential for
misuse. Our work recognizes these challenges and aims to develop more powerful medical MLLMs
that can contribute to society and human well-being. Importantly, our research does not introduce
any new, unproven clinical applications. Instead, we focus on the ethical implications and underlying
principles of utilizing MLLMs in healthcare settings, and we are committed to ensuring that future
developments in this area are guided by a careful consideration of these ethical concerns.553
554
REPRODUCIBILITY STATEMENT
555556
557
558
To ensure the reproducibility of the results in this work, all models and synthetic datasets will be
made publicly available. We also describe the data sources and detailed experimental setup in Section
4.1 and Appendix A.4.560
561
REFERENCES562
563
Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman,
Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. Gpt-4 technical report.
[arXiv preprint arXiv:2303.08774](https://arxiv.org/abs/2303.08774), 2023.565
566
Rawan AlSaad, Alaa Abd-Alrazaq, Sabri Boughorbel, Arfan Ahmed, Max-Antoine Renault, Rafat
Damseh, and Javaid Sheikh. Multimodal large language models in health care: applications,
challenges, and future outlook. *Journal of medical Internet research*, 26:e59505, 2024.569
570
Anthropic. Introducing Claude 4. <https://www.anthropic.com/news/claude-4>, may
2025.571
572
Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibo Song, Kai Dang, Peng Wang,
Shijie Wang, Jun Tang, et al. Qwen2. 5-v1 technical report. [arXiv preprint arXiv:2502.13923](https://arxiv.org/abs/2502.13923),
2025.574
575
Junying Chen, Zhenyang Cai, Ke Ji, Xidong Wang, Wanlong Liu, Rongsheng Wang, Jianye Hou,
and Benyou Wang. Huatuogpt-01, towards medical complex reasoning with llms, 2024a. URL
<https://arxiv.org/abs/2412.18925>.578
579
Junying Chen, Ruyi Ouyang, Anningzhe Gao, Shunian Chen, Guiming Hardy Chen, Xidong Wang,
Ruifei Zhang, Zhenyang Cai, Ke Ji, Guangjun Yu, Xiang Wan, and Benyou Wang. Huatuogpt-
vision, towards injecting medical visual knowledge into multimodal llms at scale, 2024b. URL
<https://arxiv.org/abs/2406.19280>.582
583
Wenhu Chen, Xueguang Ma, Xinyi Wang, and William W Cohen. Program of thoughts prompt-
ing: Disentangling computation from reasoning for numerical reasoning tasks. [arXiv preprint
arXiv:2211.12588](https://arxiv.org/abs/2211.12588), 2022.585
586
Zhe Chen, Jiannan Wu, Wenhui 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*, pp. 24185–24198, 2024c.589
590
Kanzhi Cheng, Yantao Li, Fangzhi Xu, Jianbing Zhang, Hao Zhou, and Yang Liu. Vision-language
models can self-improve reasoning via reflection. [arXiv preprint arXiv:2411.00855](https://arxiv.org/abs/2411.00855), 2024.592
593
Tianzhe Chu, Yuexiang Zhai, Jihan Yang, Shengbang Tong, Saining Xie, Dale Schuurmans, Quoc V
Le, Sergey Levine, and Yi Ma. Sft memorizes, rl generalizes: A comparative study of foundation
model post-training. [arXiv preprint arXiv:2501.17161](https://arxiv.org/abs/2501.17161), 2025.

594 Gheorghe Comanici, Eric Bieber, Mike Schaeckermann, Ice Pasupat, Noveen Sachdeva, Inderjit
595 Dhillon, Marcel Blstein, Ori Ram, Dan Zhang, Evan Rosen, et al. Gemini 2.5: Pushing the frontier
596 with advanced reasoning, multimodality, long context, and next generation agentic capabilities.
597 arXiv preprint arXiv:2507.06261, 2025.

598 Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu,
599 Shirong Ma, Peiyi Wang, Xiao Bi, et al. Deepseek-r1: Incentivizing reasoning capability in llms
600 via reinforcement learning. arXiv preprint arXiv:2501.12948, 2025.

602 Xuehai He, Yichen Zhang, Luntian Mou, Eric Xing, and Pengtao Xie. Pathvqa: 30000+ questions for
603 medical visual question answering. arXiv preprint arXiv:2003.10286, 2020.

604 Yutao Hu, Tianbin Li, Quanfeng Lu, Wenqi Shao, Junjun He, Yu Qiao, and Ping Luo. Omnimedvqa:
605 A new large-scale comprehensive evaluation benchmark for medical lvm. In Proceedings of the
606 IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 22170–22183, 2024.

608 Wenxuan Huang, Bohan Jia, Zijie Zhai, Shaosheng Cao, Zheyu Ye, Fei Zhao, Zhe Xu, Yao Hu, and
609 Shaohui Lin. Vision-r1: Incentivizing reasoning capability in multimodal large language models.
610 arXiv preprint arXiv:2503.06749, 2025.

611 Jason J Lau, Soumya Gayen, Asma Ben Abacha, and Dina Demner-Fushman. A dataset of clinically
612 generated visual questions and answers about radiology images. *Scientific data*, 5(1):1–10, 2018.

614 Bo Li, Yuanhan Zhang, Dong Guo, Renrui Zhang, Feng Li, Hao Zhang, Kaichen Zhang, Peiyuan
615 Zhang, Yanwei Li, Ziwei Liu, et al. Llava-onevision: Easy visual task transfer. arXiv preprint
616 arXiv:2408.03326, 2024a.

617 Chunyuan Li, Cliff Wong, Sheng Zhang, Naoto Usuyama, Haotian Liu, Jianwei Yang, Tristan
618 Naumann, Hoifung Poon, and Jianfeng Gao. Llava-med: Training a large language-and-vision
619 assistant for biomedicine in one day. arXiv preprint arXiv:2306.00890, 2023.

620 Wenxuan Li, Chongyu Qu, Xiaoxi Chen, Pedro RAS Bassi, Yijia Shi, Yuxiang Lai, Qian Yu, Huimin
621 Xue, Yixiong Chen, Xiaorui Lin, et al. Abdomenatlas: A large-scale, detailed-annotated, &
622 multi-center dataset for efficient transfer learning and open algorithmic benchmarking. *Medical*
623 *Image Analysis*, 97:103285, 2024b.

625 Bo Liu, Li-Ming Zhan, Li Xu, Lin Ma, Yan Yang, and Xiao-Ming Wu. Slake: A semantically-
626 labeled knowledge-enhanced dataset for medical visual question answering. In *2021 IEEE 18th*
627 *international symposium on biomedical imaging (ISBI)*, pp. 1650–1654. IEEE, 2021.

628 Yuhang Liu, Pengxiang Li, Congkai Xie, Xavier Hu, Xiaotian Han, Shengyu Zhang, Hongxia Yang,
629 and Fei Wu. Infigui-r1: Advancing multimodal gui agents from reactive actors to deliberative
630 reasoners. arXiv preprint arXiv:2504.14239, 2025a.

631 Zeyu Liu, Yuhang Liu, Guanghao Zhu, Congkai Xie, Zhen Li, Jianbo Yuan, Xinyao Wang, Qing Li,
632 Shing-Chi Cheung, Shengyu Zhang, et al. Infi-mmr: Curriculum-based unlocking multimodal
633 reasoning via phased reinforcement learning in multimodal small language models. arXiv preprint
634 arXiv:2505.23091, 2025b.

636 Zihan Liu, Zhuolin Yang, Yang Chen, Chankyu Lee, Mohammad Shoeybi, Bryan Catanzaro, and Wei
637 Ping. Acereason-nemotron 1.1: Advancing math and code reasoning through sft and rl synergy.
638 arXiv preprint arXiv:2506.13284, 2025c.

639 Ziyu Liu, Zeyi Sun, Yuhang Zang, Xiaoyi Dong, Yuhang Cao, Haodong Duan, Dahua Lin, and Jiaqi
640 Wang. Visual-rft: Visual reinforcement fine-tuning. arXiv preprint arXiv:2503.01785, 2025d.

641 Ziyu Liu, Yuhang Zang, Yushan Zou, Zijian Liang, Xiaoyi Dong, Yuhang Cao, Haodong Duan,
642 Dahua Lin, and Jiaqi Wang. Visual agentic reinforcement fine-tuning, 2025e. URL <https://arxiv.org/abs/2505.14246>.

645 Pan Lu, Liang Qiu, Jiaqi Chen, Tony Xia, Yizhou Zhao, Wei Zhang, Zhou Yu, Xiaodan Liang, and
646 Song-Chun Zhu. Iconqa: A new benchmark for abstract diagram understanding and visual language
647 reasoning. In *The 35th Conference on Neural Information Processing Systems (NeurIPS) Track*
on Datasets and Benchmarks, 2021.

648 Pan Lu, Swaroop Mishra, Tony Xia, Liang Qiu, Kai-Wei Chang, Song-Chun Zhu, Oyvind Tafjord,
 649 Peter Clark, and Ashwin Kalyan. Learn to explain: Multimodal reasoning via thought chains for
 650 science question answering. In The 36th Conference on Neural Information Processing Systems
 651 (NeurIPS), 2022a.

652 Pan Lu, Liang Qiu, Kai-Wei Chang, Ying Nian Wu, Song-Chun Zhu, Tanmay Rajpurohit, Peter
 653 Clark, and Ashwin Kalyan. Dynamic prompt learning via policy gradient for semi-structured
 654 mathematical reasoning. arXiv preprint arXiv:2209.14610, 2022b.

655 Run Luo, Lu Wang, Wanwei He, and Xiaobo Xia. Gui-r1: A generalist r1-style vision-language
 656 action model for gui agents. arXiv preprint arXiv:2504.10458, 2025.

657 Fanqing Meng, Lingxiao Du, Zongkai Liu, Zhixiang Zhou, Quanfeng Lu, Daocheng Fu, Tiancheng
 658 Han, Botian Shi, Wenhui Wang, Junjun He, et al. Mm-eureka: Exploring the frontiers of multimodal
 659 reasoning with rule-based reinforcement learning. arXiv preprint arXiv:2503.07365, 2025.

660 Sahal Shaji Mullappilly, Mohammed Irfan Kurpath, Sara Pieri, Saeed Yahya Alsejari, Shanavas
 661 Cholakkal, Khaled Aldahmani, Fahad Khan, Rao Anwer, Salman Khan, Timothy Baldwin, and
 662 Hisham Cholakkal. Bimedix2: Bio-medical expert lmm for diverse medical modalities, 2024.
 663 URL <https://arxiv.org/abs/2412.07769>.

664 Jiazhen Pan, Che Liu, Junde Wu, Fenglin Liu, Jiayuan Zhu, Hongwei Bran Li, Chen Chen, Cheng
 665 Ouyang, and Daniel Rueckert. Medvilm-r1: Incentivizing medical reasoning capability of vision-
 666 language models (vlms) via reinforcement learning. arXiv preprint arXiv:2502.19634, 2025.

667 Shuai Peng, Di Fu, Liangcai Gao, Xiuqin Zhong, Hongguang Fu, and Zhi Tang. Multimath: Bridging
 668 visual and mathematical reasoning for large language models. arXiv preprint arXiv:2409.00147,
 669 2024.

670 Yingzhe Peng, Gongrui Zhang, Miaosen Zhang, Zhiyuan You, Jie Liu, Qipeng Zhu, Kai Yang,
 671 Xingzhong Xu, Xin Geng, and Xu Yang. Lmm-r1: Empowering 3b lmm with strong reasoning
 672 abilities through two-stage rule-based rl. arXiv preprint arXiv:2503.07536, 2025.

673 Yujia Qin, Yining Ye, Junjie Fang, Haoming Wang, Shihao Liang, Shizuo Tian, Junda Zhang, Jiahao
 674 Li, Yunxin Li, Shijue Huang, et al. Uitars: Pioneering automated gui interaction with native
 675 agents. arXiv preprint arXiv:2501.12326, 2025.

676 Andrew Sellergren, Sahar Kazemzadeh, Tiam Jaroensri, Atilla Kiraly, Madeleine Traverse, Timo
 677 Kohlberger, Shawn Xu, Fayaz Jamil, Hughes, Charles Lau, et al. Medgemma technical report.
 678 arXiv preprint arXiv:2507.05201, 2025.

679 Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang,
 680 Mingchuan Zhang, YK Li, Yang Wu, et al. Deepseekmath: Pushing the limits of mathemat-
 681 ical reasoning in open language models. arXiv preprint arXiv:2402.03300, 2024.

682 Guangming Sheng, Chi Zhang, Zilingfeng Ye, Xibin Wu, Wang Zhang, Ru Zhang, Yanghua Peng,
 683 Haibin Lin, and Chuan Wu. Hybridflow: A flexible and efficient rlhf framework. arXiv preprint
 684 arXiv: 2409.19256, 2024.

685 Karan Singhal, Shekoofeh Azizi, Tao Tu, S Sara Mahdavi, Jason Wei, Hyung Won Chung, Nathan
 686 Scales, Ajay Tanwani, Heather Cole-Lewis, Stephen Pfahl, et al. Large language models encode
 687 clinical knowledge. Nature, 620(7972):172–180, 2023.

688 Karan Singhal, Tao Tu, Juraj Gottweis, Rory Sayres, Ellery Wulczyn, Mohamed Amin, Le Hou,
 689 Kevin Clark, Stephen R Pfahl, Heather Cole-Lewis, et al. Toward expert-level medical question
 690 answering with large language models. Nature Medicine, 31(3):943–950, 2025.

691 Yanzhou Su, Tianbin Li, Jiyao Liu, Chenglong Ma, Junzhi Ning, Cheng Tang, Sibo Ju, Jin Ye,
 692 Pengcheng Chen, Ming Hu, et al. Gmai-vl-r1: Harnessing reinforcement learning for multimodal
 693 medical reasoning. arXiv preprint arXiv:2504.01886, 2025.

694 Yu Sun, Xingyu Qian, Weiwen Xu, Hao Zhang, Chenghao Xiao, Long Li, Yu Rong, Wenbing Huang,
 695 Qifeng Bai, and Tingyang Xu. Reasonmed: A 370k multi-agent generated dataset for advancing
 696 medical reasoning, 2025. URL <https://arxiv.org/abs/2506.09513>.

702 Huajie Tan, Yuheng Ji, Xiaoshuai Hao, Minglan Lin, Pengwei Wang, Zhongyuan Wang, and
 703 Shanghang Zhang. Reason-rft: Reinforcement fine-tuning for visual reasoning. [arXiv preprint](https://arxiv.org/abs/2503.20752)
 704 [arXiv:2503.20752](https://arxiv.org/abs/2503.20752), 2025.

705 Gemma Team, Aishwarya Kamath, Johan Ferret, Shreya Pathak, Nino Vieillard, Ramona Merhej,
 706 Sarah Perrin, Tatiana Matejovicova, Alexandre Ramé, Morgane Rivière, et al. Gemma 3 technical
 707 report. [arXiv preprint arXiv:2503.19786](https://arxiv.org/abs/2503.19786), 2025.

708 Haozhe Wang, Chao Qu, Zuming Huang, Wei Chu, Fangzhen Lin, and Wenhui Chen. VI-rethinker:
 709 Incentivizing self-reflection of vision-language models with reinforcement learning. [arXiv preprint](https://arxiv.org/abs/2504.08837)
 710 [arXiv:2504.08837](https://arxiv.org/abs/2504.08837), 2025.

711 Xidong Wang, Nuo Chen, Junyin Chen, Yidong Wang, Guorui Chen, Chunxian Zhang, Xiangbo
 712 Wu, Yan Hu, Anningzhe Gao, Xiang Wan, et al. Apollo: A lightweight multilingual medical llm
 713 towards democratizing medical ai to 6b people. [arXiv preprint arXiv:2403.03640](https://arxiv.org/abs/2403.03640), 2024.

714 Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny
 715 Zhou, et al. Chain-of-thought prompting elicits reasoning in large language models. [Advances in](https://advances.in)
 716 [neural information processing systems](https://advances.in), 35:24824–24837, 2022.

717 Weiwen Xu, Hou Pong Chan, Long Li, Mahani Aljunied, Ruifeng Yuan, Jianyu Wang, Chenghao
 718 Xiao, Guizhen Chen, Chaoqun Liu, Zhaodonghui Li, et al. Lingshu: A generalist foundation model
 719 for unified multimodal medical understanding and reasoning. [arXiv preprint arXiv:2506.07044](https://arxiv.org/abs/2506.07044),
 720 2025.

721 Heng Yin, Yuqiang Ren, Ke Yan, Shouhong Ding, and Yongtao Hao. Rod-mllm: Towards more
 722 reliable object detection in multimodal large language models. In [Proceedings of the Computer](https://proceedings.cvcv.org)
 723 [Vision and Pattern Recognition Conference](https://proceedings.cvcv.org), pp. 14358–14368, 2025.

724 Xiang Yue, Yuansheng Ni, Kai Zhang, Tianyu Zheng, Ruoqi Liu, Ge Zhang, Samuel Stevens, Dongfu
 725 Jiang, Weiming Ren, Yuxuan Sun, et al. Mmmu: A massive multi-discipline multimodal under-
 726 standing and reasoning benchmark for expert agi. In [Proceedings of the IEEE/CVF Conference](https://proceedings.cvcv.org)
 727 [on Computer Vision and Pattern Recognition](https://proceedings.cvcv.org), pp. 9556–9567, 2024.

728 Hongbo Zhang, Junying Chen, Feng Jiang, Fei Yu, Zhihong Chen, Guiming Chen, Jianquan Li,
 729 Xiangbo Wu, Zhang Zhiyi, Qingying Xiao, et al. Huatuogpt, towards taming language model
 730 to be a doctor. In [Findings of the Association for Computational Linguistics: EMNLP 2023](https://findings.aclweb.org), pp.
 731 10859–10885, 2023a.

732 Jingyi Zhang, Jiaxing Huang, Huanjin Yao, Shunyu Liu, Xikun Zhang, Shijian Lu, and Dacheng Tao.
 733 R1-vl: Learning to reason with multimodal large language models via step-wise group relative
 734 policy optimization. [arXiv preprint arXiv:2503.12937](https://arxiv.org/abs/2503.12937), 2025.

735 Sheng Zhang, Yanbo Xu, Naoto Usuyama, Hanwen Xu, Jaspreet Bagga, Robert Tinn, Sam Preston,
 736 Rajesh Rao, Mu Wei, Naveen Valluri, et al. Biomedclip: A multimodal biomedical foundation
 737 model pretrained from fifteen million scientific image-text pairs. [arxiv 2023](https://arxiv.org/abs/2303.00915). [arXiv preprint](https://arxiv.org/abs/2303.00915)
 738 [arXiv:2303.00915](https://arxiv.org/abs/2303.00915), 2023b.

739 Xiaoman Zhang, Chaoyi Wu, Ziheng Zhao, Weixiong Lin, Ya Zhang, Yanfeng Wang, and Weidi
 740 Xie. Pmc-vqa: Visual instruction tuning for medical visual question answering. [arXiv preprint](https://arxiv.org/abs/2305.10415)
 741 [arXiv:2305.10415](https://arxiv.org/abs/2305.10415), 2023c.

742 Yaowei Zheng, Richong Zhang, Junhao Zhang, Yanhan Ye, Zheyuan Luo, Zhangchi Feng, and
 743 Yongqiang Ma. Llamafactory: Unified efficient fine-tuning of 100+ language models. In
 744 [Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics](https://proceedings.aclweb.org)
 745 [Volume 3: System Demonstrations](https://proceedings.aclweb.org), Bangkok, Thailand, 2024. Association for Computational
 746 Linguistics. URL [http://arxiv.org/abs/2403.13372](https://arxiv.org/abs/2403.13372).

747 Zijie Zhou. The table qa dataset distilled from deepseek-r1. [https://huggingface.co/](https://huggingface.co/datasets/jared-zhou/TQA-Distill-R1)
 748 [datasets/jared-zhou/TQA-Distill-R1](https://huggingface.co/datasets/jared-zhou/TQA-Distill-R1), 2025.

749 Jinguo Zhu, Weiyun Wang, Zhe Chen, Zhaoyang Liu, Shenglong Ye, Lixin Gu, Hao Tian, Yuchen
 750 Duan, Weijie Su, Jie Shao, et al. Internvl3: Exploring advanced training and test-time recipes for
 751 open-source multimodal models. [arXiv preprint arXiv:2504.10479](https://arxiv.org/abs/2504.10479), 2025.

756 Wenwen Zhuang, Xin Huang, Xiantao Zhang, and Jin Zeng. Math-puma: Progressive upward
757 multimodal alignment to enhance mathematical reasoning. In Proceedings of the AAAI Conference
758 on Artificial Intelligence, volume 39, pp. 26183–26191, 2025.

759
760 Yuxin Zuo, Shang Qu, Yifei Li, Zhang-Ren Chen, Xuekai Zhu, Ermo Hua, Kaiyan Zhang, Ning Ding,
761 and Bowen Zhou. Medxpertqa: Benchmarking expert-level medical reasoning and understanding.
762 In Forty-second International Conference on Machine Learning, 2025.

763

764

765

766

767

768

769

770

771

772

773

774

775

776

777

778

779

780

781

782

783

784

785

786

787

788

789

790

791

792

793

794

795

796

797

798

799

800

801

802

803

804

805

806

807

808

809

810
811

A APPENDIX

812
813

A.1 THE USE OF LARGE LANGUAGE MODELS

814
815
816
817
We used a Large Language Model (LLM) **exclusively for editing purposes**, focusing on correcting
grammar and typing errors. It's crucial to clarify that **the LLMs were not involved in the core
aspects of the research**, including the development or revision of key ideas and experimental design.818
819

A.2 CONSTRUCTION OF REFLECTIVE-PATTERN-INJECTED CoT

820
In this section, we present the detailed construction process of the reflective-pattern-injected CoT.821
822
823
824
825
826
For multimodal datasets with sparse information (e.g., multiple-choice questions), each query is
defined as $q = \{x, \mathcal{I}\}$, where x denotes the textual task instruction and \mathcal{I} represents one or more
images. We first employ Qwen2.5-VL-32B (Bai et al., 2025) to generate 10 candidate responses
 $\{y_i\}_{i=1}^{10}$ for each query q . Through rejection sampling, we divide these into two subsets: $\{y_i^+\}_{i=1}^m$
and $\{y_i^-\}_{i=1}^n$, where $m + n = 10$.827
For each response in $\{y_i^+\}_{i=1}^m$, we apply the following prompt to generate a score:828
829

Prompt for CoT Quality Evaluation

830
831
832
833
834
835
836
837
838
839
840
841
842
843
844
845
846
847
848
849
850
851
852
853
854
855
856
857
858
859
860
861
862
863
You are a medical reasoning evaluator. Assess the following response based on these criteria:

- 1. Clinical accuracy:** Correct incorporation of medical facts, clinical guidelines, and evidence-based practices. Accuracy, relevance, and appropriateness of clinical details.
- 2. Logical reasoning:** Coherent reasoning process, logically leading to the answer, well-supported by clinical knowledge.
- 3. Factual correctness:** All statements are factually correct and consistent with established medical knowledge.
- 4. Completeness:** Thorough coverage of all necessary aspects without missing critical information.

Question: $\{q\}$

Response: $\{y_i^+\}$

Please evaluate the response on the above criteria and ONLY provide the `Dict` object with
two keys:

`{'score': integer between 1 and 10, 'justification': concise explanation of the
score. }`

After that, we compute the pass@10 for each query q , which corresponds to the number of correct
responses among the 10 generated candidates, i.e., m . For $m \geq 6$, we directly select the y_i^+ with the
highest score as the generated CoT. If multiple y_i^+ share the highest score, we randomly choose one.

For queries with $1 \leq m \leq 5$, we synthesize a reflective-pattern-injected CoT. Specifically, we first
select one of the correct responses y_i^+ with the highest score and then randomly select one of the
incorrect responses y_i^- . The reflective-pattern-injected CoT is subsequently synthesized through the
following operation:

Synthesis of the reflective-pattern-injected CoT

$\{y_i^-\}$ Wait, perhaps we could consider it from a different perspective. Let's re-evaluate the
problem step by step to ensure accuracy. $\{y_i^+\}$

Finally, we obtain CoT data enriched with reflective patterns through the integration of the aforementioned
data, and we will release it once it is ready.

A.3 REWARD FUNCTION

The other tasks' reward functions are as follows:

864 • **Mathematical Tasks:** For tasks involving mathematical expressions or numerical answers, $R_{\text{accuracy}}(o, \text{gt})$ is determined by a specialized verification function, denoted `math_verify`($o_{\text{ans}}, \text{gt}$). This function evaluates the extracted answer from the output o , denoted o_{ans} , against the ground truth answer gt . The `math_verify` function is designed to handle nuances of mathematical evaluation, potentially allowing for symbolic equivalence or specified numerical tolerances. A successful verification yields a reward of 1; otherwise, 0.

865 • **Grounding Tasks:** For tasks where a model predicts a bounding box, we use the Intersection over

866 Union (IoU) as the reward. This score measures the overlap between the predicted and ground-truth

867 bounding boxes.

873 **A.4 DETAILS OF THE EXPERIMENTAL SETUP**

875 **Training Datasets.** In the SFT stage, we use a total of **188K** training samples from three categories: (1) multimodal general data (LLaVA-OneVision (Li et al., 2024a)), (2) multimodal medical data (VQA-RAD (Lau et al., 2018), SLAKE (Liu et al., 2021), PathVQA (He et al., 2020), PMC-VQA (Zhang et al., 2023c), and our synthetic reflective-pattern-injected CoT), and (3) text-based medical data (ReasonMed (Sun et al., 2025), Medical-R1-Distill (Chen et al., 2024a), and Medical-o1-Reasoning (Chen et al., 2024a)). During the RLVR stage, we employ **36K** samples from multimodal general and medical datasets, including VQA-RAD, SLAKE, PathVQA, PMC-VQA, GMAI-Reasoning (Su et al., 2025), IconQA (Lu et al., 2021), ScienceQA (Lu et al., 2022a), TabMWP (Lu et al., 2022b), TQA (Zhou, 2025), and ViRFT-COCO (Liu et al., 2025e).

884 A detailed description of each dataset is provided as follows:

- 885 • **LLaVA-OneVision** (Li et al., 2024a): LLaVA-OneVision is a large-scale multimodal dataset comprising 4.8 million samples collected from diverse sources. It includes single-image, multi-image, and video modalities, and is specifically designed to train vision-language models for unified visual and textual understanding.
- 886 • **VQA-RAD** (Lau et al., 2018): VQA-RAD is a medical visual question answering dataset constructed for assessing multimodal understanding of radiology. It consists of radiological images paired with manually curated question-answer pairs authored by clinical experts. The dataset includes both open-ended and binary (yes/no) questions.
- 887 • **SLAKE** (Liu et al., 2021): SLAKE is a medical visual question answering dataset comprising 642 annotated radiological images spanning 39 anatomical structures and 12 disease categories. The dataset includes conditions such as various cancers (e.g., brain, liver, kidney, lung) and thoracic diseases (e.g., atelectasis, pleural effusion, pulmonary masses, and pneumothorax).
- 888 • **PathVQA** (He et al., 2020): PathVQA is a large-scale dataset developed for medical visual question answering tasks in the domain of pathology. It comprises 32,799 expert-annotated question-answer pairs spanning seven question categories, grounded in 4,998 high-resolution pathology images. The dataset includes both binary (yes/no) and open-ended questions.
- 889 • **PMC-VQA** (Zhang et al., 2023c) PMC-VQA is a large-scale medical visual question answering dataset designed to facilitate research on multimodal understanding in the medical domain. It comprises 227K VQA pairs grounded in 149K medical images, covering a wide range of imaging modalities and disease types.
- 890 • **ReasonMed** (Sun et al., 2025): ReasonMed is the largest open-source medical textual reasoning dataset containing 370K QA examples, which is distilled and filtered from three competitive large-language models (Qwen-2.5-72B, DeepSeek-R1-Distill-Llama-70B, and HuatuoGPT-o1-70B).
- 891 • **Medical-R1-Distill** (Chen et al., 2024a): Medical-R1-Distill-Data is an SFT dataset distilled from the DeepSeek-R1, constructed on verifiable medical questions from HuaTuoGPT-o1. It provides reasoning chains for medical problems, enabling the initialization and supervision of models' reasoning processes in the medical domain.
- 892 • **Medical-o1-Reasoning** (Chen et al., 2024a): medical-o1-reasoning-SFT is a SFT dataset focused on verifiable medical problems, where candidate solutions are generated by GPT-4o and validated by a medical verifier, providing high-quality reasoning chains and answers for training medical reasoning models.
- 893 • **GMAI-Reasoning** (Su et al., 2025): GMAI-Reasoning10K is a high-quality medical visual reasoning dataset comprising 10K curated multiple-choice questions constructed from 95 publicly available medical datasets spanning 12 imaging modalities (e.g., X-ray, CT, MRI). Each question is paired with standardized visual inputs and metadata, and generated using GPT-based prompting, following rigorous preprocessing and quality control procedures.

- 918 • IconQA (Lu et al., 2021): IconQA is a large-scale dataset containing 107,439 questions designed
919 to assess abstract icon image understanding and visual language reasoning abilities.
- 920 • ScienceQA (Lu et al., 2022a): ScienceQA contains 21k multimodal questions, which align with
921 California Common Core Content Standards, covering diverse science domains, many enriched
922 with images, lectures, and explanations to support reasoning-oriented training.
- 923 • TabMWP (Lu et al., 2022b): Tabular Math Word Problems (TabMWP) is a multimodal dataset
924 designed for training models to solve math word problems using both textual and tabular data. It
925 contains 38,431 problems spanning elementary to high school levels, including both free-text and
926 multiple-choice questions.
- 927 • TQA (Zhou, 2025): Textbook Question Answering (TQA) is a multimodal dataset designed for
928 training models to answer questions using both textual and visual content from middle school
929 science textbooks. Each sample provides a question, relevant textual context, and associated images,
930 enabling models to learn to reason over multimodal inputs and generate accurate answers.
- 931 • ViRFT_COCO (Liu et al., 2025e): ViRFT_COCO is a vision-language dataset derived from COCO,
932 containing around 6,000 samples. It aims to enhance models' ability to detect all instances of a
933 given category within an image and output the corresponding bounding boxes with confidences
934 under strict formatting constraints.

935 **Implementation Details.** Our InfiMed-Series models include InfiMed-SFT-3B and InfiMed-RL-3B.

- 936 • InfiMed-SFT-3B, which is built upon Qwen2.5-VL-3B (Bai et al., 2025), is trained using LLaMA-
937 Factory (Zheng et al., 2024). We utilize 8 NVIDIA H800 GPUs. The vision tower and multimodal
938 projector are frozen during training, while the language model remains fully trainable. We use a
939 cosine learning rate scheduler with an initial learning rate of 5×10^{-6} , a warmup ratio of 0.1, and
940 train for 5 epochs. The batch size is set to 4 per device. Furthermore, we set the maximum input
941 resolution to 262,144 pixels for images, while text inputs are truncated to a maximum length of
942 4,096 tokens.
- 943 • InfiMed-RL-3B is built upon InfiMed-SFT-3B via EasyR1 (Sheng et al., 2024). For the RLVR
944 reward function $R_{\text{total}}(o, \text{gt}) = w_{\text{format}} \cdot R_{\text{format}}(o) + w_{\text{acc}} \cdot R_{\text{accuracy}}(o, \text{gt})$, we set the weights
945 $w_{\text{format}} = 0.1$ and $w_{\text{acc}} = 0.9$. All experiments were conducted using 16 NVIDIA H800 GPUs.
946 For each phase, we used a learning rate of 1.0×10^{-6} , a batch size of 256 for training updates, a
947 rollout batch size of 256, and generated 16 rollouts per sample during policy exploration.

948 **Evaluation Framework** To ensure consistency with prior work and a comprehensive, standardized
949 evaluation, we adopt MedEvalKit (Xu et al., 2025), a systematic framework that integrates mainstream
950 medical benchmarks and task types, supporting a range of question formats, including multiple-choice
951 questions, open-ended questions, and closed-ended questions. We adopt the multimodal evaluation
952 component of the framework, combining rule-based methods with the LLM-as-a-Judge strategy.

953 **Evaluation Benchmarks** We evaluate our InfiMed-Series models on seven widely used multimodal
954 medical benchmarks, assessing both their reasoning ability and their understanding of medical
955 knowledge. The detailed description of the benchmarks is as follows:

- 956 • MMMU (Yue et al., 2024): MMMU is a benchmark designed to assess the capabilities of multi-
957 modal models on large-scale, multidisciplinary tasks. It comprises 11.5K meticulously curated
958 multimodal questions drawn from university exams, quizzes, and textbooks, covering six core
959 disciplines, including Health & Medicine. The Health & Medicine includes 1,752 test ques-
960 tions—accounting for 17% of the entire benchmark—and is further subdivided into five specialized
961 domains: Basic Medical Science, Clinical Medicine, Diagnostics and Laboratory Medicine, Phar-
962 macy, and Public Health.
- 963 • VQA-RAD (Lau et al., 2018): VQA-RAD is a dataset consisting of question-answer pairs grounded
964 in radiological medical images, intended for training and evaluating medical visual question
965 answering systems. It includes both open-ended questions and binary yes/no questions. In total,
966 the dataset comprises 2,248 QA pairs linked to 315 medical images, with all annotations manually
967 curated by a team of clinicians to ensure clinical relevance and accuracy.
- 968 • SLAKE (Liu et al., 2021): SLAKE is a bilingual (Chinese-English) dataset specifically designed
969 for medical visual question answering systems. It consists of 642 medical images paired with
970 14,028 question-answer instances.
- 971 • PathVQA (He et al., 2020): PathVQA is designed for visual question answering in the field of
pathology. It comprises 4,998 pathology images collected from two pathology textbooks and the
PEIR digital library, accompanied by a total of 32,799 question-answer pairs.

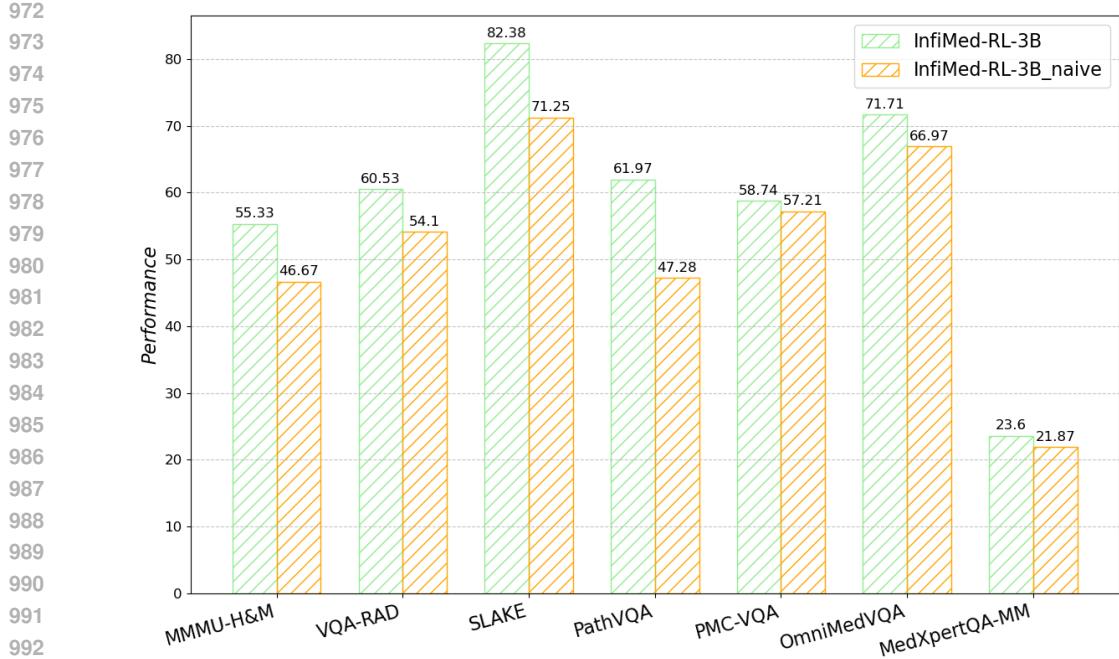


Figure 6: Performance comparison of InfiMed-RL-3B and InfiMed-RL-3B_naive on medical benchmarks. InfiMed-RL-3B_naive denotes directly utilizing RLVR upon Qwen2.5-VL-3B.

- PMC-VQA (Zhang et al., 2023c): PMC-VQA is a large-scale multimodal dataset constructed for medical visual question answering. It contains 227,000 VQA questions grounded in 149,000 medical images spanning a wide range of imaging modalities and disease types.
- OmniMedVQA (Hu et al., 2024): OmniMedVQA is a large-scale and comprehensive visual question answering benchmark tailored specifically for the medical domain. It aggregates data from 73 distinct medical datasets, comprising 118,010 images and 127,995 question-answer pairs. The benchmark encompasses 12 different medical imaging modalities and covers more than 20 anatomical regions of the human body.
- MedXpertQA (Zuo et al., 2025): MedXpertQA is a benchmark specifically designed to evaluate professional medical knowledge. It comprises 4,460 questions spanning 17 medical specialties and 11 organ systems. In our experiments, we utilize only the multimodal subset of the dataset.

A.5 ABLATION STUDY ON RLVR

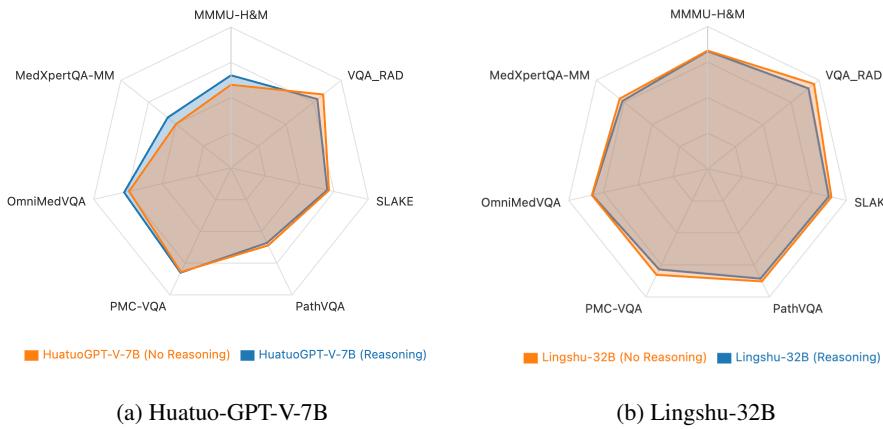
In this subsection, we present ablation studies related to RLVR, where we explore training from two different starting points: one from the Qwen2.5-VL-3B model (Bai et al., 2025), and the other from our InfiMed-SFT-3B. The results are presented in Figure 6.

Notably, in tasks requiring the integration of large amounts of domain-specific information, such as VQA-RAD, PathVQA, and SLAKE, the InfiMed-RL-3B_naive model significantly underperforms compared to InfiMed-RL-3B. This suggests that directly applying RLVR to Qwen2.5-VL-3B without incorporating domain-specific data during the SFT phase can lead to much lower performance, especially on tasks that require understanding and memorizing domain-specific knowledge. This underscores the importance of the cold start phase, where injecting relevant, knowledge-rich data during SFT is essential to building a solid foundation for the subsequent RLVR phase.

Additionally, in MMMU-H&M, InfiMed-RL-3B achieves a significantly higher score of 55.33, compared to 46.67 for InfiMed-RL-3B_naive. Given that this benchmark demands both reasoning and comprehensive multimodal understanding (general and medical), this highlights the critical role of the SFT phase in helping the model integrate complex information effectively. The results on PMC-VQA and MedXpertQA-MM further demonstrate that prior domain-specific fine-tuning improves RLVR training outcomes.

1026 **Table 3: Ablation study examining data composition during the RLVR stage.** $\Delta|\text{Data}|$ denotes
 1027 the changes of the training dataset.

Model	$\Delta \text{Data} $	Accuracy (%)						
		MMMU-H&M	VQA-RAD	SLAKE	PathVQA	PMC-VQA	OMVQA	MedXQA
<i>Base Model</i>								
Qwen2.5VL-3B	-	51.3	56.8	63.2	37.1	50.6	64.5	20.7
<i>Ablation Study in RLVR Stage</i>								
InfiMed-RL-3B	-	55.3	60.5	82.4	62.0	58.7	71.7	23.6
InfiMed-RL-3B+medical_mm	+16K	56.0	60.9	82.2	61.5	58.4	70.0	23.6
InfiMed-RL-3B-w/o-general	-10K	53.3	60.7	81.9	61.6	58.3	70.0	23.6
								58.4



1049 **Figure 7: Comparison of direct-answer and reasoning-based prompts on medical benchmarks with
 1050 larger models.**

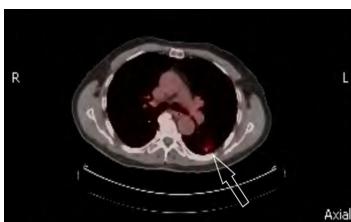
1052 We further explored whether enlarging the RLVR dataset could lead to additional performance gains.
 1053 To this end, we constructed a larger RLVR set by incorporating 16K additional medical questions
 1054 sourced from SLAKE (5K), PathVQA (9K), and VQA-RAD (2K). These questions were selected by
 1055 prioritizing cases in which InfiMed-SFT-3B produced the fewest correct responses across 10 attempts.
 1056 The experimental results in the Table 3 show that simply increasing the number of RLVR training
 1057 questions does not consistently yield performance improvements. This finding suggests that, although
 1058 additional RLVR data may still provide incremental benefits, the marginal gains are limited given the
 1059 substantial effort required to curate high-quality medical multimodal RLVR data.

1061 A.6 CASE STUDIES ON QWEN2.5-VL-3B, INFI MED-SFT-3B, AND INFI MED-RL-3B

1062 In this section, we present additional case studies to illustrate the distinct responses of Qwen2.5-VL-
 1063 3B, InfiMed-SFT-3B, and InfiMed-RL-3B. In summary, our analysis reveals that the InfiMed-SFT-3B
 1064 model already obtains self-reflective ability but is prone to delivering redundant responses, whereas
 1065 the InfiMed-RL-3B model minimizes verbosity while ensuring accuracy in its answers.

1066 This change is consistent with the characteristics of the GRPO optimization mechanism. During
 1067 GRPO training, the gradient update for each rollout response is normalized by its length, which
 1068 amplifies the learning signal for shorter and correct responses. Consequently, the reinforcement
 1069 learning stage naturally encourages the model to prefer concise and accurate answers rather than
 1070 extended reflective chains. The reflective supervision used during SFT primarily serves to enrich
 1071 the model’s reasoning search space for RLVR, rather than to enforce reflective verbosity in the final
 1072 model. After RLVR, the model shifts from imitating reflection-style formats to producing more
 1073 efficient outputs, which explains the observed reduction in explicit reflective content.

1080
1081
1082
1083
1084
1085
1086
1087
1088
1089
1090
1091
1092
1093
1094
1095
1096
1097
1098
1099
1100
1101
1102



Question:
Which lobe of the lung was affected by the nodule seen on the PET/CT scan?
A: Right middle lobe
B: Left lower lobe
C: Right upper lobe
D: Left upper lobe

Ground Truth:
B: Left lower lobe

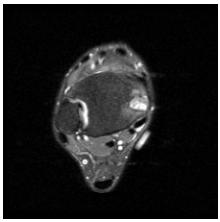
Qwen2.5-VL-3B Answer:
To determine which lobe of the lung was affected by the nodule, we need to consider the anatomical location of the nodule in relation to the lung lobes. Typically, the right lung has three lobes (upper, middle, and lower), while the left lung has two lobes (upper and lower). The nodule is located near the heart, which is positioned more towards the center of the chest. Given this information, it is more likely that the nodule is in the left lower lobe. **The final answer is B.**

InfiMed-SFT-3B Answer:
To determine which lobe of the lung is affected by the nodule seen on the PET/CT scan, let's analyze the image step by step:
1. **Orientation**: The image is labeled with "R" (right) and "L" (left), indicating the patient's right and left sides. The nodule is located towards the left side of the image.
2. **Anatomical Location**: The nodule appears to be in the central part of the chest, near the mediastinum. It is positioned in the area that corresponds to the left lung.
3. **Lung Lobes**: The right lung has three lobes: upper, middle, and lower. The left lung has two lobes: upper and lower.
- The nodule is located in the central area, closer to the mediastinum, which is consistent with the location of the left upper lobe. **Given these observations, the nodule is most likely affecting the "left upper lobe".**
Wait, perhaps we could consider it from a different perspective. Let's re-evaluate the problem step by step to ensure accuracy. The image provided is a CT scan showing a cross-sectional view of the thorax. The arrow points to a region in the lower part of the image, which corresponds to the left side of the patient. The structure indicated by the arrow appears to be in the lower lobe of the lung. Given the orientation: The left side of the image corresponds to the patient's right side. The right side of the image corresponds to the patient's left side. **Thus, the structure indicated by the arrow is in the "left lower lobe".**

InfiMed-RL-3B Answer:
The PET/CT scan shows a nodule in the left lower lobe of the lung, as indicated by the arrow pointing to the left side of the image. The orientation of the scan is such that the right side of the patient's body is on the left side of the image, and the left side of the patient's body is on the right side of the image. Therefore, **the nodule is in the left lower lobe.**

1103
1104
1105
1106
1107
1108
1109
1110
1111
1112
1113
1114
1115
1116
1117
1118
1119
1120
1121
1122
1123
1124
1125
1126
1127
1128
1129
1130
1131
1132
1133

Figure 8: Case study on Qwen2.5-VL-3B, InfiMed-Series models in medical VQA. **Red** denotes errors or irrelevant content, whereas **Green** denotes correct or important information. All of the models answered correctly; InfiMed-SFT-3B answered correctly after reflection.



Question:
What can be observed in this image?
A) Gastric ulcer
B) Pulmonary embolism
C) Cardiac arrhythmia
D) Osseous neoplasm

Ground Truth:
D) Osseous neoplasm

Qwen2.5-VL-3B Answer:
To determine what can be observed in this image, we need to consider the typical characteristics of each condition mentioned:
A) **Gastric ulcer:** This is a lesion in the stomach lining that appears as a hypoechogenic (dark) area on ultrasound.
B) **Pulmonary embolism:** This is a blockage in a blood vessel in the lungs and is not typically imaged using MRI.
C) **Cardiac arrhythmia:** This involves irregular heartbeats and is usually assessed through ECG or other imaging modalities like echocardiography.
D) **Osseous neoplasm:** This refers to a tumor in bone tissue and would appear as an abnormal mass on imaging.
Given the appearance of the image, which shows a dark area with a surrounding rim, it is most consistent with a gastric ulcer. Therefore, the answer is A.

InfiMed-SFT-3B Answer:
The image appears to be a medical scan, likely an MRI or CT scan, showing a cross-sectional view of a structure that resembles the heart. The dark central area could represent the cardiac chambers, and the surrounding structures might indicate the presence of a mass or abnormality.
Given the options:
A) **Gastric ulcer**: This is typically visualized on an upper GI series or endoscopy, not an imaging scan of the heart.
B) **Pulmonary embolism**: This would show up as a blockage in the pulmonary arteries, usually seen on a CT pulmonary angiogram or V/Q scan, not a cardiac scan.
C) **Cardiac arrhythmia**: This could be inferred from the presence of an abnormal mass or structure affecting the heart's electrical system, which is consistent with the appearance of a cardiac neoplasm.
D) **Osseous neoplasm**: This would show bone abnormalities, which are not visible in this type of scan.
The most plausible observation based on the image is related to the heart, making "cardiac arrhythmia" the least likely option, while "osseous neoplasm" is also unlikely.
Thus, the most appropriate answer is:

InfiMed-RL-3B Answer:
The image appears to be a medical scan, likely an MRI or CT scan, showing a cross-sectional view of a structure that resembles a bone or a bony structure. The presence of a well-defined, possibly abnormal, area within the structure suggests the possibility of a neoplasm (tumor) or a lesion. **The context and appearance are consistent with imaging findings of a bone tumor.**

1128
1129
1130
1131
1132
1133

Figure 9: Case study on Qwen2.5-VL-3B, InfiMed-Series models in medical VQA. **Red** denotes errors or irrelevant content, whereas **Green** denotes correct or important information. Only InfiMed-RL-3B answered correctly, providing a direct and accurate response by correctly identifying the key visual information and selecting the correct answer with a concise and focused explanation.

1134

1135

1136

1137

1138

1139

1140

1141

1142

1143

1144

1145

1146

1147

1148

1149

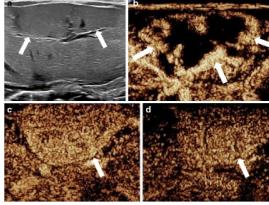
1150

1151

1152

1153

1154



Question:
What diagnostic tool was used to visualize the patient's liver lesion?
A: CT scan B: X-ray C: MRI D: Ultrasound

Ground Truth:
D: Ultrasound

Qwen2.5-VL-3B Answer:
The images provided in the figure show cross-sectional views of the liver, which is characteristic of a CT scan. **Ultrasound images typically do not show such detailed cross-sectional views of organs.** X-rays are not commonly used for visualizing liver lesions in this manner. MRI images often have different characteristics compared to the ones shown here. **The answer is A.**

InfiMed-SFT-3B Answer:
The images provided are grayscale and show a cross-sectional view with clear visualization of internal structures, including what appears to be a liver lesion. The presence of white arrows pointing to specific areas suggests a diagnostic imaging technique that allows for detailed visualization of soft tissues.
CT scan (A): Typically shows detailed anatomical structures but is not as sensitive for soft tissue differentiation as ultrasound.
X-ray (B): Primarily used for bone structures and does not provide detailed soft tissue images.
MRI (C): Provides excellent soft tissue contrast but is less commonly used for real-time imaging compared to ultrasound.
Ultrasound (D): Uses sound waves to create images and is excellent for visualizing soft tissues, including liver lesions. It is also real-time, which aligns with the dynamic nature of the images shown.
Given the characteristics of the images, the most likely diagnostic tool used is **ultrasound**.
\\boxed(D)

InfiMed-RL-3B Answer:
<think>The image provided appears to be a series of ultrasound scans, as indicated by the characteristic grayscale and the appearance of the tissue structures. **Ultrasound is commonly used for visualizing internal organs, including the liver, and is particularly useful for detecting lesions due to its ability to show detailed structural information in real-time.** CT scans and X-rays typically provide cross-sectional images and are less detailed in real-time visualization of soft tissues like the liver. MRI provides detailed images but is usually not the first choice for real-time visualization of liver lesions, and it is less likely to be the format shown here.
Therefore, the most appropriate diagnostic tool used here is an ultrasound.</think>
\\boxed(D)

Figure 10: Case study on Qwen2.5-VL-3B, InfiMed-Series models in medical VQA. **Red** denotes errors or irrelevant content, whereas **Green** denotes correct or important information. Qwen2.5-VL-3B’s response is incorrect because it fundamentally fails to recognize the visual characteristics of an ultrasound scan, leading to a flawed conclusion. InfiMed-SFT-3B provides a detailed, step-by-step reasoning process, while InfiMed-RL-3B offers a more direct and accurate answer, showcasing its improved ability to instantly recognize diagnostic imaging types.

1160

1161

1162

1163

1164

1165

1166

1167

1168

1169

1170

1171

1172

1173

1174

1175

1176

1177

1178

1179

1180

1181



Question:
What is being displayed in this ultrasound image?
A) Esophagus B) uterus C) Testis D) Prostate

Ground Truth:
(B) uterus

Qwen2.5-VL-3B Answer:
The ultrasound image shows a structure with a central area that appears to be fluid-filled, surrounded by a more solid, echogenic (bright) area. This is characteristic of a testicular ultrasound. The central area could represent the epididymis or spermatic cord, while the surrounding area represents the testicle.
Therefore, the answer is C.

InfiMed-SFT-3B Answer:
To determine what is being displayed in the ultrasound image, let's analyze the visual characteristics:
1. **Shape and Structure**: The image shows a structure that appears to be a solid, oval-shaped mass with some internal echoes. This suggests it could be a gland or organ.
2. **Location**: The structure is located in a region that is consistent with the location of the prostate gland in males.
3. **Features**:
- The image shows a relatively homogeneous structure with some internal echoes.
- There are no clear indications of the typical features of other structures like the esophagus, uterus, or testis.
Given these observations, the most likely structure being displayed is the prostate gland.
\\boxed(D)

InfiMed-RL-3B Answer:
The ultrasound image shows a structure that appears to be the uterus, given its characteristic appearance with a central area that might represent the endometrial cavity and surrounding tissue. **The structure is consistent with the typical appearance of a uterus in an ultrasound scan.**
\\boxed(B)

Figure 11: Case study on Qwen2.5-VL-3B, InfiMed-Series models in medical VQA. **Red** denotes errors or irrelevant content, whereas **Green** denotes correct or important information. Incorrectly identifies the image and chooses the wrong answer. InfiMed-SFT-3B provides a detailed analysis that correctly rules out most options but ultimately guesses the wrong answer. InfiMed-RL-3B is the only model that correctly identifies the organ shown, providing the most direct and accurate response.

1188

1189

1190

1191

1192

1193

1194

1195

1196

1197

1198

1199

1200

1201

1202

1203

1204

1205

1206

1207

1208

1209

1210

1211

1212

1213

1214

1215

1216

1217

1218

1219

1220

1221

1222

1223

1224

1225

1226

1227

1228

1229

1230

1231

1232

1233

1234

1235

1236

1237

1238

1239

1240

1241

**Question:**

An 18-year-old female athlete who plays basketball comes to the clinic with persistent knee pain that has been unresponsive to anti-inflammatory medication. The plain radiograph demonstrates a lesion located in the femoral metaphysis (figure shown). What is the most probable diagnosis?

- Ewing sarcoma
- Osteosarcoma
- Eosinophilic granuloma
- Enchondroma
- Osteoid osteoma

Ground Truth: (B) Osteosarcoma**Qwen2.5-VL-3B Answer:**

The plain radiograph shows a lytic lesion in the metaphysis of the femur, which is characteristic of Ewing sarcoma. This type of tumor typically presents with bone pain and can be aggressive, often requiring surgical intervention. Therefore, the answer is A.

InfiMed-SFT-3B Answer:

 The plain radiograph shows a lesion in the femoral metaphysis of an 18-year-old female athlete. Let's evaluate the options ... The presence of a sunburst pattern on imaging would support this diagnosis.
Wait, perhaps we could consider it from a different perspective. Let's re-evaluate the problem step by step to ensure accuracy. To determine the most probable diagnosis for the lesion in the femoral metaphysis of an 18-year-old female athlete, let's analyze the provided information and the options:
1. **Patient Demographics and Activity:** ...
2. **Lesion Characteristics:** ...
3. **Option Analysis:** ...
4. **Most Likely Diagnosis:** Given the patient's age, activity level, and the characteristics of the lesion (well-circumscribed and radiolucent), the most probable diagnosis is **enchondroma**.
Therefore, the answer is **boxed(D)**

InfiMed-RL-3B Answer:

 The patient is an 18-year-old female athlete with persistent knee pain unresponsive to medication. The plain radiograph shows a lesion in the femoral metaphysis. The options provided include various bone lesions, and the context suggests a young adult with a specific location of the lesion.
Ewing sarcoma typically affects younger individuals and often occurs in the diaphysis or metaphysis of long bones, but it is more common in the pelvis and femur.
Osteosarcoma is another common primary bone tumor, often occurring in adolescents and young adults, and it typically affects the metaphysis of long bones.
Eosinophilic granuloma is a form of Langerhans cell histiocytosis, which can occur in young adults but is less common.
Enchondroma is a benign cartilage tumor that usually occurs in the hands and feet, and it is less likely to present with such a localized lesion in the femur.
Osteoid osteoma is a benign bone tumor that typically presents with pain that improves with NSAIDs and is often located in the diaphysis of long bones.
Given the age, location, and presentation, the most probable diagnosis is **osteosarcoma**.
Therefore, the answer is **boxed(B)**

Figure 12: Case study on Qwen2.5-VL-3B, InfiMed-Series models in medical VQA. Red denotes errors or irrelevant content, whereas Green denotes correct or important information. Qwen2.5-VL-3B incorrectly identifies the lesion. InfiMed-SFT-3B initially makes a correct deduction but then overthinks the problem, leading to a self-correction that results in the wrong answer. InfiMed-RL-3B provides a detailed and accurate analysis.

Question:

A 62-year-old male patient is admitted to the hospital with symptoms including bilateral muscle aches and weakness affecting the anterior thighs and lower posterior legs, persisting for one month. He reports a 10 kg weight loss. Physical examination reveals numbness in the anterior thighs and posterior lower legs, without any skin abnormalities or abdominal tenderness. Laboratory results indicate elevated inflammatory markers. CT scans of chest, abdomen, and pelvis show no significant findings, and ANCA testing is negative. After performing abdominal angiography (images shown), which of the following conditions is most commonly associated with this patient's underlying diagnosis?

- Systemic lupus erythematosus
- Chronic lymphocytic leukemia
- HIV
- Tuberculosis
- Hepatitis B

**Ground Truth: (E) Hepatitis B****Qwen2.5-VL-3B Answer:**

 The patient's symptoms include bilateral muscle aches and weakness, weight loss, and numbness in the anterior thighs and posterior lower legs. These symptoms are consistent with vasculitis, which can be caused by various conditions such as systemic lupus erythematosus (SLE), chronic lymphocytic leukemia (CLL), HIV, tuberculosis (TB), and hepatitis B. However, the abdominal angiography images show multiple small vessel occlusions, which is more suggestive of a vasculitis that affects small vessels. The correct answer is A.

InfiMed-SFT-3B Answer:

 The images provided are angiograms showing abnormal vascular structures, likely indicative of a vascular malformation or abnormal blood flow. Given the clinical presentation of bilateral muscle aches and weakness, numbness, weight loss, and elevated inflammatory markers, the most likely diagnosis is **polyarteritis nodosa (PAN)**, which is a systemic vasculitis affecting small- to medium-sized arteries.
While SLE can cause vasculitis, it typically affects small to medium-sized vessels and is more commonly associated with other symptoms such as joint pain, malar rash, and serositis. The angiographic findings do not strongly suggest SLE.
Chronic lymphocytic leukemia (CLL) is a cancer of the white blood cells and does not typically present with the vascular abnormalities seen here. The angiographic findings are more consistent with a vasculitis rather than a hematologic malignancy.
HIV can lead to various complications, including vasculitis, but the angiographic findings are not characteristic of HIV-related vasculitis. Additionally, the patient's age and presentation do not strongly suggest HIV.
Tuberculosis can cause vasculitis, particularly in the context of granulomatous disease. However, the angiographic findings are not typical for TB-related vasculitis, and the patient's presentation does not strongly suggest TB.
Hepatitis B can cause liver-related complications and is not typically associated with the vascular abnormalities seen in the angiograms. The patient's presentation does not strongly suggest hepatitis B.
Conclusion: The angiographic findings are most consistent with **Polyarteritis nodosa (PAN)**, which is a systemic vasculitis that can affect medium-sized arteries. Among the given options, **HIV** is the least likely, as the angiographic findings do not align with typical HIV-related vasculitis. Therefore, the most appropriate answer is C.

InfiMed-RL-3B Answer:

 The patient's symptoms and clinical presentation suggest a vasculitis, likely involving the small to medium-sized vessels. The abdominal angiography images show abnormalities consistent with vasculitis, such as vessel narrowing or occlusion. The most common vasculitis associated with abdominal angiography findings in this context is polyarteritis nodosa, which is often linked to hepatitis B virus (HBV) infection. Therefore, the most likely underlying diagnosis is hepatitis B. **boxed(E)**

Figure 13: Case study on Qwen2.5-VL-3B, InfiMed-Series models in medical VQA. Red denotes errors or irrelevant content, whereas Green denotes correct or important information. Only InfiMed-RL-3B successfully synthesizes the angiographic findings and patient symptoms to provide the correct diagnosis.

1242

1243

1244

1245

1246

1247

1248

1249

1250

1251

1252

1253

1254

1255

1256

1257

1258

1259

1260

1261

1262

1263

1264

1265

1266

1267

1268

1269

1270

1271

**Question:**

A 6-year-old boy presents with a forearm deformity, and radiographs in Figure A confirm the diagnosis. What additional diagnostic tests should be performed for this patient?

- (A) CBC, renal ultrasound, and echocardiogram
- (B) CBC and MRI
- (C) CBC and renal ultrasound only
- (D) No additional testing needed
- (E) CBC

Ground Truth:

(A) CBC, renal ultrasound, and echocardiogram

Qwen2.5-VL-3B Answer:

The patient has a forearm deformity that is confirmed by the radiographs. The most common causes of forearm deformities include radial head dislocation, radial head fracture, and radial head subluxation. These conditions can be diagnosed based on the clinical presentation and physical examination findings. However, if there is any suspicion of a more serious underlying condition such as radial head dislocation or fracture, further imaging studies like MRI may be necessary to rule out other potential diagnoses. Therefore, the answer is B.

InfiMed-SFT-3B Answer:

To determine the appropriate additional diagnostic tests for a 6-year-old boy presenting with a forearm deformity, let's analyze the situation step by step:
 1. **Clinical Presentation:** The patient is a young child with a forearm deformity. This could be due to various reasons such as congenital conditions, trauma, or metabolic disorders.
 2. **Radiographic Findings:** The provided image shows radiographs of the forearm and hand. The specific findings would help in diagnosing conditions like rickets, osteogenesis imperfecta, or other bone disorders. However, without specific details from the radiographs, we cannot pinpoint the exact condition.
 3. **Diagnostic Considerations:** - **CBC (Complete Blood Count):** Useful for detecting anemia, infection, or other hematological abnormalities.
 - **Renal Ultrasound:** To evaluate kidney function and look for structural abnormalities that might be related to renal issues.
 - **Echocardiogram:** To assess cardiac function and detect any underlying cardiovascular issues that could be related to metabolic disorders.
 4. **Additional Tests:** - **CBC:** Essential to rule out systemic conditions like anemia or infection.
 - **Renal Ultrasound:** Important to check for renal abnormalities, especially if there is suspicion of renal tubular acidosis or other renal metabolic disorders.
 5. **Conclusion:** Given the clinical presentation and the need to rule out systemic conditions, the most appropriate additional tests would include a CBC, renal ultrasound, and echocardiogram.

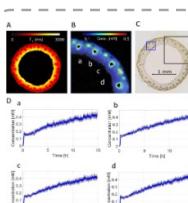
Thus, the correct answer is: **A**.

InfiMed-RL-3B Answer:

The patient is a 6-year-old boy with a forearm deformity, and the radiographs suggest a potential bone or joint issue. Given the age and presentation, it's important to rule out any underlying systemic conditions or complications. A complete blood count (CBC) is a basic test that can help identify infections, anemia, or other systemic issues. A renal ultrasound might be necessary to check for renal abnormalities or conditions that could affect bone health. An echocardiogram could be important if there are concerns about cardiac involvement, especially if the deformity is related to a systemic condition. Therefore, the most appropriate additional tests would include a CBC, renal ultrasound, and echocardiogram.

A

Figure 14: Case study on Qwen2.5-VL-3B, InfiMed-Series models in medical VQA. **Red** denotes errors or irrelevant content, whereas **Green** denotes correct or important information. Qwen2.5-VL-3B incorrectly identifies the necessary tests. Both InfiMed-SFT-3B and InfiMed-RL-3B correctly solve this case, but InfiMed-RL-3B provides a more direct and efficient answer.

**Question:**

What does Figure B show?

- (A) The concentration of CA after 7 hours of administration.
- (B) The concentration of CA after 14 hours of administration.
- (C) The concentration of CA after 21 hours of administration.
- (D) The concentration of CA after 28 hours of administration.

Ground Truth:

(B) The concentration of CA after 14 hours of administration.

Qwen2.5-VL-3B Answer:

In the figure, the x-axis is labeled 'Time [h]' and the y-axis is labeled 'Conc. [mM]', indicating that it shows the concentration of CA over time. The graphs in panel D correspond to different time points, with the graph labeled 'a' showing data from 0 to 7 hours, 'b' from 0 to 14 hours, 'c' from 0 to 21 hours, and 'd' from 0 to 28 hours.

Therefore, the correct answer is (C) 'The concentration of CA after 21 hours of administration.'

InfiMed-SFT-3B Answer:

Figure B shows the concentration of CA (likely referring to a compound like curcumin or another substance) over time, with labeled points a, b, c, and d indicating different regions or stages of concentration change. The graphs in Figure D correspond to these labeled points, showing the concentration over time. Since the graphs in Figure D represent data points at regular intervals (likely hourly), and considering the options provided, the correct answer is: **A**

Answer Choice: (A) The concentration of CA after 7 hours of administration.

InfiMed-RL-3B Answer:

Okay, let's see. The question is asking what Figure B shows. The image is divided into parts A, B, C, and D. Figure B is the middle panel, I think. Looking at Figure B, it has a color bar that goes from 0 to 0.5 mM, and there are labeled points a, b, c, d. These points are likely indicating different locations in the sample. The graph in Figure D corresponds to these points, with time on the x-axis and concentration on the y-axis. The graphs in D (a, b, c, d) show concentration over time. The concentration starts at around 0.1 mM and increases over 15 hours. Since the graphs are for the same time period, and the concentration increases linearly, it suggests that Figure B is showing the concentration profile of a substance over time. Now, the question is about Figure B specifically. Figure B is a map showing the distribution of concentration across the sample. The labeled points a, b, c, d correspond to the locations in the sample where the concentration is being measured. Since the graphs in D are for these points, and they show the concentration increasing over time, Figure B must be the concentration map after some time of administration. The options given are about the time after administration. The graphs in D show that the concentration increases over 15 hours. So Figure B, which is the concentration map, would be showing the distribution of the substance after a certain period. If the graphs in D are for 15 hours, then Figure B must be the concentration after 14 hours, as the graphs continue to the 15-hour mark.

Therefore, the answer should be B: The concentration of CA after 14 hours of administration.

Figure 15: Case study on Qwen2.5-VL-3B, InfiMed-Series models in medical-related VQA (MMU-H&M). **Red** denotes errors or irrelevant content, whereas **Green** denotes correct or important information. Qwen2.5-VL-3B incorrectly identifies the relevant data series. InfiMed-SFT-3B fails to correctly interpret the graphs and provides an incorrect analysis. Only InfiMed-RL-3B correctly synthesizes the information from both the concentration map and the graphs to provide the right answer.

1296

1297

1298

1299

1300

1301

1302

1303

1304

1305

1306

1307

1308

1309

1310

1311

1312

1313

1314

1315

1316

1317

1318

1319

1320

1321

1322

1323

1324

1325

1326

1327

1328

1329

1330

1331

1332

1333

1334

1335

1336

1337

1338

1339

1340

1341

1342

1343

1344

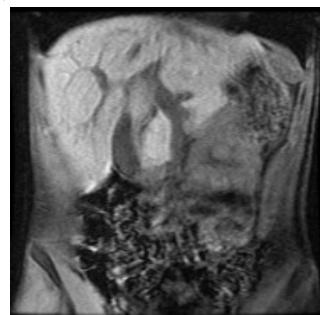
1345

1346

1347

1348

1349

**Question:**

Does this image show any variation from normal?

- It's difficult to determine if there is an abnormality based on this image.
- There is a possibility of abnormality, but further tests are needed for confirmation.
- I'm not sure, but there might be a slight abnormality.
- No, It's normal.

Ground Truth:

D. No, It's normal.

Qwen2.5-VL-3B Answer:

To determine if there is an abnormality in the image, we need to consider the typical appearance of normal tissues and structures. The image shows a cross-sectional view of what appears to be a body part, possibly the neck or throat area, with various structures visible.

Identify Normal Structures: In a normal CT scan of the neck or throat, you would expect to see structures such as the thyroid gland, larynx, trachea, and other soft tissues. These structures should have specific shapes and densities that are characteristic of their normal state.

Analyze Abnormalities: If there are any deviations from these expected shapes or densities, it could indicate an abnormality. For example, an enlarged thyroid gland, a mass in the larynx, or changes in the tracheal wall could suggest pathology.

Consider the Image: Based on the provided image, there are no obvious signs of gross abnormalities like masses, calcifications, or unusual shapes that would typically indicate pathology. However, without a detailed analysis by a radiologist who specializes in interpreting such images, it is challenging to definitively say whether there is an abnormality.

Given the options:

- Option A suggests it's difficult to determine if there is an abnormality based on this image alone.
- Option B implies there is a possibility of abnormality but requires further testing.
- Option C suggests there might be a slight abnormality, which is less definitive than option A.
- Option D states it's normal, which is too definitive given the lack of clear indicators of abnormality.

Considering all factors, the most appropriate answer is: A

InfiMed-SFT-3B Answer:

To determine if the image shows any variation from normal, let's analyze the visual characteristics:

Overall Structure: The image appears to show a cross-sectional view of a structure, possibly a part of the body like the brain or spinal cord. The symmetry and structure seem intact.

Anatomical Features: There are no obvious structural abnormalities such as holes, irregularities, or unusual shapes that would suggest a pathological condition.

Color and Texture: The colors and textures appear consistent with what would be expected in a normal anatomical structure. There are no areas that stand out as abnormal in terms of color or texture.

Comparison to Normal Expectations: Without a reference image or additional context, it is challenging to definitively say whether there is an abnormality. However, the image does not show any clear signs of pathology.

Given these observations, the most appropriate answer is: A.

Wait, perhaps we could consider it from a different perspective. Let's re-evaluate the problem step by step to ensure accuracy.

To determine if the image shows any variation from normal, let's analyze the provided MRI scan step by step:

- General Observations: The image appears to be an MRI scan of the abdomen or pelvis.
- Specific Features: The central area appears to have a relatively uniform structure, which is typical for normal tissue.
- There are no obvious signs of abnormal masses, lesions, or distortions in the immediate vicinity of the spine or surrounding structures.
- Comparison to Normal Anatomy: Normal anatomy in this region would typically show well-defined structures without significant abnormalities.
- The image does not show any clear signs of pathology such as tumors, inflammation, or other anomalies.
- Conclusion: Based on the uniform appearance and lack of visible abnormalities, the image does not show any variation from normal.

Therefore, the most appropriate answer is: D.

InfiMed-RL-3B Answer:

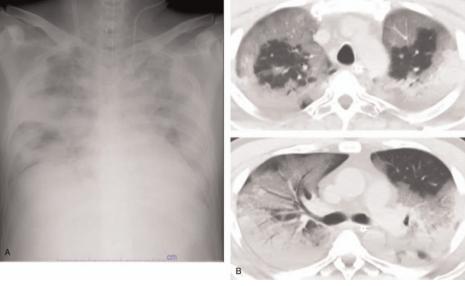
The image appears to be a medical scan, likely an MRI or CT scan, showing a cross-sectional view of the body. The presence of contrast and the structure suggest it could be a diagnostic image. However, without specific medical training or additional context, it's challenging to definitively identify any abnormalities.

The appearance could be normal, but the presence of contrast and the complexity of the image might indicate the presence of some structural or physiological changes.

Figure 16: Case study on Qwen2.5-VL-3B, InfiMed-Series models in medical VQA. **Red** denotes errors or irrelevant content, whereas **Green** denotes correct or important information. Qwen2.5-VL-3B's response is incorrect, while InfiMed-SFT-3B correctly answers after a detailed analysis and reflection, and InfiMed-RL-3B provides the most direct and accurate correct answer.

1350
1351
1352
1353
1354
1355
1356
1357
1358
1359
1360
1361
1362
1363
1364
1365
1366
1367
1368
1369
1370
1371
1372
1373

1374



Qwen2.5-VL-3B Answer:

 The chest X-ray in panel (A) shows a normal-sized heart, which is indicated by the absence of any signs of enlargement such as cardiomegaly. **The answer is B.**

InfiMed-RL-3B Answer:

 The chest X-ray in panel (A) shows a clear view of the heart. The cardiac silhouette appears to be within the normal range for the patient's thoracic cavity, suggesting a normal cardiac size. **Enlarged or abnormal cardiac size would typically be indicated by a heart that appears disproportionately large relative to the thoracic cavity.** **\boxed{B}**

Question:
 What is the cardiac size on the chest X-ray in panel (A)?
 A: Enlarged B: Normal
 C: Small D: Cannot be determined

Ground Truth:
 B: Normal

InfiMed-SFT-3B Answer:
 To determine the cardiac size on the chest X-ray in panel (A), let's analyze the image step by step:
 1. **Chest X-ray (Panel A):** The image shows a frontal view of the chest. The heart shadow is visible, and its size relative to the rest of the thoracic cavity can be assessed.
 2. **Comparison with Normal Cardiac Size:** In a normal chest X-ray, the heart should not appear disproportionately large or small compared to the thoracic cavity.
 3. **Analysis:** The heart appears to be of normal size in relation to the thoracic cavity.
 Given this analysis, the cardiac size on the chest X-ray in panel (A) appears to be normal.
\boxed{B}

Figure 17: Case study on Qwen2.5-VL-3B, InfiMed-Series models in medical VQA. **Red** denotes errors or irrelevant content, whereas **Green** denotes correct or important information. All three models correctly identify the cardiac size as normal, with Qwen2.5-VL-3B providing a concise answer, InfiMed-SFT-3B offering a detailed, step-by-step analysis, and InfiMed-RL-3B giving a direct and well-reasoned response.

A.7 LIMITATIONS AND FUTURE WORKS

Although our InfiMed-Series models achieve state-of-the-art (SOTA) performance among MLLMs with a similar number of parameters, they even outperform some MLLMs with larger parameter counts. However, it is undeniable that open-source medical multimodal data often exhibit low quality, including poor image resolution, non-uniform distribution of modalities, and errors introduced during model synthesis. Consequently, some of the results may lack full confidence, and the models' performance on more complex medical downstream tasks remains to be thoroughly explored. Moreover, how to develop reasoning steps that can be more efficient and accurate in the medical field is a critical issue that needs further study.

Additionally, we want to clarify that our training datasets are sourced exclusively from publicly available datasets, ensuring that no private or sensitive data is involved.

1392
1393
1394
1395
1396
1397
1398
1399
1400
1401
1402
1403