

# MedCoT: Medical Chain of Thought via Hierarchical Expert

Jiaxiang Liu<sup>1</sup> Yuan Wang<sup>1</sup> Jiawei Du<sup>2,3</sup> Joey Tianyi Zhou<sup>2,3</sup> Zuozhu Liu<sup>\*,1</sup>

<sup>1</sup> ZJU-Angelalign R&D Center for Intelligence Healthcare, Zhejiang University, China

<sup>2</sup> Centre for Frontier AI Research (CFAR), Agency for Science, Technology and Research (A\*STAR), Singapore

<sup>3</sup> Institute of High Performance Computing (IHPC), Agency for Science, Technology and Research (A\*STAR), Singapore  
{jiaxiang.21, zuozhuliu}@intl.zju.edu.cn

## Abstract

Artificial intelligence has advanced in Medical Visual Question Answering (Med-VQA), but prevalent research tends to focus on the accuracy of the answers, often overlooking the reasoning paths and interpretability, which are crucial in clinical settings. Besides, current Med-VQA algorithms, typically reliant on singular models, lack the robustness which usually require collaborative expert evaluation. To address these shortcomings, this paper presents MedCoT, a novel hierarchical expert verification reasoning chain method designed to enhance interpretability and accuracy in biomedical imaging inquiries. MedCoT is predicated on two principles: *The necessity for explicit reasoning paths in Med-VQA and the requirement for multi-expert review to formulate accurate conclusions.* The methodology involves an Initial Specialist proposing diagnostic rationales, followed by a Follow-up Specialist who validates these rationales, and finally, a consensus is reached through a vote among a sparse Mixture of Experts within the locally deployed Diagnostic Specialist, which then provides the definitive diagnosis. Experimental evaluations on four standard Med-VQA datasets demonstrate that MedCoT surpasses existing state-of-the-art approaches, providing significant improvements in performance and interpretability. Code is released at <https://github.com/JXLiu-AI/MedCoT>.

## 1 Introduction

Medical Visual Question Answering (Med-VQA) has recently gained significant attention (Chen et al., 2022; Gong et al., 2021; Ren and Zhou, 2020; Khare et al., 2021). As a new exploration in the medical domain, Med-VQA aims to answer medical questions in natural language based on input medical images. An effective Med-VQA system

\* Corresponding author.

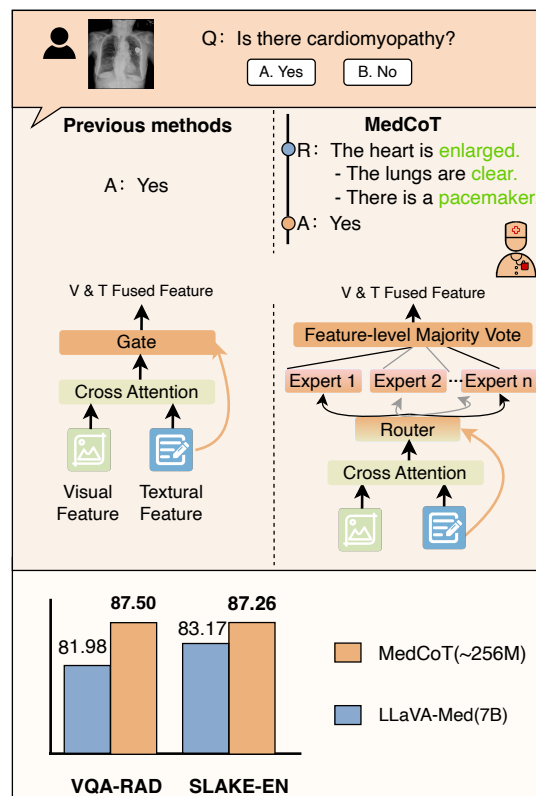


Figure 1: The upper figure shows a comparison of the outputs from the previous Med-VQA method and MedCoT, as well as the previous techniques in MMCoT (Zhang et al., 2023b) versus Sparse MoE in MedCoT. The lower figure demonstrates that MedCoT, with a model size of 256M parameters, outperforms the 7B parameter LLaVA-Med by 5.52% and 4.09% (Accuracy) on the VQA-RAD and SLAKE-EN datasets.

can assist clinicians in interpreting medical images, thereby ensuring and accelerating the diagnostic process. For patients, automated Med-VQA services can greatly satisfy the demand for personalized health consultations (Liu et al., 2023a).

In the field of Med-VQA, numerous attempts have been made using deep learning technologies (Tiong et al., 2022a; Banerjee et al., 2021; Changpinyo et al., 2022; Liu et al., 2023b; Gai et al.,

2024). For instance, [Nguyen et al. \(2019\)](#) utilized Bilinear Attention Networks (BAN) ([Kim et al., 2018](#)) and enhanced them for Med-VQA by incorporating a Mixed Enhanced Visual Feature (MEVF) setup consisting of pre-trained meta-learning modules and Convolutional Denoising Autoencoders (CDAE). Building on this, [Zhan et al. \(2020\)](#) designed a conditional reasoning framework to boost the inference capabilities of Med-VQA models. However, these approaches often underperform in many practical scenarios, primarily due to poor capabilities in extracting and integrating features from a limited number of medical images and text data ([Eslami et al., 2021](#); [Song et al., 2022](#); [Wang et al., 2022](#)). [Eslami et al. \(2021\)](#) introduced the CLIP architecture into the framework by deploying it as the visual encoder within MEVF ([Nguyen et al., 2019](#)), pre-trained on the multimodal medical dataset ROCO ([Pelka et al., 2018](#)). Their experiments demonstrated significant improvements with the CLIP. [Liu et al. \(2023a\)](#) developed VQA-Adapter, which uses a lightweight adapter and label smoothing to efficiently fine-tune the CLIP model for Med-VQA, thus reducing computational costs and mitigating overfitting. [Li et al. \(2024\)](#) proposed LLaVA-Med, which utilizes GPT-4 and a novel curriculum learning approach to efficiently train LLaVA on biomedical images, significantly enhancing Med-VQA capabilities.

However, previous Med-VQA approaches typically focused on the accuracy of the answers ([Nguyen et al., 2019](#); [Liu et al., 2023a](#); [Zhan et al., 2020](#)), where most MedVQA responses consist of a simplistic answer lacking detailed explanations or rationale, unlike in real-world scenarios where doctors not only provide answers but also explain their reasoning, professional considerations, and potential contradictions to derive a more comprehensive diagnostic insight. Besides, real-world diagnostics often rely on the combined experience of multiple doctors, as a single doctor's diagnosis may be biased by personal experience and may not be sufficiently accurate. In the multimodal Chain of Thought (CoT), answering VQA questions involves providing an answer as well as a corresponding reasoning path (rationale). The generation of this rationale helps to improve the accuracy of the language model. Inspired by real-world practices and multimodal CoT [Zhang et al. \(2023b\)](#); [Zheng et al. \(2023\)](#), integrating this paradigm into Med-VQA can enhance both the accuracy and interpretability of responses. However, implementing it

faces several challenges: (1) Previous CoT methods required manual annotation of fundamental rationales, which is time-consuming, costly, and challenging to ensure consistency and completeness ([Zhang et al., 2023b](#); [Zheng et al., 2023](#)). (2) Reliance on a single expert model can lead to misleading conclusions. (3) Multimodal CoT has limited depth in understanding the intents of images and texts, which can restrict its effectiveness in medical contexts ([Zhang et al., 2023b](#)).

To address the aforementioned issues, we introduce MedCoT, a hierarchical expert-verified model for Med-VQA. Firstly, the Initial Specialist proposes preliminary diagnostic rationale based on the medical visual and text query. The Follow-up Specialist then reviews these rationales, categorizing them as valid or invalid; valid rationales are retained, while invalid ones are reassessed. Finally, the locally implemented Diagnostic Specialist, consisting of a sparse Mixture of Experts (MoE) model functioning as a multimodal language model, casts votes to deliver the definitive diagnosis. Leveraging a hierarchy of expertise, MedCoT consistently outperforms state-of-the-art (SoTA) Med-VQA methods across four extensive datasets, demonstrating impressive generalizability and interpretability, as shown in [Figure 1](#). Our study makes three significant contributions:

- We have conducted an in-depth analysis of the challenges and insights associated with generating rationales in multimodal CoT. Our findings highlight that single specialist often fails to provide clear verifications and are more prone to errors when addressing questions about specific organs.
- Inspired by real-world diagnostics, we developed the hierarchical expert-verified MedCoT, which does not require manually annotated rationales. This involves three tiers of expert verification: initial, follow-up, and diagnosis. MedCoT not only provides more accurate answers but also offers refined rationales.
- In the diagnosis stage, we designed a sparse MoE that includes majority voting. This framework's multiple specialized experts efficiently and accurately interpret the intents of medical images and texts, enabling the Diagnostic Specialist to provide precise responses.

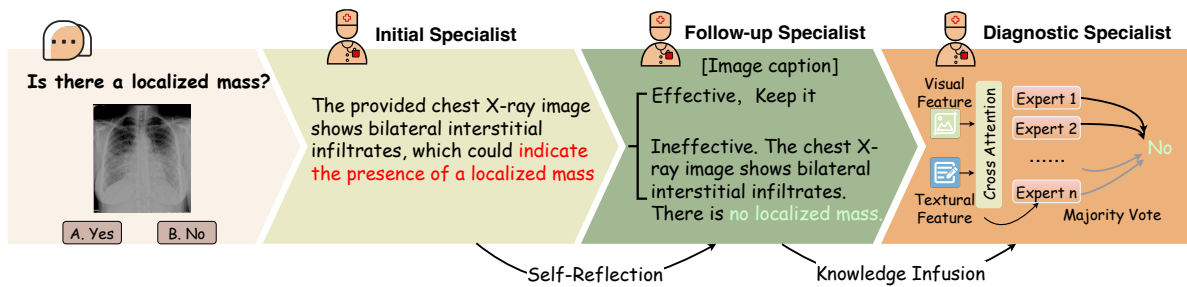


Figure 2: The MedCoT pipeline begins with an Initial Specialist receiving a medical question and image to generate a preliminary rationale. This rationale may have flaws (indicated in red), which are then reviewed by the Follow-up Specialist. If the rationale is deemed effective, it is retained; otherwise, it is reconsidered and a new rationale (indicated in green) is generated, along with an image caption. These elements are then integrated into the Diagnostic Specialist. Informed by all contexts, the Diagnostic Specialist, a multimodal language model with a designed sparse MoE structure, delivers the final diagnostic outcome (answer).

## 2 Related Work

### 2.1 Med-VQA

VQA is a multimodal task in computer vision and natural language processing, aimed at responding to queries about images in natural language (Ben Abacha et al., 2019; He et al., 2020; Ren and Zhou, 2020). It involves feature extraction, fusion, and inference to comprehend multimodal intents and manage feature processing. Med-VQA extends VQA into the medical domain, where robust medical knowledge is crucial for answering domain-specific questions (Liu et al., 2023a), thus complicating feature extraction. Innovations such as Nguyen et al.’s MEVF leverage unsupervised CDAE and meta-learning to initialize weights specifically for Med-VQA (Nguyen et al., 2019). Zhan et al. built upon this by developing a conditional reasoning framework to handle different types of questions (Zhan et al., 2020), while Eslami et al. successfully implemented the CLIP model as a visual encoder, proving its effectiveness in this context (Eslami et al., 2023). LLaVA-Med utilizes GPT-4 and a novel curriculum learning approach for training on biomedical images Li et al. (2024), significantly enhancing Med-VQA capabilities. While capable of interactive dialogue, its responses do not focus on the reasoning paths leading to the answers. MedCoT differs from the aforementioned methods by not only providing precise answers but also offering reasoning paths (rationale). Moreover, its validity is confirmed through Hierarchical Expert verification, aligning more closely with real-world medical scenarios.

### 2.2 Multimodal CoT

CoT reasoning with Large Language Models (LLMs) has shown success in natural language pro-

cessing. Multimodal CoT combines visual information with traditional textual CoT, integrating comprehensive data to perform reasoning tasks (Zhang et al., 2023b; Zheng et al., 2023). Groundbreaking works in multimodal CoT (Zheng et al., 2023; Zhang et al., 2023b; Lu et al., 2022, 2023; Zhang et al., 2023a) are first examined on the ScienceQA dataset. ScienceQA includes multimodal scientific questions along with annotated rationales (Lu et al., 2022). MM-CoT developed a two-stage framework based on ScienceQA that trains models to generate rationales from annotations, which are then used to form final answers (Lu et al., 2022). With the increasing integration of open-world knowledge in LLMs, research is focusing on equipping these models with visual modalities to tackle complex visual and multimodal challenges. For instance, DD-CoT (Zheng et al., 2023), introduces role-specific Chains of Thought that decompose questions into subproblems and use LLMs to recombine principles, enhancing accuracy and addressing language illusions in multimodal contexts. Inspired by these advancements, we aim to adapt multimodal CoT reasoning to the medical field, aiming to improve the explainability and accuracy of Med-VQA.

### 2.3 MoE

MoE optimizes learning and prediction by combining multiple expert networks and using a gating network to determine which experts are activated based on the given input (Zhang et al., 2024; Fedus et al., 2022b). Sparse MoE, a variant of the MoE model, activates only a few experts during each prediction, thus efficiently utilizing computational resources and enhancing scalability (Shazeer et al., 2016). Sparse MoE models have been independently explored within the context of conditional

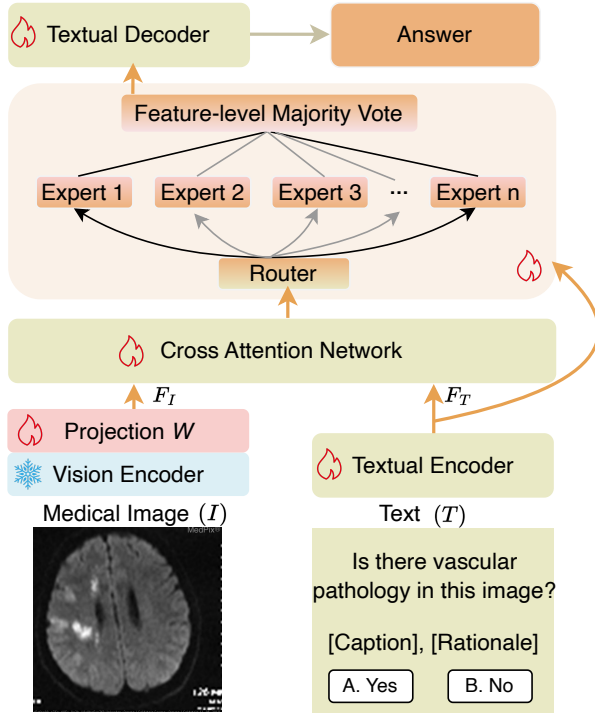


Figure 3: Diagnostic Specialist Pipeline. After passing through a visual encoder, medical images yield visual features. Contextual textual information—including captions, rationales, and options—is processed by a text encoder to obtain textual features. These are then subjected to cross-attention for feature integration, producing combined features. These integrated features, along with textual features, are input into a Sparse MoE structure. Here, multiple specialized experts thoroughly understand the intents of both the image and text. The insights are then fed into a textual decoder, which decodes the information to produce the final answer.

computation in both computer vision and natural language processing domains [Jacobs et al. \(1991\)](#); [Fedus et al. \(2022a\)](#). Conditional computation aims to increase the number of model parameters without proportionally increasing computational costs. This is achieved by selectively activating only the relevant parts of the model based on input-specific factors ([Shazeer et al., 2016](#)). Sparse MoE models employ a learned gating mechanism that activates only a subset of experts, specifically  $k$  out of  $N$  experts, for a given input. This allows for the selection of either all experts or just a sparse mix, optimizing resource usage [Lepikhin et al. \(2020\)](#).

### 3 Methodology

#### 3.1 Preliminaries

Throughout this paper, we model the Med-VQA task within a multimodal CoT framework as follows: The framework takes an image  $I$  and a ques-

tion  $Q$  as inputs, and outputs a reasoning rationale  $R$ . This rationale  $R$  is subsequently used to generate an answer  $A$ . This paradigm ensures that the process is transparent, providing a traceable path from input to conclusion, which is essential for both validating the results and improving user trust in the framework’s diagnostic capabilities. We can model the Med-VQA task within a multimodal CoT as follows:

$$\min_{f,g} \mathbb{E}_{(I,Q,A^*) \sim \text{Data}} [L(g(f(I,Q), I, Q), A^*)]. \quad (1)$$

$f$  is responsible for generating a rational and helpful reasoning rationale  $R$  (Initial and Follow-up Specialists), while  $g$  uses this rationale to generate the final answer  $A$  (Diagnostic Specialist). The rationale  $R$  is derived from the Initial Specialist assessments and self-reflection by the Follow-up Specialist. The final answer  $A$  is determined by a Diagnostic Specialist through a loss function  $L$ , which measures the discrepancy between the predicted answer  $A$  and the true answer  $A^*$ .

#### 3.2 Initial Specialist

In the initial diagnosis phase, we cue the LLMs to act as the primary rationale Diagnostic Specialist. We prompt the LLMs with the instruction: "Please proceed with a step-by-step analysis and provide a rationale" ( $prompt_i$ ). This is done to guide the LLMs in performing a detailed, step-by-step reasoning process. The textual rationale obtained from this is represented as  $R_i = LLM_s(T, I, prompt_i)$ , where  $T$  and  $I$  denote the text and image inputs, respectively.  $T$  includes textual context such as the question  $Q$  and options.  $prompt_i$  is the specific prompting strategy used to elicit the rationale. For further technical details about the prompt, please refer to the appendix [H](#).

For instance, as shown in [Figure 2](#), for the question "Is there a localized mass?", we obtain a highly interpretable rationale (for the final diagnostic outcome): "The provided chest X-ray image shows bilateral interstitial infiltrates, which could indicate the presence of a localized mass".

#### 3.3 Follow-up Specialist

In the follow-up diagnosis phase, we instruct LLMs to conduct self-reflection reasoning and test within the problem’s context to identify effective rationales, retain them, and reconstruct ineffective ones to generate accurate rationales. Specifically, we prompt the LLMs with: "Please judge whether this

rationale is effectively valid for the question and image. If it is effective..., If the existing rationale is Ineffective..." (prompt  $\hat{f}$ ). For the complete prompt, please refer to the appendix H.2. We can define the Self-Reflection reasoning of the Follow-up Specialist using the following formula:

$$R_{\hat{f}} = \begin{cases} R_i & \text{if } R_i = \text{Effective} \\ \text{LLMs}(T, I, \text{prompt}_{\hat{f}}) & \text{if } R_i = \text{Ineffective}, \end{cases} \quad (2)$$

where  $R_{\hat{f}}$  is Follow-up Specialist rationale. This process helps us obtain the textual rationale needed for the diagnostic analysis, as shown in Figure 2.

To infuse the Diagnostic Specialist with more knowledge and bridge the gap between image and text, we utilize the Follow-up Specialist to generate image captions. This process helps to reduce the modality gap, effectively channeling this knowledge into the Diagnostic Specialist. For detailed caption prompts, please refer to the appendix H.3.

### 3.4 Diagnostic Specialist

We employ the designed model based on multimodal T5 combined with sparse MoE to serve as the Diagnostic Specialist, as shown in Figure 3. The Diagnostic Specialist receives enriched textual context and medical imaging information to generate the final diagnostic outcome.

#### 3.4.1 Multimodal T5

Figure 3 shows the structure of multimodal T5, including the *TextualEncoder*, *VisualEncoder*, *Cross-Attention Network*, sparse MoE, and the *TextualDecoder*. Here are the network details:

*TextualEncoder* transforms natural language input  $T$  into the textual feature space  $F_T \in \mathbb{R}^{n \times d}$ , and *VisualEncoder* converts the input image  $I$  into visual features  $F_I \in \mathbb{R}^{m \times d}$ . Here,  $n$  signifies the length of the input language text,  $d$  the dimensionality of hidden features, and  $m$  the count of image patches. Upon obtaining the textual representation  $F_T$  and visual representation  $F_I$ , our model leverages the *Cross-Attention Network* for modality interaction. This network computes the attention-guided visual feature  $H_V^{\text{att}} \in \mathbb{R}^{n \times d}$ , which selectively captures relevant visual features in response to the textual query, as delineated in the operation:

$$H_V^{\text{att}} = \text{Softmax} \left( \frac{QK^\top}{\sqrt{d}} \right) V, \quad (3)$$

where  $Q$ ,  $K$ ,  $V$  correspond to the query, key, and value, derived from  $F_T$ ,  $F_I$ ,  $F_I$ , respectively.

Once the attention-guided visual feature  $H_V^{\text{att}}$  and the textual representation  $F_T$  are obtained, we construct the MoE to dynamically amalgamate them, resulting in  $F_F = \text{MoE}(H_V^{\text{att}}, F_T)$ . Details of the MoE are provided in the following section.  $F_F$  is input into the *TextualDecoder* to generate answer  $A = \text{TextualDecoder}(F_F)$ , as shown in Figure 3.

In the training, refinements enable predicted answers ( $A$ ) to more accurately approximate label answers. Specifically, The model  $f$  with input maximizes the likelihood of the correct sequence  $Y = A$ . The loss function  $L$ , which is the negative log-likelihood over all tokens, is given by:  $L = -\sum_{n=1}^N \log p(Y_n|X, Y_1^{n-1})$ , where  $N$  is the number of tokens, and  $p(Y_n|X, Y_1^{n-1})$  is the probability of predicting the correct  $n$ -th token in  $Y$ .

#### 3.4.2 MoE

In the multimodal CoT, a crucial step is understanding the intent of both the image and the text and responding accordingly. However, previous methods primarily utilized gates for integration, where the gate function  $\lambda = \text{Sigmoid}(W_l F_T + W_v H_V^{\text{att}})$  weights the importance of the image relative to the source text, with  $W_l$  and  $W_v$  as learnable parameters (see Appendix B.1) (Zhang et al., 2023b; Zheng et al., 2023). Which, according to our experiments, shows that the gate is insufficient (subsection 4.3). Therefore, MedCoT proposes constructing a MoE for the integration process.

The Sparse MoE implements a top-k sparse mixture of experts (Fedus et al., 2022b), leveraging multiple Sparse Experts to specialize in processing complex Med-VQA data. This module dynamically selects the top-k experts for each input based on gating scores, as shown in Figure 3.

After obtaining the outputs from the experts, we use Feature-level Majority Vote to aggregate their outputs. The weight of each expert is calculated using the following formula:

$$W_i = \text{softmax}(V^{\text{top } k})_i = \frac{e^{V_i^{\text{top } k}}}{\sum_{j=1}^k e^{V_j^{\text{top } k}}}, \quad (4)$$

where  $W_i$  is the weight of the  $i$ -th selected expert, and  $V_i^{\text{top } k}$  is the score of the  $i$ -th selected expert. For each feature  $F_f$ , the final result of Feature-level Majority Vote is calculated by weighted averaging the outputs of all selected experts:

$$E_{F_f} = \sum_{i=1}^k W_i \cdot E_{i, F_f}, \quad (5)$$

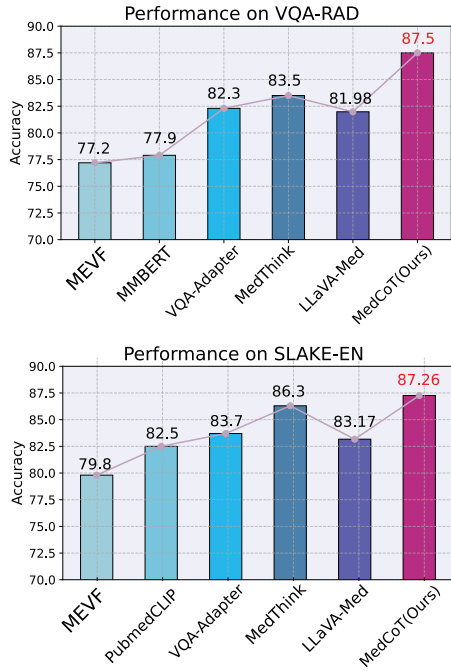


Figure 4: MedCoT is compared with various SoTA methods on closed questions on the VQA-RAD and SLAKE-EN datasets. MedCoT not only achieves SoTA accuracy in answers but also provides reasoning paths (rationale). The metric used is Accuracy (%).

where  $E_{F_f}$  is the value of the final result for feature  $F_f$ , and  $E_{i,F_f}$  is the output of the  $i$ -th selected expert for feature  $F_f$ . Then,  $\lambda = \text{Sigmoid}(E_{F_f})$ . Finally, this results in  $F_f$  are as follows:

$$F_f = (1 - \lambda) \cdot F_T + \lambda \cdot H_V^{\text{att}}. \quad (6)$$

The sparse MoE network allows each selected expert to handle data they specialize in, as demonstrated in Figure 6, which shows experts proficient in addressing head-related issues.

## 4 Experiments

### 4.1 Experimental Setting

In MedCoT framework, the encoder and decoder from Flan-T5 (Khashabi et al., 2020; Raffel et al., 2020) are integrated as TextualEncoder(·) and TextualDecoder(·), respectively. Additionally, DETR (Carion et al., 2020) is employed as VisualEncoder(·). Our Diagnostic Specialist model was trained 100 epochs with a learning rate of  $8e - 5$  and a batch size of 8. To demonstrate the effects of MedCoT, four benchmark datasets are used for validation in the medical VQA domain: VQA-RAD (Lau et al., 2018), SLAKE-EN (Liu et al., 2021), Med-VQA-2019 (Abacha et al., 2019b), and PathVQA (He et al., 2020), with detailed statistics

provided in Appendix Table 4. All experiments were conducted using PyTorch (Paszke et al., 2019) and HuggingFace (Wolf et al., 2020), implemented on 4 NVIDIA GEFORCE RTX 3090 GPUs. Accuracy is utilized as the evaluation metric. For LLMs, Gemini Pro 1.5 version is used for our Initial Specialist and Follow-up Specialist. The more experimental details can be found in the Appendix B.

### 4.2 Main Results

We evaluate the performance of MedCoT on the VQA-RAD and SLAKE-EN datasets, benchmarking them against established models like MEVF (Nguyen et al., 2019), MMBERT (Tiong et al., 2022b), PubMedCLIP (Eslami et al., 2023), VQA-Adapter (Liu et al., 2023a), MedThink (Gai et al., 2024), LLaVA-Med (Li et al., 2024) (Table 2).

Our performance evaluation is divided into two parts, focusing separately on closed-end and open-end questions. Closed-end questions, structured as multiple-choice questions with a single correct answer, are assessed using accuracy as the performance metric, as shown in Figure 4. In facing closed-end questions, MedCoT surpasses a range of SoTA methods on the VQA-RAD and SLAKE-EN datasets. Notably, MedCoT achieved improvements of 27.21% and 14.66% over Gemini on the two datasets, demonstrating the unreliability of a single model. Besides, MedCoT, with a fine-tuning size of approximately 256M parameters, outperforms the 7B parameter LLaVA-Med (trained on extensive medical data), exceeding it by 5.52% and 4.09% on two datasets, respectively. Moreover, compared to previous methods, MedCoT clearly displays the reasoning paths (rationale), as illustrated in Figure 2. More comparative method results can be seen in Table 2.

In contrast, open-end questions allow for a range of answers due to their inherent nature. The answers generated by MedCoT are difficult to match precisely against the dataset. Therefore, we employ text generation metrics such as Rouge and BLEU to evaluate MedCoT’s performance. We conducted experiments on the open-end VQA-RAD and SLAKE-EN, with results shown in the Appendix Table 3. MedCoT demonstrated higher Rouge and BLEU scores on the VQA-RAD and SLAKE-EN dataset, surpassing MedThink (Gai et al., 2024). Besides, MedCoT also showed higher scores on the SLAKE-EN.

Additionally, we evaluated MedCoT’s performance on the Med-VQA-2019 and PathVQA

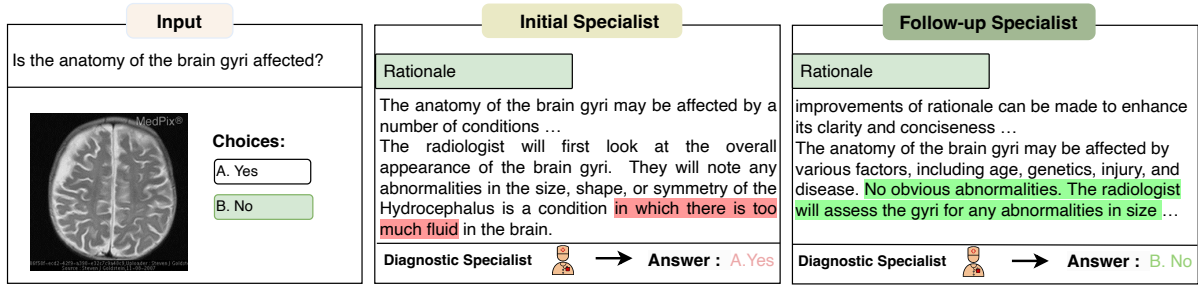


Figure 5: The MedCoT pipeline begins with an Initial Specialist receiving a medical question and image to generate a preliminary rationale. This rationale may have flaws (indicated in red), which are then reviewed by the Follow-up Specialist. If the rationale is deemed effective, it is retained; otherwise, it is reconsidered and a new rationale (indicated in green) is generated, along with an image caption. These elements are then integrated into the Diagnostic Specialist. Informed by all context, the Diagnostic Specialist, a multimodal language model with a designed sparse MoE structure, delivers the final diagnostic outcome (answer).




Figure 6: The Diagnostic Specialist’s sparse MoE shows varying accuracy levels for different organ-related questions in VQA-RAD. ‘ABD’ represents abdominal-related questions, ‘Head’ refers to head-related questions, and ‘Chest’ refers to chest-related questions. It can be observed that head-related questions saw an improvement of nearly 10%. We visualized the weights of the experts (right figure). Notably, in the top 2 expert selections, the model chose Expert 0 and Expert 5 to understand the intents of the "head" image and text.

Table 1: Ablation Study on MedCoT

Follow-up	MoE	VQA-RAD	SLAKE-EN
		77.57	83.17
✓		82.72	86.05
	✓	80.88	83.65
✓	✓	<b>87.50</b>	<b>87.26</b>

datasets, as shown in Table 2. The results indicate that MedCoT consistently achieves SoTA results compared to the majority of SoTA methods.

### 4.3 Ablation Study

**Effects of Follow-up Specialist** To validate the effectiveness of the Follow-up Specialist, we compared the results of experiments involving only the initial and diagnostic specialists with those from the complete MedCoT. As shown in Table 1, across

two medical datasets, there is a significant performance loss when the Follow-up Specialist is removed. For instance, on the VQA-RAD dataset, performance dropped from 87.50% to 80.88%, a decrease of 6.62%. This demonstrates the effectiveness of the Follow-up Specialist.

Besides, we conducted experiments involving only the initial and diagnostic specialists, bypassing the self-reflection of the Follow-up Specialist. In all cases involving varying numbers of experts, the results without the self-reflection were consistently lower than those with rationales refined by the Follow-up Specialist’s reflection, and even lower than those from a Diagnostic Specialist that had undergone self-reflection but was lacking the MoE component, as shown in Figure 7. This underscores the importance of the self-reflection provided by the Follow-up Specialist. Additionally, we conducted zero-shot experiments using both the initial and Follow-up Specialist. As shown in the appendix Figure 12, these results further confirm the effectiveness of the Follow-up Specialist.

**Effects of MoE** To validate the effectiveness of the MoE, we compared the performance with and without the MoE. As shown in Table 1, there is a significant performance drop across all datasets without MoE. For instance, in the VQA-RAD, the performance decreased from 87.50% to 82.72%, a loss of 4.78%. This indicates that MoE plays a crucial role in Diagnostic Specialist. As can also be seen from Figure 7, lacking MoE, in most expert number scenarios, the performance is weaker compared to MedCoT equipped with Sparse MoE.

Additionally, we conducted experiments for each organ-related question category within the VQA-RAD and SLAKE-EN, as shown in Figure 6 and




Figure 7: The expert number grid search on two datasets. The blue line represents the results from training with Initial Specialist rationales and grid search of expert numbers in the Diagnostic Specialist. The purple line represents results from using the Follow-up Specialist rationales and grid searching expert numbers. The gray line represents the results of the Diagnostic Specialist using Follow-up Specialist rationales, conducted without the sparse MoE.

**Figure 9.** It is evident that in the majority of organ-related questions, methods employing MoE outperform those using the gating mechanism. Notably, the Gate mechanism, resembling as a single-expert system, tends to falter with head-related questions, where it performs the worst. For such questions in the VQA-RAD, methods using MoE exceeded those with gates by 10%, further emphasizing MoE’s effectiveness. We visualized the weights of MoE, as shown in the **Figure 6** (right figure), revealing that Experts 0 and 5 primarily handle head-related issues. This demonstrates that these two experts dynamically process and understand the intents of medical images and texts more effectively than the gating. Similar results can also be observed in the experiments conducted on the SLAKE-EN, as shown in **Figure 9**.

**Grid Search** We conducted a parameter search experiment for the hyperparameters in the sparse MoE, such as the number of experts and the  $k$  value. The results are shown in **Figure 6**, **Figure 9** and **Figure 11**. The experiment revealed that the optimal number of experts varies for different datasets. Specifically, the best number of experts for VQA-RAD, SLAKE-EN, Med-2019 and PathVQA are 6, 10, 5, and 5, respectively. Regarding the  $k$  value, the optimal value for all datasets was consistently 2, as illustrated in **Figure 10**.

#### 4.4 Discussion

**Figure 2** and **Figure 5** illustrate cases where the Initial Specialist provides a rationale, the Follow-up Specialist makes corrections, and the Diagnostic Specialist delivers the final, accurate diagnosis. For instance, in **Figure 5**, the Initial Specialist, influenced by the illusions of the LLMs, mistakenly

observes non-existent brain fluid and diagnoses the brain as being affected by gyri. However, after the self-reflection by the Follow-up Specialist, it is clarified that no clear fluid was observed. Ultimately, the Diagnostic Specialist, using the rationale from the Follow-up Specialist and considering the full context, arrives at the correct diagnosis.

**Figure 8** provides an example where the limitations of LLMs affect the ability to accurately diagnose certain cases. The question posed is whether there is pneumomediastinum. The Initial Specialist, based on observations, affirms its presence, and the Follow-up Specialist concurs, leading to a unanimous agreement. However, due to the limitations of the LLMs, these rationales are incorrect, ultimately leading to an erroneous answer.

## 5 Conclusion

In this paper, we propose an effective hierarchical expert reasoning chain method for Med-VQA, named MedCoT. This method is based on two insights: 1) Med-VQA should have a clear reasoning path; 2) Med-VQA scenarios should be reviewed by multiple experts to arrive at a conclusion. Specifically, the process involves initial experts providing preliminary diagnostic rationales based on medical visual questions. Follow-up experts then review these rationales for validity, retaining the effective ones and reassessing the ineffective ones. Finally, a locally deployed Diagnostic Specialist, consisting of a sparse MoE that conducts a vote, then provides the definitive diagnosis. Experimental results on multiple Med-VQA datasets show that MedCoT outperforms existing SoTA techniques, significantly surpasses recent methods, and demonstrates excellent interpretability for final diagnosis.



## Limitation

A limitation is that the performance of MedCoT is influenced by the hallucinations of the LLMs used by the Initial and Follow-up Specialist. Although self-reflection and Hierarchical Expert design can mitigate some issues with LLMs' hallucinations, it must be acknowledged that the problem is not completely resolved. As shown in [Figure 8](#), MedCoT is still susceptible to hallucination risks. Researching methods to suppress hallucinations is a potential topic for further study. In this work, the Gemini-Pro model was employed. If Med-Gemini becomes available, MedCoT could be further enhanced. Moreover, MedCoT could inspire future paradigms that integrate proprietary commercial LLMs with local models. By utilizing desensitized information to prompt the extensive knowledge and reasoning capabilities of LLMs, the generated rationales could be combined with local models for further diagnostic analysis, enhancing both interpretability and accuracy.

Another limitation is that compared to single-model methods, MedCoT may be more time-consuming. However, the hierarchical expert approach aligns more closely with real-world medical diagnostics and provides clear diagnostic pathways as well as more accurate answers, making the additional time worthwhile.

## Acknowledgements

This work is supported by the National Natural Science Foundation of China (Grant No. 62106222), the Natural Science Foundation of Zhejiang Province, China (Grant No. LZ23F020008), and the Zhejiang University-Angelalign Inc. R&D Center for Intelligent Healthcare. This work is also supported by Jiawei Du's A\*STAR Career Development Fund (CDF) C233312004.

## References

Asma Ben Abacha, Sadid A. Hasan, Vivek Datla, Joey Liu, Dina Demner-Fushman, and Henning Müller. 2019a. [Vqa-med: Overview of the medical visual question answering task at imageclef 2019](#). In *Conference and Labs of the Evaluation Forum*.

Asma Ben Abacha, Sadid A Hasan, Vivek V Datla, Joey Liu, Dina Demner-Fushman, and Henning Müller. 2019b. Vqa-med: Overview of the medical visual question answering task at imageclef 2019. *CLEF (working notes)*, 2(6).

Pratyay Banerjee, Tejas Gokhale, Yezhou Yang, and Chitta Baral. 2021. Weaqa: Weak supervision via captions for visual question answering. *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*.

Yakoub Bazi, Mohamad Mahmoud Al Rahhal, Laila Bashmal, and Mansour Zuair. 2023. Vision-language model for visual question answering in medical imagery. *Bioengineering*, 10(3):380.

Asma Ben Abacha, Sadid A Hasan, Vivek V Datla, Dina Demner-Fushman, and Henning Müller. 2019. Vqa-med: Overview of the medical visual question answering task at imageclef 2019. In *Proceedings of CLEF (Conference and Labs of the Evaluation Forum) 2019 Working Notes*. 9-12 September 2019.

Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. 2020. End-to-end object detection with transformers. In *European conference on computer vision*, pages 213–229. Springer.

Soravit Changpinyo, Doron Kukliansy, Idan Szpektor, Xi Chen, Nan Ding, and Radu Soricut. 2022. All you may need for vqa are image captions. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1947–1963.

Zhihong Chen, Guanbin Li, and Xiang Wan. 2022. Align, reason and learn: Enhancing medical vision-and-language pre-training with knowledge. In *Proceedings of the 30th ACM International Conference on Multimedia*, pages 5152–5161.

Sedigheh Eslami, Gerard de Melo, and Christoph Meinel. 2021. Does clip benefit visual question answering in the medical domain as much as it does in the general domain? *arXiv preprint arXiv:2112.13906*.

Sedigheh Eslami, Christoph Meinel, and Gerard de Melo. 2023. [PubMedCLIP: How much does CLIP benefit visual question answering in the medical domain?](#) In *Findings of the Association for Computational Linguistics: EACL 2023*, pages 1181–1193, Dubrovnik, Croatia. Association for Computational Linguistics.

William Fedus, Jeff Dean, and Barret Zoph. 2022a. A review of sparse expert models in deep learning. *arXiv preprint arXiv:2209.01667*.

William Fedus, Barret Zoph, and Noam Shazeer. 2022b. Switch transformers: Scaling to trillion parameter models with simple and efficient sparsity. *Journal of Machine Learning Research*, 23(120):1–39.

Xiaotang Gai, Chenyi Zhou, Jiayang Liu, Yang Feng, Jian Wu, and Zuozhu Liu. 2024. Medthink: Explaining medical visual question answering via multimodal decision-making rationale. *arXiv preprint arXiv:2404.12372*.

- Haifan Gong, Guanqi Chen, Sishuo Liu, Yizhou Yu, and Guanbin Li. 2021. Cross-modal self-attention with multi-task pre-training for medical visual question answering. In *Proceedings of the 2021 International Conference on Multimedia Retrieval*, pages 456–460.
- Xuehai He, Yichen Zhang, Luntian Mou, Eric Xing, and Pengtao Xie. 2020. [Pathvqa: 30000+ questions for medical visual question answering](#). Preprint, arXiv:2003.10286.
- Lei Huang, Weijiang Yu, Weitao Ma, Weihong Zhong, Zhangyin Feng, Haotian Wang, Qianglong Chen, Weihua Peng, Xiaocheng Feng, Bing Qin, et al. 2023. A survey on hallucination in large language models: Principles, taxonomy, challenges, and open questions. *arXiv preprint arXiv:2311.05232*.
- Xiaofei Huang and Hongfang Gong. 2023. A dual-attention learning network with word and sentence embedding for medical visual question answering. *IEEE Transactions on Medical Imaging*.
- Robert A Jacobs, Michael I Jordan, Steven J Nowlan, and Geoffrey E Hinton. 1991. Adaptive mixtures of local experts. *Neural computation*, 3(1):79–87.
- Yash Khare, Viraj Bagal, Minesh Mathew, Adithi Devi, U Deva Priyakumar, and CV Jawahar. 2021. Mmbert: multimodal bert pretraining for improved medical vqa. In *2021 IEEE 18th International Symposium on Biomedical Imaging (ISBI)*, pages 1033–1036. IEEE.
- Daniel Khashabi, Sewon Min, Tushar Khot, Ashish Sabharwal, Oyvind Tafjord, Peter Clark, and Hannaneh Hajishirzi. 2020. Unifiedqa: Crossing format boundaries with a single qa system. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 1896–1907.
- Jin-Hwa Kim, Jaehyun Jun, and Byoung-Tak Zhang. 2018. Bilinear attention networks. *Advances in neural information processing systems*, 31.
- Jason J Lau, Soumya Gayen, Asma Ben Abacha, and Dina Demner-Fushman. 2018. A dataset of clinically generated visual questions and answers about radiology images. *Scientific data*, 5(1):1–10.
- Dmitry Lepikhin, HyoukJoong Lee, Yuanzhong Xu, Dehao Chen, Orhan Firat, Yanping Huang, Maxim Krikun, Noam Shazeer, and Zhifeng Chen. 2020. Gshard: Scaling giant models with conditional computation and automatic sharding. In *International Conference on Learning Representations*.
- Chunyuan Li, Cliff Wong, Sheng Zhang, Naoto Usuyama, Haotian Liu, Jianwei Yang, Tristan Naumann, Hoifung Poon, and Jianfeng Gao. 2024. Llava-med: Training a large language-and-vision assistant for biomedicine in one day. *Advances in Neural Information Processing Systems*, 36.
- Bo Liu, Li-Ming Zhan, Li Xu, Lin Ma, Yan Yang, and Xiao-Ming Wu. 2021. Slake: A semantically-labeled knowledge-enhanced dataset for medical visual question answering. In *2021 IEEE 18th International Symposium on Biomedical Imaging (ISBI)*, pages 1650–1654. IEEE.
- Jiaxiang Liu, Tianxiang Hu, Yan Zhang, Yang Feng, Jin Hao, Junhui Lv, and Zuozhu Liu. 2023a. Parameter-efficient transfer learning for medical visual question answering. *IEEE Transactions on Emerging Topics in Computational Intelligence*.
- Jiaxiang Liu, Tianxiang Hu, Yan Zhang, Xiaotang Gai, YANG FENG, and Zuozhu Liu. 2023b. A chatgpt aided explainable framework for zero-shot medical image diagnosis. In *ICML 3rd Workshop on Interpretable Machine Learning in Healthcare (IMLH)*.
- Yunyi Liu, Zhanyu Wang, Dong Xu, and Luping Zhou. 2023c. Q2atransformer: Improving medical vqa via an answer querying decoder. In *International Conference on Information Processing in Medical Imaging*, pages 445–456. Springer.
- Pan Lu, Swaroop Mishra, Tanglin Xia, Liang Qiu, Kai-Wei Chang, Song-Chun Zhu, Oyvind Tafjord, Peter Clark, and Ashwin Kalyan. 2022. Learn to explain: Multimodal reasoning via thought chains for science question answering. *Advances in Neural Information Processing Systems*, 35:2507–2521.
- Pan Lu, Baolin Peng, Hao Cheng, Michel Galley, Kai-Wei Chang, Ying Nian Wu, Song-Chun Zhu, and Jianfeng Gao. 2023. Chameleon: Plug-and-play compositional reasoning with large language models. *arXiv preprint arXiv:2304.09842*.
- Binh D Nguyen, Thanh-Toan Do, Binh X Nguyen, Tuong Do, Erman Tjiputra, and Quang D Tran. 2019. Overcoming data limitation in medical visual question answering. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 522–530. Springer.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. 2019. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32.
- Obioma Pelka, Sven Koitka, Johannes Rückert, Felix Nensa, and Christoph M Friedrich. 2018. Radiology objects in context (roco): a multimodal image dataset. In *Intravascular Imaging and Computer Assisted Stenting and Large-Scale Annotation of Biomedical Data and Expert Label Synthesis: 7th Joint International Workshop, CVII-STENT 2018 and Third International Workshop, LABELS 2018, Held in Conjunction with MICCAI 2018, Granada, Spain, September 16, 2018, Proceedings 3*, pages 180–189. Springer.
- Zhangyang Qi, Ye Fang, Mengchen Zhang, Zeyi Sun, Tong Wu, Ziwei Liu, Dahua Lin, Jiaqi Wang, and Hengshuang Zhao. 2023. Gemini vs gpt-4v: A preliminary comparison and combination of vision-language models through qualitative cases. *arXiv preprint arXiv:2312.15011*.

- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *The Journal of Machine Learning Research*, 21(1):5485–5551.
- Fuji Ren and Yangyang Zhou. 2020. Cgmvqa: A new classification and generative model for medical visual question answering. *IEEE Access*, 8:50626–50636.
- Noam Shazeer, Azalia Mirhoseini, Krzysztof Maziarz, Andy Davis, Quoc Le, Geoffrey Hinton, and Jeff Dean. 2016. Outrageously large neural networks: The sparsely-gated mixture-of-experts layer. In *International Conference on Learning Representations*.
- Noah Shinn, Federico Cassano, Edward Berman, Ashwin Gopinath, Karthik Narasimhan, and Shunyu Yao. 2023. [Reflexion: Language agents with verbal reinforcement learning](#). Preprint, arXiv:2303.11366.
- Haoyu Song, Li Dong, Weinan Zhang, Ting Liu, and Furu Wei. 2022. Clip models are few-shot learners: Empirical studies on vqa and visual entailment. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 6088–6100.
- Anthony Meng Huat Tiong, Junnan Li, Boyang Li, Silvio Savarese, and Steven C.H. Hoi. 2022a. [Plug-and-play VQA: Zero-shot VQA by conjoining large pre-trained models with zero training](#). In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 951–967, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Anthony Meng Huat Tiong, Junnan Li, Boyang Li, Silvio Savarese, and Steven C.H. Hoi. 2022b. [Plug-and-play VQA: Zero-shot VQA by conjoining large pre-trained models with zero training](#). In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 951–967, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Tom Van Sonsbeek, Mohammad Mahdi Derakhshani, Ivona Najdenkoska, Cees GM Snoek, and Marcel Worring. 2023. Open-ended medical visual question answering through prefix tuning of language models. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 726–736. Springer.
- Zhecan Wang, Bin Xiao, Noel Codella, Jianwei Yang, Yen-Chun Chen, Luowei Zhou, Shih-Fu Chang, Xiyang Dai, Haoxuan You, and Lu Yuan. 2022. Clip-td: Clip targeted distillation for vision-language tasks. In *International Conference on Learning Representations*.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, et al. 2020. Transformers: State-of-the-art natural language processing. In *Proceedings of the 2020 conference on empirical methods in natural language processing: system demonstrations*, pages 38–45.
- Rongwu Xu, Zehan Qi, and Wei Xu. 2024. [Preemptive answer "attacks" on chain-of-thought reasoning](#). Preprint, arXiv:2405.20902.
- Li-Ming Zhan, Bo Liu, Lu Fan, Jiaxin Chen, and Xiaoming Wu. 2020. Medical visual question answering via conditional reasoning. In *Proceedings of the 28th ACM International Conference on Multimedia*, pages 2345–2354.
- Renrui Zhang, Jiaming Han, Aojun Zhou, Xiangfei Hu, Shilin Yan, Pan Lu, Hongsheng Li, Peng Gao, and Yu Qiao. 2023a. Llama-adapter: Efficient fine-tuning of language models with zero-init attention. *arXiv preprint arXiv:2303.16199*.
- Ruiyuan Zhang, Jiayang Liu, Zexi Li, Hao Dong, Jie Fu, and Chao Wu. 2024. Scalable geometric fracture assembly via co-creation space among assemblers. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 7269–7277.
- Zhuosheng Zhang, Aston Zhang, Mu Li, Hai Zhao, George Karypis, and Alex Smola. 2023b. Multi-modal chain-of-thought reasoning in language models. *arXiv preprint arXiv:2302.00923*.
- Ge Zheng, Bin Yang, Jiajin Tang, Hong-Yu Zhou, and Sibe Yang. 2023. Ddcot: Duty-distinct chain-of-thought prompting for multimodal reasoning in language models. In *Thirty-seventh Conference on Neural Information Processing Systems*.

## Appendix

In this section, we present additional implementation details, experiment results, and supplements. The content structure is outlined as follows:

- Section **A** - MedCoT on Four Datasets
- Section **B** - Method Details
  - Section **B.1** - Gate Mechanism
  - Section **B.2** - MedCoT Method Details
- Section **C** - Datasets
- Section **D** - The Effect of Initial Specialist
- Section **E** - Self-Reflection in Follow-up Specialist
  - Section **E.1** - Effectiveness of Self-Reflection in Follow-up Specialist
  - Section **E.2** - Error Analysis on Self-Reflection
- Section **F** - Computing Resource Costs
- Section **G** - Assessing the Impact of Different Prompts
- Section **H** - Prompt Template

### A MedCoT on Four Datasets

**Table 2** presents a comparison of methods performances across various datasets on closed-end questions. MedCoT, not only achieves superior performance but also demonstrates significant efficiency in model size compared to SoTA models. Despite being more lightweight, MedCoT consistently outperforms larger models across different datasets.

- **Efficiency in Size:** MedCoT has a model size of  $\sim 256\text{M}$ , which is significantly smaller than many other models such as Prefix T (1.5B), and LLaVA variants (7B). Note that the MedCoT model has approximately 256M parameters with 5 experts, 257M with 6 experts, and 261M with 10 experts.
- **Superior Performance:** Despite its smaller size, MedCoT achieves the highest or near-highest scores across multiple datasets, with notable performance on VQA-RAD (87.50%), SLAKE-EN (87.26%), VQA-Med 2019 (82.81%), and PathVQA (90.37%). For example, on the VQA-RAD dataset, MedCoT improved by 27.21% over Gemini and achieved a 3.31% increase compared to the strongest performance of LLaVA, which has 7B training parameters and was trained on

extensive medical data. These results indicate that our self-reflection and Hierarchical Expert design effectively reduce errors and ensure the accuracy of multimodal reasoning in LLMs. Open-end questions allow for a range of answers due to their inherent nature. The answers generated by MedCoT are difficult to match precisely against the dataset. Therefore, we employ text generation metrics such as Rouge and BLEU to evaluate MedCoT’s performance, as exhibited in **Table 3**. We conducted experiments on the open-end VQA-RAD and SLAKE-EN, with results shown in the Appendix **Table 3**. The Rouge score, similar to "recall," emphasizes the completeness of the generated text, while the BLEU score, akin to "precision," highlights its accuracy. MedCoT demonstrated higher Rouge and BLEU scores on the VQA-RAD dataset, with Rouge-1 and BLEU-1 reaching 66.30% and 61.29%, respectively, surpassing MedThink (Gai et al., 2024). Besides, MedCoT also showed higher scores on the SLAKE-EN.

### B Method Details

#### B.1 Gate Mechanism

The *Gated Dense Layer* fuses the textual representation  $F_T$  and the attention-guided visual feature  $H_V^{\text{att}}$ , deriving the fusion coefficient  $\lambda$  through a sigmoid-activated linear combination of these modalities:

$$\lambda = \text{Sigmoid}(W_l F_T + W_v H_V^{\text{att}}), \quad (7)$$

where  $W_l$  and  $W_v$  are the model parameters learned during training which optimize the blend of information from the textual and visual streams. The integrated output  $F_I \in \mathbb{R}^{n \times d}$  is then computed as a weighted sum of  $F_T$  and  $H_V^{\text{att}}$ , moderated by  $\lambda$ :

$$F_I = (1 - \lambda) \cdot F_T + \lambda \cdot H_V^{\text{att}}. \quad (8)$$

The gate  $\lambda \in [0, 1]$  is to weight the expected importance of image for source text. The gate dense layer is trainable.  $W_l$  and  $W_v$  are the learnable parameters to fuse the input  $F_T \in \mathbb{R}^{n \times d}$  and  $H_V^{\text{att}} \in \mathbb{R}^{n \times d}$  for output  $F_I$ .

#### B.2 MedCoT Method Details

For Diagnostic Specialist, we adopt the Flan-T5 encoder-decoder architecture in its Base version. The loss function is configured in accordance with

Table 2: Comparison of model performances across different datasets on closed-end question: All results are in %, the best ones are in **bold**.

Model	Size	VQA-RAD	SLAKE-EN	VQA-Med 2019	PathVQA
CGMVQA Ens. (Ren and Zhou, 2020)	-	-	-	78.10	-
MEVF (Nguyen et al., 2019)	-	77.20	79.80	-	-
MMBERT (Khare et al., 2021)	-	77.90	-	78.10	-
WDAN (Huang and Gong, 2023)	-	76.50	-	81.20	-
Gemini Pro (Qi et al., 2023)	-	60.29	72.60	60.22	70.30
PubmedCLIP (Eslami et al., 2023)	-	79.50	82.50	-	-
VL Encoder-Decoder (Bazi et al., 2023)	-	82.47	-	-	85.61
Q2ATransformer (Liu et al., 2023c)	-	81.20	-	-	88.85
VQA-Adapter (Liu et al., 2023a)	-	82.30	83.70	-	-
MedThink (Gai et al., 2024)	223M	83.50	86.30	-	-
Prefix T (Van Sonsbeek et al., 2023)	1.5B	-	82.01	-	87.00
LLaVA (Li et al., 2024)	7B	65.07	63.22	-	63.20
LLaVA-Med (From LLaVA) (Li et al., 2024)	7B	84.19	85.34	-	91.21
LLaVA-Med (From Vicuna) (Li et al., 2024)	7B	81.98	83.17	-	<b>91.65</b>
LLaVA-Med (BioMed CLIP) (Li et al., 2024)	7B	83.09	86.78	-	91.09
MedCoT (Ours)	~ 256M	<b>87.50</b>	<b>87.26</b>	<b>82.81</b>	90.37

Table 3: Evaluation Metrics for SLAKE-EN and VQA-RAD on Open-end datasets.

		Rouge1	Rouge-L	Rouge-Lsum	BLEU-1
SLAKE-EN	MedCoT (ours)	<b>80.86</b>	<b>80.14</b>	<b>80.12</b>	<b>78.33</b>
	MedThink	80.12	79.91	79.93	77.94
VQA-RAD	MedCoT (ours)	<b>66.30</b>	<b>65.78</b>	<b>65.98</b>	<b>61.29</b>
	MedThink	58.10	58.09	58.10	51.76

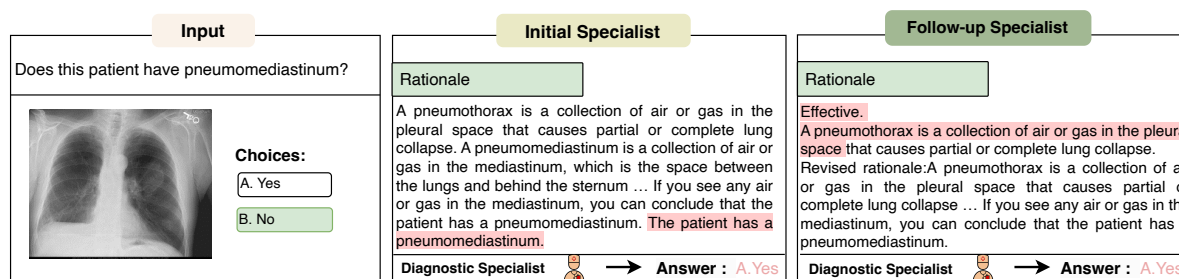


Figure 8: The MedCoT pipeline begins with an Initial Specialist receiving a medical question and image to generate a preliminary rationale. This rationale may have flaws (indicated in red), which are then reviewed by the Follow-up Specialist. If the rationale is deemed effective, it is retained; otherwise, it is reconsidered and a new rationale (indicated in green) is generated, along with an image caption. These elements are then integrated into the Diagnostic Specialist. Informed by all context, the Diagnostic Specialist, a multimodal language model with a designed sparse MoE structure, delivers the final diagnostic outcome (answer).

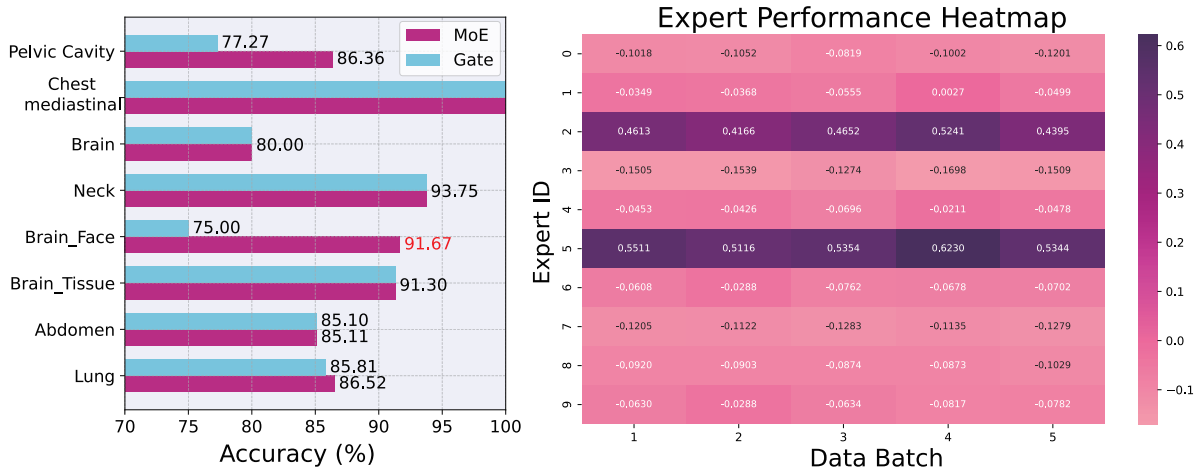


Figure 9: The Diagnostic Specialist’s sparse MoE demonstrates varying accuracy levels for different organ-related questions in the SLAKE-EN dataset. It can be observed that questions related to "Brain Face" organs saw an improvement of nearly 16%. We visualized the weights of the experts (right figure). Notably, in the top 2 expert selections, the model chose Expert 2 and Expert 5 to understand the intents of the "face" image and text. Experts 2 and 5 can be considered as "Face" specialists, proficient in diagnosing issues related to "Face".

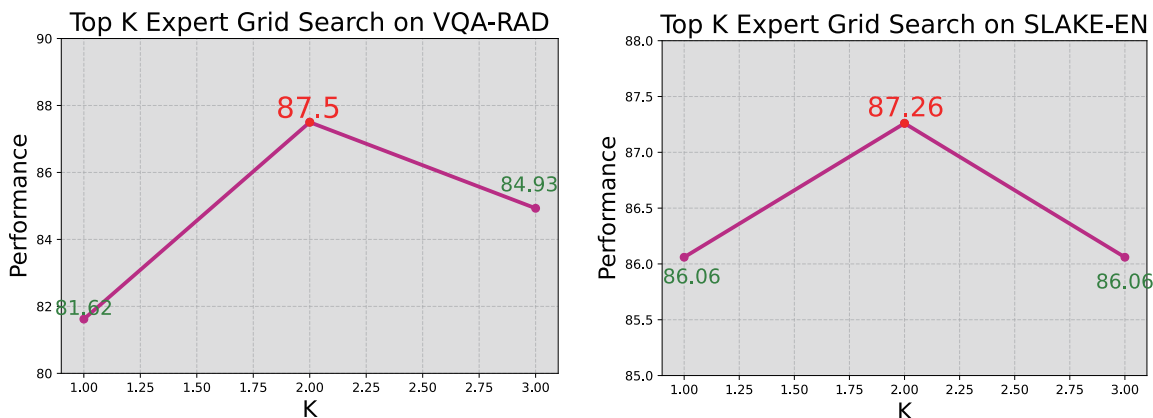


Figure 10: The top  $k$  expert grid search experiment on the VQA-RAD and SLAKE-EN.

the settings from T5 (Zhang et al., 2023b). The dimension of the vision features (processed by VisualEncoder DETR: *detr<sub>resnet101dc5</sub>*) is (100,256).

### C Datasets

Four well-known Med-VQA datasets are used in MedCoT: VQA-RAD (Lau et al., 2018), SLAKE-EN (Liu et al., 2021), VQA-2019 (Abacha et al., 2019b), and PathVQA (He et al., 2020). Each question and Image can be classified into closed-end questions and open-end questions (see Table 4).

### D The Effect of Initial Specialist

To evaluate the effect of Initial Specialist, we randomly sampled 100 question sets that had undergone Initial Specialist CoT prompting. These

Table 4: Details on datasets: The distribution of the closed-end and open-end attributes of questions in the four datasets.

Dataset	Images	QA pairs
VQA-RAD (Lau et al., 2018)	0.5k	3.5k
PathVQA (He et al., 2020)	5k	32.8k
SLAKE-EN (Liu et al., 2021)	0.7k	14k
VQA-2019 (Abacha et al., 2019a)	5k	13k

sets were then assembled into <Question, Options, Rationale, Answer> quadruples. Initial Specialist demonstrated robust performance across these quadruples, indicating that the CoT approach enhances the performance of MedVQA (see Table 5).

Besides, to validate the generalization efficacy of Initial Specialist, we randomly sampled data from

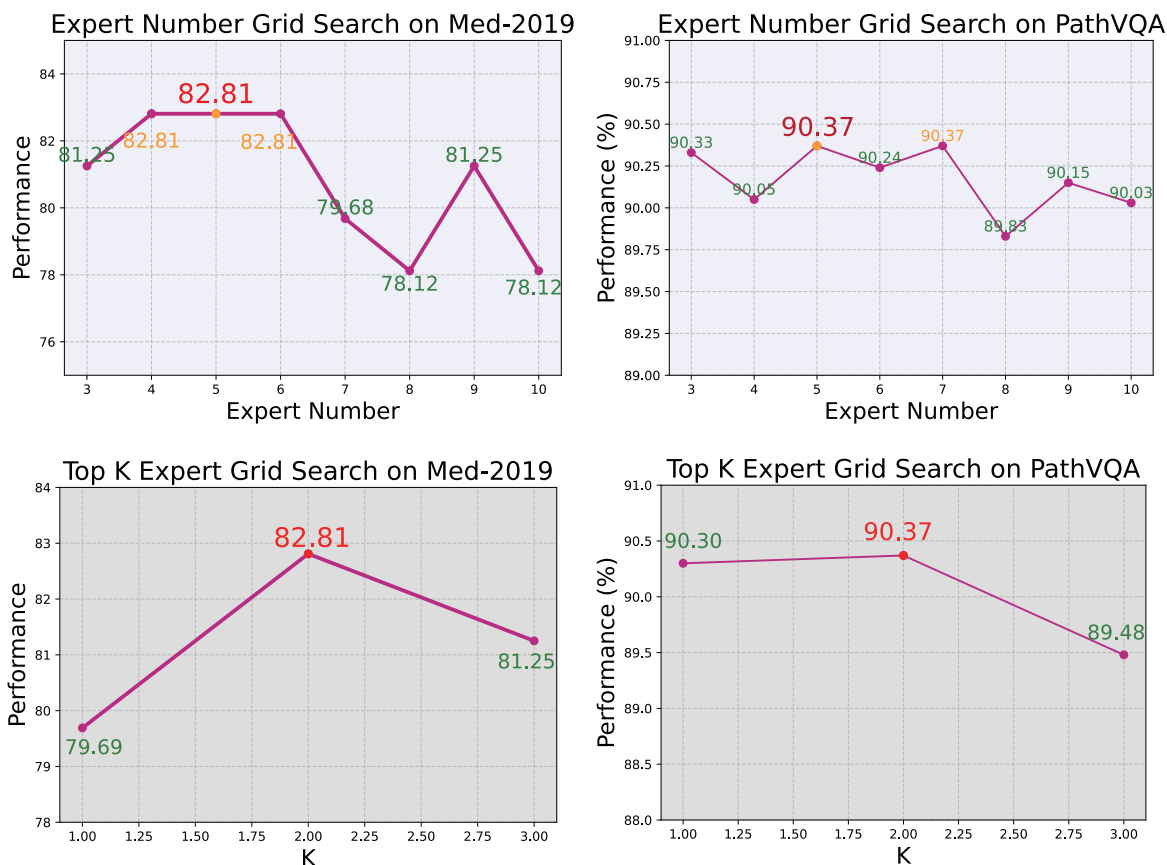


Figure 11: In the Expert Number grid search experiment on the Med-VQA-2019 and PathVQA, the purple line represents results from using rationales provided by the Follow-up Specialist and grid searching expert numbers in the Diagnostic Specialist thereafter (upper figure). The top  $k$  expert grid search experiment on the Med-VQA-2019 and PathVQA (lower figure).

four datasets to construct a validation set. The results, presented in Table 5, provide empirical evidence supporting the robustness and effectiveness in Initial Specialist.

## E Self-Reflection in Follow-up Specialist

Self-Reflection, introduced (Shinn et al., 2023; Xu et al., 2024), is a technique initially designed to assist LLMs like Gemini Pro/GPT-4 in addressing hallucinations and optimizing planning. Self-Reflection in LLMs, akin to human metacognition, involves the model self-assessing its outputs to identify and correct errors, enhance reasoning and justification, and integrate feedback and knowledge updates, thereby improving learning and problem-solving capabilities.

### E.1 Effectiveness of Self-Reflection in Follow-up Specialist

To better demonstrate the impact of model self-reflection on the quality of rationales and its in-

fluence on the final answer judgment, we employed a zero-shot prompting method. By forming <Question, Options, Rationale, Answer> quadruples as prompts and utilizing the Gemini-pro-1.5 API for evaluation, we obtained the data presented in Figure 12. The Follow-up Specialist, by employing Self-Reflection after the Initial Specialist, achieves significant improvements in the accuracy of closed-end questions across the majority of the four datasets.

However, the improvement is relatively modest, and the performance of the Follow-up Specialist on the SLAKE-EN fell below expectations. This may be due to the hallucination phenomenon experienced by LLMs during self-reflection, leading to toxic reflection Huang et al. (2023). This phenomenon can cause previously correct answers from the initial to generate incorrect rationales in follow-up specialist, thus resulting in errors.

Table 5: Verifying Initial Specialist Utility. The numbers in the table represent the count of correctly answered questions out of the 100 sampled.

	VQA-RAD	SLAKE-EN	VQA-2019	PathVQA	Mixture Set
Initial Specialist CoT prompting	72	70	68	76	73
LLMs prompting w/o rationale	47	44	37	49	46

Table 6: Example of Self-Reflection Rationale: Category can be Effective, Not Effective, Not Mention.

Category	Example
Effective	The rationale is (generally) effective.
	The (existing) rationale is effective.
	Yes, the rationale is effective for the question and image.
	The existing rationale is (generally) effective in explaining.....
Not Effective	The existing rationale is insufficient (because)
	The rationale is not effective for the question and image.
	The existing rationale is insufficient and irrelevant.
Not mention	Summary/Analysis of the existing rationale:



Figure 12: The zero-shot ability of Initial Specialist and Follow-up Specialist in Med-VQA tasks.

## E.2 Error-Analysis on Self-Reflection

To better understand the model’s reflective attitude, we categorized the rationales into three groups: "Effective," "Not Effective," and "Not Mentioned," as detailed in Table 6. This classification clarifies how the model views the initial specialist’s rationale and whether the reflection process modified it. From the statistics conducted across four datasets (shown in Table 7), the model appears to adopt a conservative approach, primarily affirming the initial specialist’s rationale. This indicates a significant influence of the initially provided results on the model’s responses.

To validate the effectiveness of Follow-up rationale on the results of closed-end questions, we employed a zero-shot prompting method. Using the <Question, Options, Rationale, Answer> quadruples from the test sets of the datasets as prompts, we allowed the LLMs to make initial judgments.

We analyzed the error cases in zero-shot prompting across four datasets. By categorizing the errors based on the reflection attitude classification proposed earlier, we analyzed their proportion relative to the total dataset.




Figure 13: Error set percentages across different datasets: the proportion of the error set in zero-shot setting was classified based on the attitude of self-reflection (The better the performance, the lower the bar).

As seen in Figure 13, the "Not Effective" category has the lowest error proportion in four datasets. This indicates that when the model deems the rationale from the Initial Specialist insufficient to answer the question reasonably, modifications made during the reflection stage can prevent some errors, resulting in the lowest proportion in the error set. Therefore, this experiment demonstrates that Self-Reflection can enhance the model’s ability to judge and solve problems in Follow-up Specialist, providing support for subsequent experiments and model improvements (An complete example in Figure 14).



Table 7: Self-evaluation in Follow-up: the solution can be divided into three categories: Effective, Insufficient, and no direct evaluation.

	VQA-RAD	SLAKE-EN	VQA-2019	PathVQA
Effective	243	353	62	2807
Not Effective	18	41	1	346
Not Mention	11	22	1	94

## F Computing Resource Costs

We compared the cost and latency of some baselines on 4 NVIDIA GEFORCE RTX 3090 GPUs, as shown in Table 8. Although MedCoT incurs higher costs and latency compared to Vanilla-CoT (which derives answers by querying rationales using Gemini Pro) and MedThink, it exhibits a significant leap in performance. This performance improvement is particularly crucial in medical scenarios, enhancing reliability.

Table 8: Cost, Latency, and VQA-RAD Accuracy for MedCoT, Vanilla (CoT), and MedThink.

	MedCoT (Ours)	Vanilla (CoT)	MedThink
Cost (tokens/sample)	1550	1493	-
Latency (seconds)	11.23	5.02	7.25
VQA-RAD (accuracy)	87.50%	60.29%	83.50%

## G Assessing the Impact of Different Prompts

MedCoT, places significant emphasis on prompt engineering, which is inherently sensitive to minor variations in wording and structure. To assess the impact of these prompt variations on the generated rationales, we conducted an experiment where we modified the prompt wording and generated a new version of the rationale. The details of these prompt modifications are outlined in Table 9.

A comparison of the original and modified rationales was conducted by calculating their cosine similarity. The results indicated that the rationales remained semantically close, with an average similarity score exceeding 70%. In the SLAKE-EN dataset, All the rationales had cosine similarity scores above 50%, with an average of 74%. Similarly, in the VQA-RAD dataset, 98.4% of the rationales achieved similarity scores above 50%, with an average score of 78.1%. These findings demonstrate that changes in prompt wording did not lead to significant semantic divergence.

To evaluate the practical implications of these modifications, we performed subsequent zero-shot

experiments using both the original and modified rationales to query Gemini Pro. The resulting zero-shot scores were comparable across both versions, with performance differences constrained to within a 3% margin. This consistency suggests that while variations in prompt wording can influence the generated rationales, their effect on model performance is relatively modest. The detailed results of this comparison are presented in Table 10.

## H Prompt Template

### H.1 Initial Specialist CoT prompt

In this work, we introduce the templates for the Initial Specialist and the Follow-up Specialist Prompt for CoT, which utilize LLMs to generate rationales and optimize VQA pairs, thereby achieving hierarchical expert authentication. Furthermore, to enhance image information and mitigate potential loss of visual data, we employ LLMs to supplement image captions.

#### Initial Specialist CoT Prompt ( $prompt_i$ )

As an agent/assistant of an experienced doctor, the next steps are required.

- Provide a reasonable rationale for the questions: question and image entered.
- Please proceed with a step-by-step analysis and provide a rationale.

Table 9: Prompt Modifications for Initial and Follow-up Specialists

Role	Type	Prompt
Initial Specialist	Original	As ... <b>experienced doctor</b> , next steps are required. Provide a reasonable rationale for the questions: {question} ...
Initial Specialist	Change	Acting ... assistant to a <b>seasoned doctor</b> , ... explain why each question {question}...
Follow-up Specialist	Original	...Evaluate the effectiveness of the given rationale. <b>If effective</b> , summarize and refine it. If insufficient, ...
Follow-up Specialist	Change	As an assistant to an experienced doctor, ... Check if the rationale: {rationale} is <b>useful</b> ...

Table 10: Performance Variations across Different Prompts.

Dataset	Original Setting	Change Setting	Absolute Change
VQA-RAD	60.29%	57.45%	2.84%
SLAKE-EN	72.40%	73.04%	0.64%

## H.2 Follow-up Specialist Prompt

### Follow-up Specialist Self-Reflection Prompt ( $prompt_{\hat{f}}$ )

Please as an agent/assistant of an experienced doctor: Question and Image is also included, it is useful to provide the information


- Existing Rationale: {rationale}; Task: Please judge whether this rationale is effectively valid for question and image.
- If it is effective, summarize and refine the rationale to highlight key points, else the existing rationale is insufficient or irrelevant, generate a clear, more precise, and concise explanation.

## H.3 Follow-up specialist Image Caption Prompt

### Image Caption Prompt

Your task is to add a caption for the image, the example is

- The image features a black and white X-ray of a person's chest, revealing their lungs and rib cage. ...[Few-shot Example of an Image]... Overall, the X-ray provides a comprehensive view of the person's respiratory system.
- So this image caption can be:[caption]



**Problem:** Is the heart enlarged?  
**Choices:**  Yes  No

**Initial specialist rationale Gemini**

**Rationale:** The heart is located in the upper left quadrant of the chest. The heart is enlarged if it takes up more than half of the space in the upper left quadrant of the chest. **Analysis:** The heart takes up more than half of the space in the upper left quadrant of the chest. **Therefore, the heart is not enlarged.** The next step is to determine if it is the cause of the enlarged heart. The doctor will need to order more tests, such as an echocardiogram, to determine whether it is enlarged heart.

**Gemini Initial specialist Judgment**  No

**Follow-up specialist rationale with Image Caption Gemini**

**Caption:** The image features a black and white X-ray of a person's chest, revealing their lungs and ribs. The person's upper body is partially visible, with the chest area being the main focus of the X-ray. The lungs and ribs are clearly visible in the image, providing a detailed view of the person's respiratory system. The X-ray is likely taken as part of a medical examination or for diagnostic purposes.

**Rationale:** The existing rationale is **generally effective, but it can be improved.** The rationale can be summarized as follows: The heart is located in the upper left quadrant of the chest. The heart is enlarged if it takes up more than half of the space in the upper left quadrant of the chest. ... **Therefore, the heart is enlarged.** The next step is to determine the cause of the enlarged heart. ... The rationale can be refined by providing more details about the location of the heart and the size of the heart. For example, the rationale could state that the heart is located between the lungs and that it is about the size of a fist. The rationale could also provide more information about **the causes of an enlarged heart.** This could be caused by a lot of factors, such as high blood pressure, heart disease, and obesity.

**Diagnostic specialist Judgment**  Yes  **Case: The initial specialist judgment is wrong, while it is corrected by follow-up reflection.**

Figure 14: Detailed Procedure of Specialist Self-Reflection and Diagnostic Correction: It starts with the initial problem statement, followed by the Initial Specialist's rationale and judgment. The Follow-up Specialist provides a detailed image caption and refines the rationale by Self-Reflection. The final diagnostic judgment confirms the corrected assessment, demonstrating the importance of self-reflection for accurate medical diagnoses.