

# MMXU: A Multi-Modal and Multi-X-ray Understanding Dataset for Disease Progression

Linjie Mu<sup>1,3</sup>, Zhongzhen Huang<sup>1,3</sup>, Shengqian Qin<sup>1,3</sup>, Yakun Zhu<sup>1,3</sup>,  
Shaoting Zhang<sup>1†</sup>, Xiaofan Zhang<sup>1,2,3†</sup>

<sup>1</sup>Shanghai Jiao Tong University, <sup>2</sup>SII, <sup>3</sup>SPIRAL Lab

Correspondence: xiaofan.zhang@sjtu.edu.cn

## Abstract

Large vision-language models (LVLMs) have shown great promise in medical applications, particularly in visual question answering (MedVQA) and diagnosis from medical images. However, existing datasets and models often fail to consider critical aspects of medical diagnostics, such as the integration of historical records and the analysis of disease progression over time. In this paper, we introduce MMXU (Multimodal and MultiX-ray Understanding), a novel dataset for MedVQA that focuses on identifying changes in specific regions between two patient visits. Unlike previous datasets that primarily address single-image questions, MMXU enables multi-image questions, incorporating both current and historical patient data. We demonstrate the limitations of current LVLMs in identifying disease progression on MMXU-test, even those that perform well on traditional benchmarks. To address this, we propose a MedRecord-Augmented Generation (MAG) approach, incorporating both global and regional historical records. Our experiments show that integrating historical records significantly enhances diagnostic accuracy by at least 20%, bridging the gap between current LVLMs and human expert performance. Additionally, we fine-tune models with MAG on MMXU-dev, which demonstrates notable improvements. We hope this work could illuminate the avenue of advancing the use of LVLMs in medical diagnostics by emphasizing the importance of historical context in interpreting medical images. Our dataset is released at [github\\*](https://github.com/linjiemu/MMXU).

## 1 Introduction

Stemming from the ever-increasing number of parameters and large-scale training corpus, large vision-language models (LVLMs) (Zhu et al., 2023; Bai et al., 2023; Liu et al., 2023; Achiam et al.,

**Question:** What change is observed in the **mediastinum region** between the two sets of chest CXR images?

**Options:** A: There is a new rightward mediastinal shift.  
B: There is no mediastinal shift observed.  
C: There is a new leftward mediastinal shift.  
D: The mediastinal shift has improved.

**Answer:** A: There is a new rightward mediastinal shift.

**Images:**





Figure 1: The upper part of the figure presents an example from our constructed dataset, which includes two images and a question regarding their changes. The lower part illustrates the performance of human experts and several models on three types of disease progression questions, emphasizing a notable gap between the models and human experts.

2023; Lu et al., 2024; Chen et al., 2024c) have demonstrated remarkable capabilities in general visual scene perception and understanding, as well as in generating textual descriptions. As the development of LVLMs accelerates, this paradigm has spurred significant advancements in the medical field (Li et al., 2023b; Wu et al., 2024a; Li et al., 2024a; Chen et al., 2024a), particularly in the analysis and diagnosis of medical images.

Although current medical LVLMs have shown strong performance in downstream tasks such as medical visual question answering (MedVQA) (Hu et al., 2024; Sun et al., 2024; Saeed, 2024) and medical report generation (Zhou and Wang, 2024; Bu

<sup>†</sup>Corresponding authors

\* <https://github.com/linjiemu/MMXU>

et al., 2024; Huang et al., 2025; Yin et al., 2025) on public benchmarks, they are still limited in responding to basic visual questions (i.e., those involving a single image and a brief description (Bae et al., 2024; Liu et al., 2024)). In real-world scenarios, diagnosis often requires physicians to integrate both case history and current evidence (Lorkowski and Pokorski, 2022). Case history serves as an invaluable source of evidence, encompassing factors such as previous medical conditions, treatments, and patient demographics (Liu et al., 2024). Identifying the differences between symptoms, clinical signs, and diagnostic outcomes across the timeline is key to accurate diagnosis. Based on these considerations, we ask: Can LVLMs enhance the diagnostic process by identifying these critical factors?

In this paper, we introduce MMXU (Multimodal and MultiX-ray Understanding), leveraging the rich resources of patient electronic medical records (EMRs) from the MIMIC-CXR dataset. As shown in Figure 1, unlike previous works (Bae et al., 2024; Liu et al., 2024), which focus solely on the current image or are restricted to simple questions, this benchmark is specifically designed to inquire about differences in specific regions between a patient’s two visits. MMXU contains two splits: *test* and *dev*. MMXU-*test* consists of 3,000 entries from 1,201 patients and 2,469 studies, while MMXU-*dev* contains 118K QA pairs involving 114K images. We evaluate a range of open-source and closed-source models and conduct a user study to assess human performance. The results indicate that LVLMs struggle to identify differences between two visits, even the LVLM that demonstrates notable performance on other medical benchmarks exhibits a significant performance gap—nearly 40%—when compared to human performance.

Recognizing the limitations of current LVLMs in such scenarios, we take steps toward bridging this gap. Since physicians often rely on patient history, previous treatments, and other contextual information during the diagnostic process, we propose a novel approach, MedRecord-Augmented Generation (MAG) to facilitate LVLMs. We introduce two types of records as contextual information for diagnosis: global records (e.g., patient demographics and overall medical report) and regional records (e.g., specific regional details and diagnostic results). Our experiments show that incorporating global or regional medical records leads to a significant accuracy improvement of at least 20% for several well-known models. Acknowledging the in-

adequate training on such tasks, we apply the MAG method for fine-tuning the MMXU-*dev* dataset and substantiate the effectiveness of both the dataset and the MAG approach in enhancing overall performance.

In summary, this paper presents three key contributions. First, we introduce MMXU, the first multi-image MedVQA dataset designed to investigate differences in specific regions of CXR images between a patient’s two visits, addressing the gap between current MedVQA benchmarks and real-world clinical scenarios. Second, our evaluation results show that current large vision-language models (LVLMs), including proprietary models such as GPT-4o, face significant challenges in identifying disease progression between two visits. Third, we propose the MedRecord-Augmented Generation (MAG) method to enhance the diagnosis of X-rays over time by leveraging the patient’s global and regional historical records. Experimental results demonstrate the effectiveness of MAG, underscoring the potential of incorporating contextual information in diagnostic processes.

## 2 Related Work

**Large Vision-Language Models** Large Vision-Language Models, which integrate vision encoders, connectors, and large language models to enhance cross-modal understanding, have emerged as powerful frameworks that combine visual and textual information for a wide range of tasks. These models can be systematically categorized based on the type of connector. The first category comprises approaches utilizing cross-attention-based connectors, such as Flamingo (Alayrac et al., 2022) and CogVLM (Wang et al., 2023; Hong et al., 2024), which exploit attention mechanisms to facilitate the exchange of information between the vision encoder and the language model. The second category includes methods that employ query-based connectors, such as BLIP-2 (Li et al., 2023a), Instruct-BLIP (Dai et al., 2023), mPLUG-owl2 (Ye et al., 2024), and Qwen-VL (Bai et al., 2023), wherein queries are leveraged to orchestrate the interaction between visual and textual modalities, thereby enhancing the alignment and coherence of visual and linguistic representations. Furthermore, projection-based connector methods, exemplified by LLaVA (Liu et al., 2023), Mini-GPT4 (Zhu et al., 2023), DeepSeek-VL (Lu et al., 2024), and Mini-Gemini (Li et al., 2024b), project visual data into a

Dataset	#Images	# QA Pairs	#Question Type	#Regional	#Comparative	#Complex
MIMIC-CXR-VQA (Bae et al., 2024)	142K	377K	Single	✓	✗	✓
GEMeX (Liu et al., 2024)	151K	1.6M	Single	✓	✗	✓
Medical-Diff-VQA (Hu et al., 2023)	164K	700K	Single&Multiple	✗	✓	✗
MMXU-dev (Ours)	114K	118K	Multiple	✓	✓	✓

Table 1: Comparison of the MedVQA dataset constructed on MIMIC-CXR. In the #Question Type column, “Single” refers to questions about a single image, and “Multiple” refers to questions about multiple images. Most existing datasets primarily focus on observations in a single image. While Medical-Diff-VQA contains 131k QA pairs that focus on multi-image changes, it has a simple structure and only considers global-level differences. Our MMXU-dev dataset is the first to focus on complex changes in the same regions across multiple images at the regional level.

shared embedding space, thereby fostering seamless integration with textual information. These innovations offer a range of solutions for cross-modal understanding, driving the potential applications of intelligent systems in multi-task learning.

**MedVQA Dataset on radiology** Medical visual question answering (MedVQA) datasets play a pivotal role in advancing AI-driven clinical decision-making. VQA-RAD (Lau et al., 2018), as an early pioneering work, introduces a meticulously curated dataset for radiology images, featuring clinician-generated questions and corresponding answers tailored to clinically relevant tasks. SLAKE (Liu et al., 2021) stands out as a large, bilingual dataset, enriched with extensive semantic annotations and spanning a wide range of radiological modalities.

MIMIC-CXR (Johnson et al., 2019) provides a vast collection of 371,920 chest X-rays from 65,079 patients, serving as the foundation for numerous subsequent studies. The comparison of these datasets is shown in Table 1. MIMIC-CXR-VQA (Bae et al., 2024) seamlessly integrates chest X-rays with Electronic Health Records (EHRs), facilitating multi-modal question answering with an emphasis on region-specific queries. The Medical-Diff-VQA (Hu et al., 2023) is notable for its inclusion of seven distinct question types, particularly focusing on the comparative analysis of current and reference images for diagnostic purposes. GEMeX (Liu et al., 2024) offers a large-scale, explainable VQA benchmark, complete with detailed visual and textual explanations, thus addressing the growing need for a diverse array of clinical questions. Most of these datasets primarily focus on observations from a single image. While Medical-Diff-VQA contains 131k QA pairs that address multi-image changes, its structure is relatively simple and only accounts for global-level differences. Our MMXU-dev dataset is the first to understand complex changes in the same regions across multi-

ple chest X-ray images of the same patient at the regional level, spanning several visits.

### 3 Dataset Construction

In this section, we outline the pipeline of constructing MMXU, as shown in Figure 2. The process starts with the Chest ImaGenome dataset (Wu et al., 2021), which includes the *silver\_dataset* section containing annotations for 243,310 images from 63,945 patients. These annotations cover bounding boxes for 29 anatomical regions, along with corresponding region-level report *phrases*, labeled *attributes*, and *relationships*. The entire method consists of four distinct phases: (1) Comparative Sentences Extraction (§3.1), (2) Comparative Targets Selection (§3.2), (3) QA pairs Generation (§3.3), and (4) Post-Processing (§3.4).

#### 3.1 Comparative Sentences Extraction

In the first stage, our objective is to identify and extract sentences that contain comparative information, forming the foundation for generating question-answer pairs in subsequent stages. For example, the sentence “Previously seen ill-defined peribronchial lower lobe opacity seen on lateral view has resolved,” along with its associated *relationships* label “comparisonyeslimproved” explicitly indicates a comparison with prior conditions, highlighting the resolution of the lower lobe opacity and thus signaling an improvement in the patient’s condition. In total, we extract 232,247 comparative sentences, encompassing 22,770 patients and 102,606 reports. A more detailed example of such a sentence is provided in Appendix C.1.

#### 3.2 Comparative Targets Selection

At this stage, we categorize comparative sentences into three groups based on disease progression: “Worsen” (108,734), “Improved” (91,084), and “No Change” (264,676), with the majority concentrated

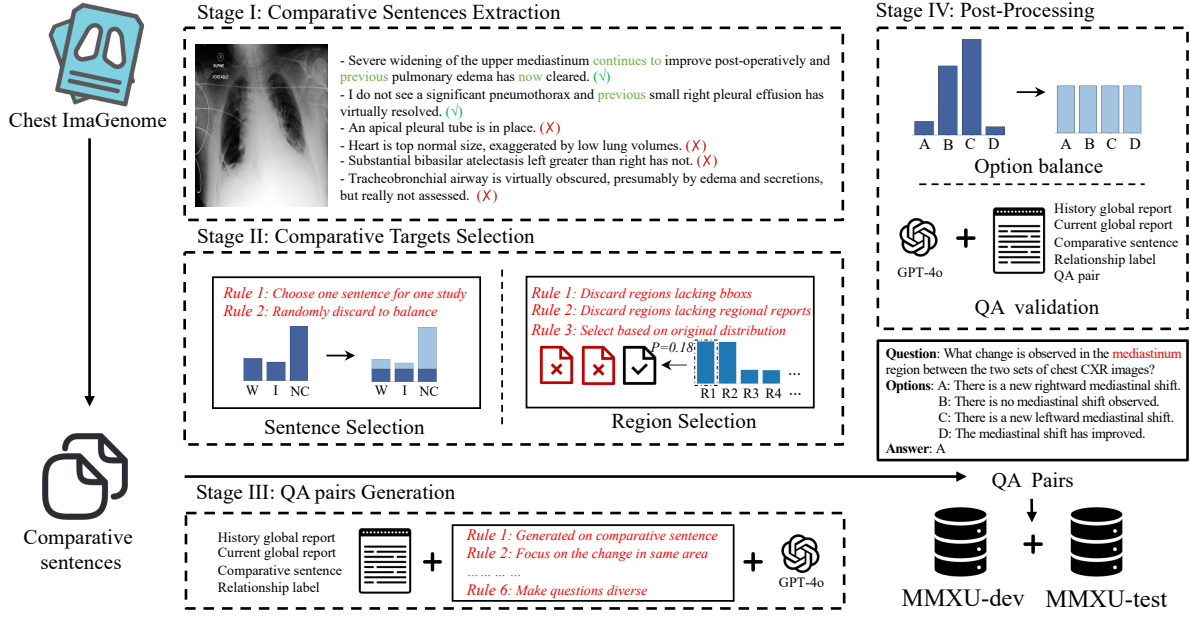


Figure 2: The pipeline of constructing MMXU. In the first stage, we extract all comparative sentences and their associated labels from the Chest ImaGenome dataset. In the second stage, we refine our selection by applying precise filtering rules and identifying regions of interest within these sentences. The third stage involves leveraging GPT-4o with meticulously crafted prompts to generate region-level question-answer pairs. In the final stage, we further refine and filter generated QA pairs, thereby constructing the MMXU-test and MMXU-dev.

in the “No Change” category. Given that the prevalence of medical conditions often follows a long-tail distribution (Wu et al., 2024b), it is essential to achieve a more balanced dataset for a robust evaluation of LVLMs. To this end, we implement two key principles for sentence selection: (1) Retaining only one comparative sentence per report, prioritizing those indicative of change, to enhance data diversity and reduce redundancy. (2) Randomly discarding a portion of “No Change” sentences to equalize their count with that of the “Improved” and “Worsen” categories, thereby mitigating class imbalance.

Next, we implement a regional selection process to improve query accuracy, where one region is chosen for each sentence to question. In real-world applications, a comparative statement may reference multiple anatomical regions simultaneously. For instance, the sentence “In comparison with the study of \_\_\_, there has been worsening of the increased opacification at the left base with silhouetting of the hemidiaphragm and blunting of the costophrenic angle” concurrently describes the “left lung”, “left lower lung zone”, “left costophrenic angle” and “left hemidiaphragm”. Consequently, it is crucial to select the precise regions mentioned in the sentence to eliminate ambiguity. Our approach is as follows:

- (1) Exclude regions in both current and historical images that lack a clearly defined bounding box.
- (2) Eliminate regions that are not referenced in prior reports, as it is infeasible to retrieve historical region-level reports.
- (3) In cases where multiple regions remain, randomly select one based on the original distribution.

### 3.3 QA pairs Generation

In the third stage, we employ GPT-4o as the generator, leveraging comparative sentences, *relationship* labels, and reports from two visits as foundational data to guide the model in producing question-answer pairs. Our observations suggest that a single sentence can encapsulate multiple changes. For instance, the sentence “Bilateral pleural effusions are again seen, and atelectasis is present” concurrently describes alterations in both “pleural effusions” and “atelectasis”. To fully harness this data, we direct GPT-4o to generate up to three distinct question-answer pairs, ensuring maximum diversity. Additionally, to facilitate the future validation of data accuracy, we require the model to provide justifications for the generated pairs.

In total, we have established six rules to guide GPT-4o generating QA pairs that capture changes in the same area across the two reports. Appendix



Here are two chest X-RAY images reports of the same patient. Previous chest X-RAY is Image-1. Current chest X-RAY is Image-2. Your task is to evaluate the differences between the two images based on the provided report and question.  
The report of Image-1 is: %s  
The report of Image-2 is: %s  
Question:%s  
Options:  
A: %s  
B: %s  
C: %s  
D: %s  
Answer with the option's letter from the given choices directly.

Figure 3: Text template during the data post-processing.

C.2 presents our prompt template and a detailed example during the QA generation process.

### 3.4 Post-Processing

Finally, we conduct post-processing to further ensure the quality of generated QA-pairs. Since imbalanced answer options can mislead the model during training and result in unfair evaluations during testing, we first balance the distribution of four answer choices in the QA pairs generated by GPT-4o. Moreover, we employ GPT-4o to answer the generated questions with the prompt detailed in Figure 3, which contains the information used during question generation. Questions that were answered incorrectly are considered either excessively difficult or erroneous and are removed. As illustrated in Table 2, 95.1% of the questions were answered correctly. The final candidate set comprises 121,800 QA pairs, including 47,000 instances of “Worsen”, 41,000 instances of “Improved” and 35,000 instances of “No Change”.

	#Examples	#Correctness	Rate
<b>Worsen</b>	50,305	45,395	90.2%
<b>Improved</b>	40,647	39,956	98.3%
<b>No change</b>	37,124	33,449	90.1%
<b>Overall</b>	128,076	121,800	95.1%

Table 2: Post-processing validation results of QA pairs using GPT-4o. Here, #Total and #Correctness represent the total number of original QA pairs and the number of correctly answered pairs, respectively.

**Data Splitting** To facilitate a rigorous evaluation for LVLMS, we carefully curate the MMXU-

test benchmark comprising 1,000 data samples for each of the three categories—“Worsen”, “Improved” and “No Change”. The remaining data forms the MMXU-dev dataset. During the selection process, we meticulously ensured that questions derived from the same patient did not appear in both the dev and test sets, thereby preventing data leakage. Furthermore, we maintained a balanced regional distribution across the training and test sets to preserve the integrity of the evaluation. Finally, the MMXU-test contains 3,000 QA pairs, and MMXU-dev contains 118,800 QA pairs.

## 4 Benchmark Results

To ensure the professionalism and accuracy of MMXU-test benchmark, we recruited a panel of 5 board-certified chest radiology experts to assess it. Following that, we evaluated the performance of several prominent open-source and closed-source LVLMS capable of supporting multi-image VQA on the MMXU-test benchmark. Since all the questions from our benchmark are single-choice, we use accuracy as the metric.

### 4.1 Evaluation Models

For the general domain, we evaluated the following models:

**Open-source LVLMS:** Qwen2-VL 2B& 7B (Bai et al., 2023), DeepSeek-VL 1.3B& 7B (Lu et al., 2024), InternVL2 1B&2B& 4B&8B (Chen et al., 2024c), IDEFICS2 8B (Laurençon et al., 2024) and Llama3.2-Vision 11B (Touvron et al., 2023)  
**Closed-source LVLMS:** GPT-4o (Achiam et al., 2023) and Claude-3-5-sonnet

For the medical domain, we evaluated the following models:

**Medical LVLMS:** LLaVA-Med-v1.5 (Li et al., 2024a), Med-Flamingo (Moor et al., 2023) and HuatuoGPT-vision (Chen et al., 2024b)

### 4.2 Human Expert Evaluation

To evaluate the quality of the MMXU-test benchmark, we conduct the human expert evaluation with five radiologists. The data from the MMXU-test benchmark was randomly divided into five parts, containing 500, 500, 500, 750, and 750 questions, respectively. We ensured that the three question categories were distributed as evenly as possible within each subset. The evaluation results are presented in Table 4. Except for Expert 3, all experts achieved an accuracy rate of at least 96.0%, with

Source	LVLMs	Size	VQA-RAD			MMXU-test			
			Closed	Open	Overall	Worsen	Improved	No change	Overall
Open	Qwen2-VL	2B	0.594	0.380	0.499	<b>0.712</b>	0.258	0.284	0.418
		7B	<b>0.745</b>	0.430	<b>0.605</b>	0.331	0.554	0.500	0.458
	DeepSeek-VL	1.3B	0.566	0.250	0.426	0.292	0.294	0.288	0.291
		7B	0.582	0.300	0.457	0.389	0.310	<u>0.606</u>	0.435
	InternVL2	1B	0.490	0.265	0.390	0.553	0.214	0.113	0.293
		2B	0.641	0.350	0.512	0.560	0.182	0.559	0.434
		4B	0.649	0.370	0.525	0.476	0.571	0.315	0.454
		8B	0.665	0.480	0.583	0.423	0.483	0.495	0.467
Closed	IDEFICS2	8B	<u>0.673</u>	0.450	0.574	0.234	0.570	<b>0.668</b>	0.491
	Llama3.2-vision	11B	0.649	<b>0.515</b>	<u>0.590</u>	<u>0.596</u>	<u>0.608</u>	0.356	<u>0.520</u>
	GPT-4o	-	0.578	0.480	0.534	0.480	<b>0.675</b>	0.559	<b>0.571</b>
	Claude-3.5	-	0.622	<u>0.510</u>	0.572	0.494	0.518	0.493	0.502

Table 3: Evaluation results of several mainstream open-source and closed-source LVLMs supporting multi-image question answering on the VQA-RAD and MMXU-test benchmarks. For open-ended questions in VQA-RAD, we used GPT-4o to evaluate LVLMs’ responses. The results in **bold** and underlined represent the best and the second-best results, respectively.

Expert ID	#Examples	#Correctness	Rate
1	500	487	97.4%
2	500	480	96.0%
3	500	453	90.6%
4	750	732	97.6%
5	750	723	96.4%
<b>Total</b>	3,000	2,875	95.8%

Table 4: Evaluation results of five experienced chest radiology human experts on the MMXU-test benchmark.

an overall accuracy reaching 95.3%. These findings demonstrate that the MMXU-test benchmark is both highly professional and well-structured.

### 4.3 Model Performance

We evaluated the performance of several well-known open-source and closed-source models on two benchmarks: VQA-RAD, which focuses on single-image visual question answering, and MMXU-test, which emphasizes multi-image difference analysis. The experimental results are presented in Table 3. For open-ended questions in VQA-RAD (Lau et al., 2018), we utilized GPT-4o to compare the model-generated answers with the ground truth, incorporating a certain level of tolerance to assess correctness (He et al., 2024).

Comparing the performance of these models on single-image and multi-image benchmarks, it is clear that **the accuracy of the same model on**

Model	Wors.	Impr.	Noch.	Over.
Med-Flamingo-7B	0.289	0.304	0.251	0.281
Llava-med-7B	0.293	0.334	0.323	0.317
HuatuoGPT-Vision-7B	0.629	0.502	0.354	0.495

Table 5: Performance comparison of various medical vision-language models.

**overall accuracy on MMXU-test is lower than its VQA-RAD closed-ended questions.** For example, with the Qwen2-VL 7B model, the accuracy on VQA-RAD closed questions reaches an impressive 74.5%. However, its performance on MMXU-test drops significantly to just 45.8%. This discrepancy highlights the gap between current MedVQA benchmarks and the demands of real-world scenarios, suggesting that models that perform well on public benchmarks may not be effective in supporting clinical diagnosis. Furthermore, as shown in Table 4, even the best-performing model lags behind human experts by nearly 40%. This highlights the existing limitations of these models in multi-image MedVQA tasks.

Additionally, we observe that nearly all open-source models **demonstrate substantial disparities in accuracy across the three question types** in the MMXU-test, indicating inherent biases in how these models interpret disease progression. This phenomenon is especially pronounced

in smaller models—for instance, Qwen2-VL 2B achieves an accuracy of 71.2% on “Worsen” questions, yet performs below 30% on the remaining two types. Evidently, the model exhibits a strong inclination toward outcomes associated with disease deterioration. In contrast, such biases are less prominent in closed-source models. **With increasing model size, overall accuracy improves**, and the extent of progression-related bias tends to decrease. Despite being general-purpose in nature, these models still produce divergent results. Notably, smaller models tend to yield relatively consistent performance, whereas discrepancies become more pronounced once the model size surpasses 7B. For example, models such as DeepSeek-VL 1.3B and InternVL2 1B, or Qwen2-VL 2B and InternVL2 2B, exhibit similar capabilities. Conversely, larger models like DeepSeek-VL 7B, InternVL2 8B, and IDEFICS2 8B demonstrate marked differences in their performance.

As shown in Table 5, we also evaluated three medical models. For models restricted to single-image inputs, we concatenated two images to facilitate evaluation. The results, summarized in the table below, indicate that HuatuoGPT-Vision outperformed other models of similar size and achieved performance comparable to the commercial model Claude 3.5.

## 5 MedRecord-Augmented Generation

Through prior experiments, we pinpointed that current LVLMs face significant challenges in identifying disease progression. To address this challenge, we propose a novel approach, MedRecord-Augmented Generation (MAG). In routine diagnostic practice, physicians often rely on a patient’s historical records to inform their analysis and diagnosis of current conditions. To replicate this process, we integrate historical records directly into the prompt as contextual information. Our study investigates the effectiveness of global reports derived from historical images, alongside regional reports related to the specified questions, which we categorize as global and regional historical records.

### 5.1 Effectiveness of MAG

To more comprehensively evaluate the effectiveness of MAG, we conducted tests across general-domain open-source models, commercial models, and medical models. The outcomes of these experiments are summarized in Table 6. It is evident

Method	Worsen	Improved	No change	Overall
<i>InternVL2 8B</i>				
-	0.423	0.483	0.495	0.467
+ G-MRec	0.440	0.644	0.838	0.641
+ R-MRec	<b>0.483</b>	<b>0.830</b>	<b>0.842</b>	<b>0.718</b>
<i>Qwen2-VL 7B</i>				
-	0.331	0.544	0.500	0.458
+ G-MRec	0.427	0.683	0.784	0.631
+ R-MRec	<b>0.458</b>	<b>0.755</b>	<b>0.805</b>	<b>0.673</b>
<i>GPT-4o</i>				
-	<b>0.480</b>	0.675	0.559	0.571
+ G-MRec	0.380	0.760	0.629	0.590
+ R-MRec	0.374	<b>0.802</b>	<b>0.765</b>	<b>0.647</b>
<i>HuatuoGPT-Vision 7B</i>				
-	0.629	0.502	0.354	0.495
+ G-MRec	0.468	0.660	0.792	0.640
+ R-MRec	<b>0.490</b>	<b>0.712</b>	<b>0.775</b>	<b>0.659</b>
<i>Llava-med-v1.5-mistral 7B</i>				
-	0.293	0.334	0.323	0.317
+ G-MRec	0.378	0.468	<b>0.708</b>	<b>0.518</b>
+ R-MRec	<b>0.396</b>	<b>0.493</b>	0.647	0.512

Table 6: Results of MedRecord-augmented generation on the MMXU-test benchmark without fine-tuning. G-MRec and R-MRec denote generation augmentation using global and regional historical records, respectively.

that providing historical records significantly enhances overall model accuracy, regardless of the model type. With access to historical records, open-source models achieved performance comparable to that of closed-source commercial models.

With the exception of GPT-4o and Llava-med-v1.5, incorporating medical records led to improvements across all three question types, suggesting that these models are capable of understanding and reasoning from historical data rather than merely repeating it. Furthermore, it is noticeable that regional historical records yield greater improvements than global records in nearly all models. This may be due to the models’ limited ability to analyze contextual and region-specific nuances, which makes it challenging to pinpoint the most relevant information. By supplying precise, localized historical data, the models are better equipped to extract meaningful insights.

The extent to which historical records contribute to performance improvement varies across different question categories. As illustrated in Table 6, historical records notably enhance performances in the “Improved” and “No Change” categories, while providing only minimal benefit in the “Worsen” category. In fact, for GPT-4o, performance in this category even deteriorated. This indicates that the use of historical records may **not effectively address**

Method	Worsen	Improved	No change	Overall
<i>InternVL2 8B + 20% MMXU-dev</i>				
-	0.834	0.802	0.810	0.815
+ G-MRec	0.856	0.841	0.843	0.847
+ R-MRec	0.866	0.830	0.848	0.848
<i>InternVL2 8B + 40% MMXU-dev</i>				
-	0.838	0.813	0.839	0.830
+ G-MRec	0.867	0.857	0.867	0.864
+ R-MRec	0.868	0.845	0.874	0.862
<i>InternVL2 8B + 60% MMXU-dev</i>				
-	0.855	0.822	0.841	0.839
+ G-MRec	0.884	0.855	0.873	0.871
+ R-MRec	0.880	0.855	0.869	0.868
<i>InternVL2 8B + 80% MMXU-dev</i>				
-	0.857	0.814	0.866	0.846
+ G-MRec	0.881	0.861	0.884	0.876
+ R-MRec	0.873	0.869	0.881	0.875
<i>InternVL2 8B + 100% MMXU-dev</i>				
-	0.860	0.840	0.852	0.851
+ G-MRec	0.887	0.878	0.882	0.883
+ R-MRec	0.884	0.871	0.888	0.881

Table 7: Fine-tuning Results of InternVL2 8B on MMXU-dev dataset with MAG Method. G-MRec and R-MRec denote generation augmentation using global and regional historical records, respectively.

biases in a high-quality and robust manner.

## 5.2 MAG Fine-tuning on MMXU-dev

Furthermore, we assess the efficacy of MMXU-dev dataset and our proposed MAG method by fine-tuning the InternVL2 8B. The model was fine-tuned using 20%, 40%, 60%, 80%, and 100% of the MMXU-dev dataset. We evaluate the performance on MMXU-test benchmark. The corresponding fine-tuning results are presented in Table 7.

We observed that even with just 20% of the MMXU-dev, the model achieved substantial improvements in accuracy across all three problem categories, and **the bias was almost eliminated**. As the dataset size expanded, the model’s overall performance continued to improve. When utilizing the full 100% of MMXU-dev, the performance gap between the model and human experts narrowed to approximately 10%. This clearly underscores the efficacy of the MMXU-dev dataset we developed. Additionally, we found that the MAG method continues to deliver significant improvements after model fine-tuning. Notably, as the volume of training data increases, **the enhancement effect of global historical records surpasses that of regional historical records**. Fine-tuning with larger datasets enables the model to more effectively capture relevant information, resulting in more comprehensive and refined final outputs.

Method	Worsen	Improved	No Change	Overall
InternVL2 8B	0.423	0.483	0.495	0.467
+ G-MRec	0.443	0.644	0.838	0.641
+ G-MRec w/o PI	0.415	0.510	0.835	0.573
+ R-MRec	0.483	0.830	0.842	0.718
+ R-MRec w/o PI	0.364	0.593	0.823	0.593

Table 8: Ablation study on the impact of prior-image information and historical records. Here, w/o PI denotes the removal of the prior image when using the MAG method.

## 5.3 Ablation Study on MAG

To assess whether the MAG model disproportionately relies on report content, we evaluate the performance of InternVL2 8B using only global and regional historical records—excluding imaging data—for question answering. As shown in Table 8, the experimental results indicate a marked performance decline in the zero-shot setting when relying solely on textual input, as compared to utilizing both textual and visual modalities.

This degradation is expected and underscores the critical role of imaging data. Several factors contribute to this dependency: clinical reports often lack detailed descriptions of prior findings or may contain inconsistencies. In cases involving disease progression or resolution, textual narratives alone frequently fall short of conveying the nuance required to assess subtle changes. Image comparison is indispensable for accurately gauging such progression. For example, if a prior report notes an anomaly still visible in the current scan, the text alone cannot determine whether the condition has improved, deteriorated, or remained stable. Furthermore, if a subtle lesion was previously overlooked and thus omitted from the report, the model might incorrectly classify it as a novel finding, despite its earlier presence in the images.

## 5.4 Case Study

In Figure 4, we present two answer samples provided by GPT-4o, InternVL2 8B, Qwen2-VL 7B, and fine-tuned InternVL2 8B, evaluated under various historical record strategies. Case 1 illustrates a successful scenario where all models adjusted their responses accurately when regional historical records were incorporated. However, the zero-shot InternVL model misinterpreted critical information, primarily due to the global records containing terms such as “unchanged” and “minimally,” which led to an erroneous conclusion. Fine-tuning effec-



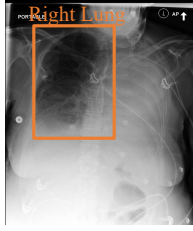
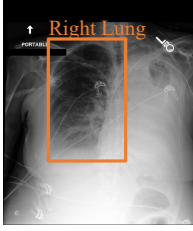
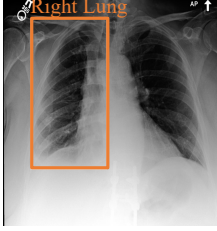
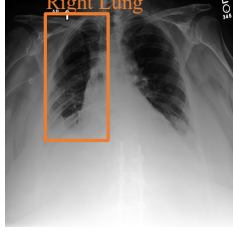
	CASE 1				CASE 2			
								
	(Prior Image)	(Current Image)			(Prior Image)	(Current Image)		
Question	(Worsen) In the chest images, what change is observed in the right lung area?				(Improved) How has the condition of the pleural effusion changed in the right lung when comparing the two CXR images?			
Options	A: The right lung shows new pulmonary infiltrates. B: There is a mild increase of pleural effusion partially obliterating the diaphragmatic contours. (✓) C: The right lung appears unchanged with minimal pleural effusion. D: The right lung is completely opacified with no visible infiltrates.				A: Right pleural effusion has increased in size. B: Right pleural effusion remains unchanged. C: Right pleural effusion has decreased in size. (✓) D: Right pleural effusion has resolved completely.			
Historical Record	As compared to the previous radiograph, the right pleural effusion has minimally increased. (regional report of right lung) Otherwise, the right lung looks unchanged. On the left, there is subtotal opacity of the hemithorax, with just a minimal apical lung portion that is ventilated. Mediastinal shift to the left persists.				___ year old woman with bilateral pleural effusions, s/p right chest tube ___. In the interval, the right chest tube was removed. The right PICC line is in stable position. Stable appearance of the cardiac silhouette. Unchanged bilateral small pleural effusions. No evidence of pneumothorax (regional report of right lung) .			
	Qwen2-VL	GPT-4o	InternVL2	InternVL2 (lora)	Qwen2-VL	GPT-4o	InternVL2	InternVL2 (lora)
	A (X)	C (X)	A (X)	C (X)	C (✓)	C (✓)	C (✓)	C (✓)
+ G-MRec	B (✓)	B (✓)	C (X)	B (✓)	C (✓)	C (✓)	B (X)	B (X)
+ R-MRec	B (✓)	B (✓)	B (✓)	B (✓)	B (X)	B (X)	B (X)	B (X)

Figure 4: Two Examples from MMXU answered by some LVLMS. The left image shows successful cases with medical records enhancing the answers. The right image shows failed cases where historical records mislead the answers.

tively mitigated this issue. In Case 2, historical records resulted in misleading predictions. When historical information was absent, all models delivered correct answers. However, the current right pleural effusion had diminished in size, and prior diagnoses indicated no changes. The inclusion of historical records resulted in an incorrect response, and fine-tuning did not yield an improvement in this case. For more case studies, please refer to Appendix E.

## 6 Conclusion

In this paper, we introduce MMXU, a dataset designed for multi-modal and multi-X-ray understanding in MedVQA. First, we propose a benchmark, MMXU-test, and invite five chest X-ray experts to evaluate the performance. Then, we conduct evaluations using several well-known open-source and closed-source large vision-language models (LVLMS) that support multi-image VQA. The experimental results indicate that even the best-performing models exhibit a significant performance gap—nearly 40%—compared to human experts. To bridge this gap, we propose the MAG method, which leverages historical records to en-

hance the understanding of disease progression. We further evaluate its performance both without fine-tuning and with fine-tuning on MMXU-dev and the experiment confirm the effectiveness.

## Limitation

Although we have carefully designed our MMXU-test benchmark and established the MMXU-dev dataset, there are still some limitations: 1) Our dataset is based on MIMIC-CXR, a chest X-ray dataset, which somewhat limits its generalization when applied to other datasets. 2) We have proposed the MAG method, inspired by clinical scenarios, to validate the effectiveness of historical historical records in enhancing LVLMS’ medical responses. Although we aim to replicate clinical scenarios as closely as possible, the scarcity of data means we can only use previous reports as historical records for research purposes.

## Acknowledgment

This work was supported by the National Natural Science Foundation of China (No. 62301311) and the Fundamental Research Funds for the Central Universities (No. YG2022QN029).

## References

- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altschmidt, Sam Altman, Shyamal Anadkat, et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katherine Millican, Malcolm Reynolds, et al. 2022. Flamingo: a visual language model for few-shot learning. *Advances in neural information processing systems*, 35:23716–23736.
- Seongsu Bae, Daeun Kyung, Jaehye Ryu, Eunbyeol Cho, Gyubok Lee, Sunjun Kweon, Jungwoo Oh, Lei Ji, Eric Chang, Tackeun Kim, et al. 2024. Ehrxqa: A multi-modal question answering dataset for electronic health records with chest x-ray images. *Advances in Neural Information Processing Systems*, 36.
- Jinze Bai, Shuai Bai, Shusheng Yang, Shijie Wang, Sinan Tan, Peng Wang, Junyang Lin, Chang Zhou, and Jingren Zhou. 2023. Qwen-vl: A versatile vision-language model for understanding, localization, text reading, and beyond. *arXiv preprint arXiv:2308.12966*, 1(2):3.
- Shenshen Bu, Yujie Song, Taiji Li, and Zhiming Dai. 2024. Dynamic knowledge prompt for chest X-ray report generation. In *LREC:2024:main*, pages 5425–5436, Torino, Italia. ELRA and ICCL.
- Junying Chen, Zhenyang Cai, Ke Ji, Xidong Wang, Wanlong Liu, Rongsheng Wang, Jianye Hou, and Benyou Wang. 2024a. Huatuoqpt-o1, towards medical complex reasoning with llms. *arXiv preprint arXiv:2412.18925*.
- 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. 2024b. [Huatuoqpt-vision, towards injecting medical visual knowledge into multimodal llms at scale](#). *Preprint*, arXiv:2406.19280.
- Zhe Chen, Jiannan Wu, Wenhai Wang, Weijie Su, Guo Chen, Sen Xing, Muyan Zhong, Qinglong Zhang, Xizhou Zhu, Lewei Lu, et al. 2024c. Internvl: Scaling up vision foundation models and aligning for generic visual-linguistic tasks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 24185–24198.
- Wenliang Dai, Junnan Li, D Li, AMH Tiong, J Zhao, W Wang, B Li, P Fung, and S Hoi. 2023. Instructblip: Towards general-purpose vision-language models with instruction tuning. *arxiv 2023. arXiv preprint arXiv:2305.06500*, 2.
- Jinlong He, Pengfei Li, Gang Liu, Zixu Zhao, and Shenjun Zhong. 2024. Pefomed: Parameter efficient fine-tuning on multimodal large language models for medical visual question answering. *arXiv preprint arXiv:2401.02797*.
- Wenyi Hong, Weihang Wang, Ming Ding, Wenmeng Yu, Qingsong Lv, Yan Wang, Yean Cheng, Shiyu Huang, Junhui Ji, Zhao Xue, et al. 2024. Cogvlm2: Visual language models for image and video understanding. *arXiv preprint arXiv:2408.16500*.
- Xinyue Hu, L Gu, Q An, M Zhang, L Liu, K Kobayashi, T Harada, R Summers, and Y Zhu. 2023. Medical-diff-vqa: a large-scale medical dataset for difference visual question answering on chest x-ray images. *PhysioNet*, 12:13.
- Yutao Hu, Tianbin Li, Quanfeng Lu, Wenqi Shao, Junjun He, Yu Qiao, and Ping Luo. 2024. Omnimed-vqa: A new large-scale comprehensive evaluation benchmark for medical lvlm. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 22170–22183.
- Xiyang Huang, Yingjie Han, Yx L, Runzhi Li, Pengcheng Wu, and Kunli Zhang. 2025. [CmEAA: Cross-modal enhancement and alignment adapter for radiology report generation](#). In *COLING:2025:main*, pages 8546–8556, Abu Dhabi, UAE. acl.
- Alistair EW Johnson, Tom J Pollard, Seth J Berkowitz, Nathaniel R Greenbaum, Matthew P Lungren, Chih-ying Deng, Roger G Mark, and Steven Horng. 2019. MIMIC-CXR, a de-identified publicly available database of chest radiographs with free-text reports. *Scientific data*, 6(1):317.
- Jason Lau, Soumya Gayen, Asma Ben, and Dina Demner-Fushman. 2018. [A dataset of clinically generated visual questions and answers about radiology images](#). *Scientific Data*, 5(1):180251.
- Hugo Laurençon, Léo Tronchon, Matthieu Cord, and Victor Sanh. 2024. What matters when building vision-language models? *arXiv preprint arXiv:2405.02246*.
- Chunyuan Li, Cliff Wong, Sheng Zhang, Naoto Usuyama, Haotian Liu, Jianwei Yang, Tristan Naumann, Hoifung Poon, and Jianfeng Gao. 2024a. Llava-med: Training a large language-and-vision assistant for biomedicine in one day. *Advances in Neural Information Processing Systems*, 36.
- Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. 2023a. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. In *International conference on machine learning*, pages 19730–19742. PMLR.
- Yanwei Li, Yuechen Zhang, Chengyao Wang, Zhisheng Zhong, Yixin Chen, Ruihang Chu, Shaoteng Liu, and Jiaya Jia. 2024b. Mini-gemini: Mining the potential of multi-modality vision language models. *arXiv preprint arXiv:2403.18814*.
- Yunxiang Li, Zihan Li, Kai Zhang, Ruilong Dan, Steve Jiang, and You Zhang. 2023b. Chatdoctor: A medical chat model fine-tuned on a large language model meta-ai (llama) using medical domain knowledge. *Cureus*, 15(6).

- 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.
- Bo Liu, Ke Zou, Liming Zhan, Zexin Lu, Xiaoyu Dong, Yidi Chen, Chengqiang Xie, Jiannong Cao, Xiao-Ming Wu, and Huazhu Fu. 2024. Gemex: A large-scale, groundable, and explainable medical vqa benchmark for chest x-ray diagnosis. *arXiv preprint arXiv:2411.16778*.
- Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. 2023. Visual instruction tuning.
- Jacek Lorkowski and Mieczyslaw Pokorski. 2022. Medical records: A historical narrative. *Biomedicines*, 10(10):2594.
- Haoyu Lu, Wen Liu, Bo Zhang, Bingxuan Wang, Kai Dong, Bo Liu, Jingxiang Sun, Tongzheng Ren, Zhuoshu Li, Hao Yang, et al. 2024. Deepseek-vl: towards real-world vision-language understanding. *arXiv preprint arXiv:2403.05525*.
- Michael Moor, Qian Huang, Shirley Wu, Michihiro Yasunaga, Yash Dalmia, Jure Leskovec, Cyril Zakkka, Eduardo Pontes Reis, and Pranav Rajpurkar. 2023. Med-flamingo: a multimodal medical few-shot learner. In *Machine Learning for Health (ML4H)*, pages 353–367. PMLR.
- Hareem Nisar, Syed Muhammad Anwar, Zhifan Jiang, Abhijeet Parida, Ramon Sanchez-Jacob, Vishwesh Nath, Holger R. Roth, and Marius George Linguraru. 2025. D-rax: Domain-specific radiologic assistant leveraging multi-modal data and expert model predictions. In *Foundation Models for General Medical AI*, pages 91–102, Cham. Springer Nature Switzerland.
- Samyam Rajbhandari, Jeff Rasley, Olatunji Ruwase, and Yuxiong He. 2020. Zero: Memory optimizations toward training trillion parameter models. In *SC20: International Conference for High Performance Computing, Networking, Storage and Analysis*, pages 1–16. IEEE.
- Nadia Saeed. 2024. [MediFact at MEDIQA-M3G 2024: Medical question answering in dermatology with multimodal learning](#). In *CLINICALNLP:2024:1*, pages 339–345, Mexico City, Mexico. acl.
- Guohao Sun, Can Qin, Huazhu Fu, Linwei Wang, and Zhiqiang Tao. 2024. [Self-training large language and vision assistant for medical question answering](#). In *EMNLP:2024:main*, pages 20052–20060, Miami, Florida, USA. acl.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. 2023. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*.
- Weihan Wang, Qingsong Lv, Wenmeng Yu, Wenyi Hong, Ji Qi, Yan Wang, Junhui Ji, Zhuoyi Yang, Lei Zhao, Xixuan Song, et al. 2023. Cogvlm: Visual expert for pretrained language models. *arXiv preprint arXiv:2311.03079*.
- Chaoyi Wu, Weixiong Lin, Xiaoman Zhang, Ya Zhang, Weidi Xie, and Yanfeng Wang. 2024a. Pmc-llama: toward building open-source language models for medicine. *Journal of the American Medical Informatics Association*, page ocae045.
- Joy T Wu, Nkechinyere N Agu, Ismini Lourentzou, Arjun Sharma, Joseph A Paguio, Jasper S Yao, Edward C Dee, William Mitchell, Satyananda Kashyap, Andrea Giovannini, et al. 2021. Chest imagenome dataset for clinical reasoning. *arXiv preprint arXiv:2108.00316*.
- Zheng Wu, Kehua Guo, Entao Luo, Tian Wang, Shoujin Wang, Yi Yang, Xiangyuan Zhu, and Rui Ding. 2024b. Medical long-tailed learning for imbalanced data: bibliometric analysis. *Computer Methods and Programs in Biomedicine*, page 108106.
- Qinghao Ye, Haiyang Xu, Jiabo Ye, Ming Yan, Anwen Hu, Haowei Liu, Qi Qian, Ji Zhang, and Fei Huang. 2024. mplug-owl2: Revolutionizing multi-modal large language model with modality collaboration. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13040–13051.
- Wenwu Ye, Jin Yao, Hui Xue, and Yi Li. 2020. [Weakly supervised lesion localization with probabilistic-cam pooling](#). *Preprint*, arXiv:2005.14480.
- Heng Yin, Shanlin Zhou, Pandong Wang, Zirui Wu, and Yongtao Hao. 2025. [KIA: Knowledge-guided implicit vision-language alignment for chest X-ray report generation](#). In *COLING:2025:main*, pages 4096–4108, Abu Dhabi, UAE. acl.
- Yuanpin Zhou and Huogen Wang. 2024. [Divide and conquer radiology report generation via observation level fine-grained pretraining and prompt tuning](#). In *EMNLP:2024:main*, pages 7597–7610, Miami, Florida, USA. acl.
- Deyao Zhu, Jun Chen, Xiaoqian Shen, Xiang Li, and Mohamed Elhoseiny. 2023. Minigpt-4: Enhancing vision-language understanding with advanced large language models. *arXiv preprint arXiv:2304.10592*.

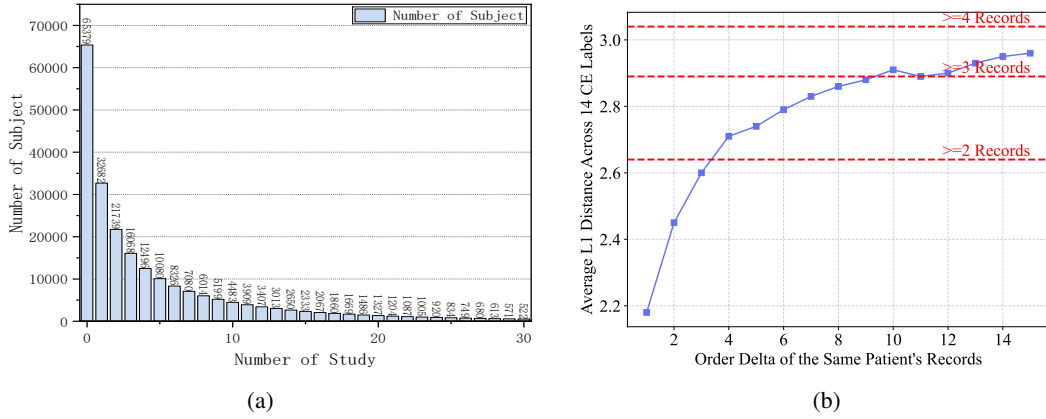


Figure 5: Figure (a) shows the distribution of the number of patients corresponding to different study counts. We can see that a large number of patients have only one study. Figure (b) shows the blue line representing the changes between two studies of the same patient in MIMIC-CXR as the study order increases. The red dashed lines indicate the differences across patients with a number of studies no less than specific values.

## A Pilot Studies on Medical Diff-VQA

### A.1 Motivation

In the clinical diagnostic process, medical experts integrate a patient’s medical historical records with current medical evidence to guide their diagnostic decisions. To explore whether this approach can enhance the diagnostic capabilities of large vision-language models in medical settings, we fine-tuned and evaluated InternVL2 8B and Qwen2-VL 8B on the publicly available Medical-Diff-VQA dataset (Hu et al., 2023), incorporating historical medical records.

Assume you are a professional thoracic doctor. You are now provided with a chest radiology image. There is also a certain probability that you will be given the patient’s previous treatment image and report. Based on the provided information, please answer the corresponding questions. The current study image is <image>. The reference image is <image>. The report of reference image is: {historical records} The question is: {question}

Figure 6: Text template on Medical-Diff-VQA using historical records.

The Medical-Diff-VQA dataset comprises seven categories of questions: abnormality, location, type, level, view, presence, and difference. However, we excluded the view questions, as they do not require historical information. Therefore, we selected

Assume you are a professional thoracic doctor. You are now provided with a chest radiology image and a question. Based on the provided information, please answer the corresponding questions. The study image is <image>. The question is: {question}

Figure 7: Text template on Medical-Diff-VQA without historical records.

six categories of questions from the Medical-Diff-VQA dataset for our study: **abnormality, location, type, level, presence, and difference**. The distribution of the training and test sets for each category is detailed in Table 9. For each category, two distinct templates were used as model inputs: one incorporating historical records and the other devoid of additional information, as illustrated in Figure 6 and Figure 7. These templates were employed to fine-tune and evaluate the model performance. It is worth noting that not all questions are associated with historical information; some are constructed based on records from a patient’s initial visit, seen in Figure 5(a). In such instances, we directly employed the template without additional information as the model input.

For historical records, we utilized the report in MIMIC-CXR dataset, which contains data from 65,079 subjects, each representing an individual patient. A subject may have multiple studies, where each study corresponds to a patient’s visit and includes multiple chest X-ray (CXR) images along



	Abnormality	Level	Location	Presence	Type	Difference
Training set	116,394	53,728	67,187	124,654	22,067	131,563
Test set	14,515	6,846	8,496	15,523	2,702	16,389

Table 9: The number of training and test samples for the six categories of questions in Medical-Diff-VQA.




Figure 8: Visualization of experimental results for different settings on the Medical-Diff-VQA dataset.

with an associated medical report. As shown in Figure 5(b), the blue line shows the average L1 distance variation between the Chexpert (Ye et al., 2020) labels of different studies within the same subject. The red dashed line indicates the average L1 distance between patients who meet certain conditions in terms of the number of studies. It is evident that as the study order increases, the differences between studies grow significantly. Furthermore, under the same conditions, the differences between patients are much greater than those within a single patient. Based on these observations, we **selected the most recent frontal CXR and report** as the historical records for our study.

## A.2 Experimental

**Experimental Setting** We benchmarked our approach against the DRax method (Nisar et al., 2025), employing distinct evaluation metrics: **token recall** for open-ended questions and **accuracy** for close-ended ones. The models utilized in the experiments were InternVL2 8B and Qwen2-VL 8B. During the fine-tuning process, the LoRA rank was set to 16, the learning rate to  $4e-5$ , and a cosine learning rate scheduling strategy was implemented. Training was conducted using the DeepSpeed Zero Stage 2 (Rajbhandari et al., 2020) distributed strategy on a 8-GPU RTX 3090 server.

**Experimental Results** Figure 8 shows the vi-

sualization of experimental results on 5 types of single-image questions (including closed-ended questions and open-ended questions). The results clearly indicate that integrating historical records as references significantly enhances VQA performance on both Qwen2-VL 8B and InternVL2 8B.

For the sixth category, the difference questions, we also investigated the effect of providing reference images along with historical records on the outcomes. These results are shown in Table 10. The responses to “difference” questions consist primarily of two components: the 14 ChexPert labels (Ye et al., 2020) that are either missing or added in the main image compared to the reference. The “Missed Labels” and “Added Labels” sections of the table employ accuracy, recall, and F1 scores for the quantitative assessment of the model outputs. Our findings indicate that the inclusion of historical records resulted in a substantial improvement across all evaluation metrics. (The 14 clinical efficacy labels in chexpert include: atelectasis, cardiomegaly, consolidation, edema, enlarged cardio-mediastinum, fracture, lung lesion, lung opacity, no finding, pleural effusion, pleural other, pneumonia, pneumothorax, support devices.)

However, the Medical-Diff-VQA dataset does not fully capture the nuanced relationship between these questions and the evolution of historical historical records. This shortcoming complicates the determination of whether the model is truly leveraging historical information or simply reiterating prior data. Such behavior falls short of the expectations for real-world clinical applications. As a result, **there is an urgent need for a dataset that seamlessly integrates and balances historical medical records**, thereby supporting their meaningful utilization in practical healthcare settings.

## B Detailed Data Case

Figure 9 illustrates a detailed example from the MMXU dataset we developed, which is specifically tailored to analyze changes between two historical records of a patient related to chest X-rays. Each sample includes a single-choice question focused on changes in the condition of a particular

Method	NLG metrics			Missed Labels			Added Labels			Average
	BLUE-1	Token Recall	Token F1	ML Acc	ML Recall	ML F1	AL Acc	AL Recall	AL F1	
w/ MedRecord	0.464435	0.614229	0.59154	0.884127	0.267652	0.259081	0.885768	0.283076	0.275101	0.502779
w/o MedRecord	<b>0.550117</b>	<b>0.684575</b>	<b>0.674029</b>	<b>0.947046</b>	<b>0.629887</b>	<b>0.615927</b>	<b>0.901559</b>	<b>0.345302</b>	<b>0.334627</b>	<b>0.631452</b>

Table 10: Evaluation results for the differences category in Medical-Diff-VQA using different methods.

```
{
  "qid": 53804,
  "question": "Considering the changes observed, what can be concluded about the left lung in the recent CXR images?",
  "question_type": "single choice",
  "content_type": "worsen",
  "options": {
    "A": "Worsened with increased effusion.",
    "B": "Stable with no change.",
    "C": "Improved or resolving condition.",
    "D": "Complete resolution."
  },
  "answer": "A",
  "region_name": "left lung",
  "current_image_path": "files/p19/p19839145/s55247427/79b946ab-8d938e59-52027cad-8b4a4268-dab951d0.jpg",
  "regional_current_bbox": [1637, 122, 2796, 2196],
  "history_image_path": "files/p19/p19839145/s59400960/36b4f554-95744e28-8b7baa81-fa621df8-41ad666f.jpg",
  "regional_history_bbox": [1514, 463, 2442, 2278],
  "regional_history_report": "Opacities in the left upper lobe have markedly improved. Moderate bilateral effusions larger on the left side with adjacent atelectasis are grossly unchanged. There is no pneumothorax.",
  "global_history_report": "Mild cardiomegaly is stable. The aorta is tortuous. Cardiomeidiastinum is shifted to the right as before. Moderate bilateral effusions larger on the left side with adjacent atelectasis are grossly unchanged. Opacities in the left upper lobe have markedly improved. There is no pneumothorax."
}
```

Figure 9: A detailed data case sampled from the MMXU dataset.

region, along with comprehensive records from the patient’s two visits. For the initial visit, the dataset provides images, bounding boxes marking the relevant regions, region-level reports, and overall-level reports. The questions are divided into three categories: “Worsened”, “Improved”, and “No Change”, based on the comparative statements within the raw report. This dataset serves as a valuable resource for advancing research in multi-image comparison, historical record integration, and the application of visual grounding.

## C Detailed of Dataset Construction

In this section, we will introduce some additional information during the construction process.

### C.1 Comparative Sentences Extraction

Figure 10 presents a detailed example of a comparative sentence extracted during the dataset construction process. The example includes the compara-

tive sentence itself, the specific regions it pertains to, and the two corresponding images. Additionally, the bounding boxes of the relevant regions within the images, along with the comparative relationship, are illustrated. This detailed data structure is crucial for building a robust dataset that facilitates accurate comparisons during QA generation.

### C.2 QA Pairs Generation

The version of GPT-4o we utilized is gpt-4o-2024-08-06. We employ the prompt template shown in Figure 11, where the “explanatory information” is derived from the details associated with the *relationship* label. Figure 16 illustrates a detailed example of the input prompt and output content for the question-answer pair model using GPT-4o. We provided six specific rules and supplied GPT-4o with comparative sentences, *relationship* labels, reports from two visits, and the region name of interest. Notably, the answers generated by GPT-

```

{
  "cmp_id": 3,
  "subject_id": "14731346",
  "study_id": "54684841",
  "cur_image_path": "files/p14/p14731346/s54684841/56ec2474-cfdbbdcd-5b133f15-6c0bc409-17435e
c2.jpg",
  "prior_image_path": "files/p14/p14731346/s51369333/33585d85-edf829eb-51303048-7bd9062a-b6ee3
a5b.jpg",
  "cur_report": "Right lower quadrant pain. In comparison with the study of ___, there has been
worsening of the increased opacification at the left base with silhouetting of the hemidiaphragm
and blunting of the costophrenic angle. These findings are consistent with a combination of
volume loss in the left lower lobe and pleural effusion. Right lung is clear, and there is no
evidence of vascular congestion.",
  "prior_report": "Hyperglycemia, intubated for airway protection, please assess NG tube and
ET tube placement. AP radiograph of the chest was compared to prior study obtained the same
day earlier. The ET tube tip is approximately 4.3 cm above the carina and slightly impinging
the right wall of the trachea and should be repositioned. The NG tube tip is in the stomach.
Right subclavian line tip is at the level of cavoatrial junction. Heart size and mediastinum
are unremarkable. There is no pneumothorax. Minimal bibasilar opacities most likely reflect
areas of atelectasis. No pulmonary edema is seen.",
  "comp_sent": "FINDINGS: In comparison with the study of ___, there has been worsening of the
increased opacification at the left base with silhouetting of the hemidiaphragm and blunting of
the costophrenic angle.",
  "relationships": {
    "left lung": ["comparison|yes|worsened"],
    "left lower lung zone": ["comparison|yes|worsened"],
    "left costophrenic angle": ["comparison|yes|worsened"],
    "left hemidiaphragm": ["comparison|yes|worsened"]
  },
  "related_region_names": ["left lung", "left lower lung zone", "left costophrenic angle",
"left hemidiaphragm"],
  "cur_image_size": [2544, 3056],
  "cur_image_bboxes": {
    "left lung": [1186, 777, 2087, 2210],
    "left lower lung zone": [1241, 1705, 2087, 2210],
    "left costophrenic angle": [1719, 1719, 1991, 1991],
    "left hemidiaphragm": [1186, 1691, 2605, 2769]
  },
  "prior_image_size": [[2544, 3056]],
  "prior_image_bboxes": {
    "left lung": [[1268, 327, 2537, 2878]],
    "left lower lung zone": [[1487, 1705, 2537, 2878]],
    "left costophrenic angle": [[[]],
    "left hemidiaphragm": [[[]]
  }
}

```

Figure 10: A detailed data case about **comparative sentence** during dataset construction.

4o are not limited to simple responses regarding improvements, deteriorations, or stability. Instead, we emphasize capturing more detailed changes in the regions.

## D Evaluation on MMXU-test benchmark

**Experimental Settings** We respectively utilized specialized prompt templates for three distinct scenarios: direct evaluation, enhancement through the integration of global and regional historical reports as historical records. For all generation processes, we set do\_sample to False and temperature to 0. During the InternVL2 8B fine-tuning pro-

cess, we set epochs to 1, max\_dynamic\_patch to 6, down\_sample\_ratio to 0.5, lora\_rank to 16, and the learning\_rate to 4e-5. The learning rate scheduler type is set to "cosine," with a warmup ratio of 0.03. The fine-tuning is carried out using the DeepSpeed Zero Stage 2 (Rajbhandari et al., 2020) distributed strategy on an 8-GPU NVIDIA RTX 3090 server.

**Text Prompt** To assess the impact of historical records on medical diagnosis, we performed evaluations on our MMXU-test dataset under three distinct conditions: without historical records, with global reports as historical records, and with regional reports as historical records. Figure 14 il-

You are a chest x-ray assistant and you are presented with a sentence involving a comparison in the original and current reports of a patient's two visits before and after, as well as some additional related explanatory information about the sentence. Please generate up to 3 single choice questions about the changes in the images from the two visits, based primarily on the sentence involving the comparison and the additional information, with the aid of referencing the two reports. Format the output into JSON format.

---

**Here are some rules:**

- (1) Your questions are generated primarily around the KEY COMPARISON SENTENCE IN THE CURRENT REPORT and related explanatory information, with the content of the two reports as a secondary aid.
- (2) Your question should be whether the chest CXR images have changed in some way in the same area, including abnormality, disease, location, severity, etiology, etc.
- (3) For each question, you should clearly identify the area in which the question is asked and provide the correct answer, along with an explanation of why that answer was chosen.
- (4) Your question is intended to answer the case where only images of the two diagnoses are known. Therefore, please do not include the words report comparison in your question; what should be included is image comparison.
- (5) The JSON in related explanatory information consists of a "region" and a "list of explanatory information". For example, "right lung": ["comparison|yes|improved"] indicates that the right lung region has improved, and "left lung": ["comparison|yes|worsen"] indicates that the left lung region has worsened.
- (6) Avoid asking questions that do not involve a change in two pictures. Ensuring all the questions are as diverse as possible.

---

**Here is one example:**

{...}

---

**Here is some information:**

{comparison sentence}, {explanatory information}, {current report}, {prior report}

Figure 11: Text prompt for generating multiple image-based question-answer pairs in MMXU using GPT-4o

Image-1: <image>

Image-2: <image>

Here are two chest X-RAY images of the same patient. Previous chest X-RAY is Image-1. And current chest X-RAY is Image-2. Your task is to evaluate the differences between the two images based on the provided report and question.

The report of Image-1 about is:

%s

Question:

%s

Options:

A: %s

B: %s

C: %s

D: %s

Answer with the option's letter from the given choices directly.

Figure 12: Text prompt template for evaluation on MMXU-test benchmark with full report as historical records

illustrates the prompt template used in the absence of historical records, where only the basic instructions, questions, and options are provided. Figure 12 depicts the prompt template incorporating global reports as historical records, in which we enhance the original template by including the global report.

Figure 13 showcases the prompt template with regional reports as historical records. In contrast to the global reports, these regional reports are tailored to the specific region relevant to the question, with the region's name explicitly indicated.

When generating region-level reports as histori-



```

Image-1: <image>
Image-2: <image>
Here are two chest X-RAY images of the same patient. Previous chest X-RAY is Image-1. And
current chest X-RAY is Image-2. Your task is to evaluate the differences between the two images
based on the provided regional report and question.

The report of Image-1 about region %s is:
%s

Question:
%s

Options:
A: %s
B: %s
C: %s
D: %s

Answer with the option's letter from the given choices directly.

```

Figure 13: Text prompt template for evaluation on MMXU-*test* benchmark with regional report as historical records

```

Image-1: <image>
Image-2: <image>
Here are two chest X-RAY images of the same patient. Previous chest X-RAY is Image-1. And
current chest X-RAY is Image-2. Your task is to evaluate the differences between the two images
based on the provided question.

Question:
%s

Options:
A: %s
B: %s
C: %s
D: %s

Answer with the option's letter from the given choices directly.

```

Figure 14: Text prompt template for direct evaluation on MMXU-*test* benchmark without historical records

cal record enhancements, we distinguish regions using the 29 anatomical names provided in the Chest ImaGenome dataset, including: right lung, right upper lung zone, right mid lung zone, right lower lung zone, right hilar structures, right apical zone, right costophrenic angle, right hemidiaphragm, left lung, left upper lung zone, left mid lung zone, left lower lung zone, left hilar structures, left apical zone, left costophrenic angle, left hemidiaphragm, trachea, spine, right clavicle, left clavicle, aortic arch, mediastinum, upper mediastinum, svc, cardiac silhouette, cavoatrial junction, right atrium, carina, and abdomen.

## E More Case Study

Figure 15 shows two challenging examples from the "Worsen" category in MMXU. In CASE 1, none

of the models without fine-tuning were able to correctly answer the question, and only the fine-tuned model using the MAG method provided the correct answer. This may be due to the difficulty of the question itself, as well as the challenge of extracting useful information to answer the question from the historical records. In CASE 2, only GPT-4o was able to answer the question correctly after providing historical records. For other models, since R-MRec reiterated previously unchanged views, all incorrectly chose option A (no change), indicating that these models failed to correctly understand the change and simply repeated the previous information.

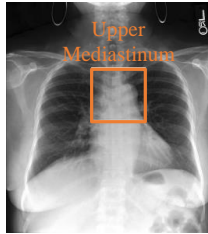
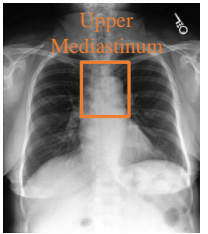


	CASE 1				CASE 2			
	 (Prior Image)		 (Current Image)		 (Prior Image)		 (Current Image)	
Question	(Worsen) In terms of the upper mediastinum, how have the findings changed between the prior and current chest x-ray images?				(Worsen) What change is observed in the right lung region when comparing the chest CXR images from the two visits?			
Options	A: Findings have improved. B: Findings remain stable. C: Findings have worsened. D: The region is not mentioned in either report.				A: No change in the right lung. B: Improvement in the right lung. C: Worsening in the right lung. (✓) D: New consolidation.			
Historical Record	Lungs are well expanded and clear. There is no pleural abnormality. Heart size is normal. There is greater fullness in the right lower paratracheal station of the mediastinum. If there is lymph node enlargement, there is no compromise of the airway, and therefore adenopathy would not explain shortness of breath.				Cardiomediastinal contours are within normal limits and without change. Minimal bibasilar atelectasis is present, but there are no new areas of consolidation to suggest the presence of a new site of pneumonia. Nasogastric tube continues to terminate in the stomach, but side port is in close proximity to the gastroesophageal junction.			
	Qwen2-VL	GPT-4o	InternVL2	InternVL2 (lora)	Qwen2-VL	GPT-4o	InternVL2	InternVL2 (lora)
+ G-MRec	B (✗)	B (✗)	B (✗)	B (✗)	C (✓)	B (✗)	D (✗)	C (✓)
+ R-MRec	B (✗)	D (✗)	B (✗)	C (✓)	C (✓)	C (✓)	A (✗)	A (✗)
	B (✗)	B (✗)	B (✗)	C (✓)	A (✗)	C (✓)	A (✗)	A (✗)

Figure 15: Two hard examples from “Worsen” category of MMXU answered by some LVLMS.

**USER:** You are a chest x-ray assistant and you are presented with a sentence involving a comparison in the original and current reports of a patient's two visits before and after, as well as some additional related explanatory information about the sentence. Please generate at most 3 single choice questions about the changes in the images from the two visits, based primarily on the sentence involving the comparison and the additional information, with the aid of referencing the two reports. Format the output into JSON format.

Here are some rules:

- (1) Your questions are generated primarily around the KEY COMPARISON SENTENCE IN THE CURRENT REPORT and related explanatory information, with the content of the two reports as a secondary aid.
- (2) Your question should be whether the chest CXR images have changed in some way in the same area, including abnormality, disease, location, severity, etiology, etc.
- (3) For each question, you should clearly identify the area in which the question is asked and provide the correct answer, along with an explanation of why that answer was chosen.
- (4) Your question is intended to answer the case where only images of the two diagnoses are known. Therefore, please do not include the words report comparison in your question; what should be included is image comparison.
- (5) The JSON in related explanatory information consists of a "region" and a "list of explanatory information". For example, "right lung": ["comparison|yes|improved"] indicates that the right lung region has improved, and "left lung": ["comparison|yes|worsen"] indicates that the left lung region has worsened.
- (6) Avoid asking questions that do not involve a change in two pictures. Ensuring all the questions are as diverse as possible.

Here is one example:

```
{
  "question": ...,
  "region": ...,
  "content_type": ..., # worsen, improved or no change
  "options": {
    "A": ...,
    "B": ...,
    "C": ...,
    "D": ... }
  "answer": ...,
  "reason": ...
}
```

Here is some information:

KEY COMPARISON SENTENCE IN THE CURRENT REPORT: "Widening mediastinum has improved."

Explanatory Information: "upper mediastinum": ["comparison|yes|improved"]

Current Report: Moderate to severe cardiomegaly is stable. Widening mediastinum has improved. Bilateral effusions are small. Opacities in the lower lobes bilaterally larger on the left side are likely atelectasis, superimposed infection cannot be totally excluded. There is no evident pneumothorax. Sternal wires are aligned. Degenerative changes in the thoracic spine are mild. There is mild vascular congestion. Bilateral healed rib fractures are again noted. Small bilateral effusions with adjacent opacities, these are likely atelectasis but superimposed infection cannot be totally excluded. Mild vascular congestion.

Prior Report: IN COMPARISON TO \_\_\_ RADIOGRAPH, CARDIOMEGALY IS ACCOMPANIED BY IMPROVING PULMONARY EDEMA. A MORE CONFLUENT OPACITY IN THE LEFT RETROCARDIAC REGION HAS WORSENER, AND MAY REFLECT A COMBINATION OF PLEURAL effusion and adjacent atelectasis or consolidation. No other relevant change.

---

**GPT-4o:**

```
{
  "question": "What change has been observed in the upper mediastinum region of the chest CXR images?",
  "region": "upper mediastinum",
  "content_type": "improved",
  "options": {
    "A": "The widening has worsened.",
    "B": "There is new vasculature congestion.",
    "C": "The widening has improved.",
    "D": "There is no notable change."
  },
  "answer": "C",
  "reason": "The related explanatory information indicates that there is an improvement in the upper mediastinum, corresponding to the general improvement of the mediastinum's widening mentioned in the key comparison sentence."
}
```

Figure 16: A detailed example of the input and output content for the question-answer pair model using GPT-4o