A STRUCTURED, TAGGED, AND LOCALIZED VISUAL QUESTION ANSWERING DATASET WITH FULL SENTENCE ANSWERS AND SCENE GRAPHS FOR CHEST X-RAY IMAGES

Anonymous authors

Paper under double-blind review

ABSTRACT

Visual Question Answering (VQA) enables targeted and context-dependent analysis of medical images, such as chest X-rays (CXRs). However, existing VQA datasets for CXRs are typically constrained by simplistic and brief answer formats, lacking localization annotations (e.g., bounding boxes) and structured tags (e.g., region or radiological finding/disease tags). To address these limitations, we introduce *MIMIC-Ext-CXR-QBA* (abbr. *CXR-QBA*), a large-scale CXR VQA dataset derived from MIMIC-CXR, comprising 42 million QA-pairs with multi-granular, multi-part answers, detailed bounding boxes, and structured tags. We automatically generated our VQA dataset from scene graphs (also made available), which we constructed using LLM-based information extraction from radiology reports. After automatic quality assessment, we identified 31M pre-training and 7.5M fine-tuning grade QA-pairs, providing the largest and most sophisticated VQA dataset for CXRs to date. Tools for using our dataset and the construction pipeline are available at https://anonymous.4open.science/r/mimic-ext-cxr-qba/.

1 Introduction

With the emergence of Large Language Models (LLMs) and Large Multimodal Models (LMMs), interactive and conversational tasks have gained popularity in medical image analysis, particularly in the context of *chest X-ray* (*CXR*) interpretation (Chen et al., 2024; Müller et al., 2025; Tu et al., 2024; Xie et al., 2025). A prominent example of such interactive tasks is *Visual Question Answering* (*VQA*), where a model is presented with an image and a corresponding textual question, and is tasked with generating an answer. Unlike conventional medical imaging approaches, which always produce the same output (such as classification labels, bounding boxes, or textual reports) for a given image, VQA enables users to interactively explore and interpret images in a context-dependent manner. Training robust VQA models for medical applications requires high-quality, large-scale training datasets. Existing CXR VQA datasets suffer from several limitations: (i) they often contain only short and simplistic answers, (ii) they lack localization information (such as bounding boxes), and (iii) they provide little structured metadata (e.g., region and finding/disease annotations, or uncertainty estimates). Additionally, their relatively small size constrains their utility for pretraining.

To address these challenges, we propose a pipeline for automatic VQA dataset creation and apply it to construct a new large-scale CXR VQA dataset. Unlike prior datasets, each question-answer (QA) pair includes multi-granular, multi-part answers composed of full sentences in the style of radiology reports. Furthermore, our dataset provides detailed bounding boxes and additional structured tags (e.g., findings and regions), enhancing interpretability and facilitating the development of more advanced and transparent medical VQA models. Fig. 1 shows examples of our generated QA-pairs.

Our contributions are as follows:

• We propose a dataset construction pipeline, co-designed with radiologists, for the automatic generation of *large-scale medical VQA datasets*, offering structured, tagged, and localized answers not yet provided in existing datasets.

060 061 062

063 064 065

066 067

068

071

073

074

075

076

077 078 079

081

082

084

085

090

092

094

095

096

098

100

101

102

103

104

105

106

107



(a) Indication question.

(b) Study abnormality question.



(c) Region abnormality question.

(d) Finding question.

Figure 1: Examples of question-answer (QA) pairs for each of our four different types of questions. For each question (for a given chest X-ray), a detailed answer with sentences in the style of free-text radiology reports is given, supplemented by bounding boxes (for both positive and negative answers), and a set of tags (e.g. regions, findings, certainty, etc.). For more examples, we refer to Sec. A.

- To provide a reliable data source for this pipeline, we develop an automatic scene-graph construction method, leveraging LLMs, semantic entity mapping, and localization models.
- Building on this, we introduce *MIMIC-Ext-CXR-QBA* (abbr. *CXR-QBA*), a CXR VQA dataset of 42M question–answer pairs derived from MIMIC-CXR (Johnson et al., 2024b).
- We automatically assess the quality of the generated QA-pairs, identifying 31.2M pairs as pre-training grade and 7.5M of these as fine-tuning grade.
- We perform extensive quality controls, including validations against expert annotations and assessment of the automatic ratings, ensuring the robustness of our pipeline.
- We introduce the *structured VQA* task for CXRs and provide a baseline model to demonstrate our dataset's utility and guide future model development.

2 RELATED WORK

VQA Datasets for Chest X-Rays VQA datasets (shown in Tab. 1) are scarce in the medical imaging domain, with most notable examples being VQA-RAD (Lau et al., 2018) and SLAKE (Liu et al., 2021), which are hand-labeled but limited in size. On the other hand, VQA-Med at ImageCLEF 2019 (Ben Abacha et al., 2019) was automatically constructed using QA templates based on image annotations, which may limit its answer quality. To improve the quality, PMC-VQA (Zhang et al., 2024) used an LLM to generated QA-pairs based on provided captions. VQA datasets for chest X-rays include MIMIC-Ext-MIMIC-CXR-VQA (Bae et al., 2024) and Medical-CXR-VQA (Hu et al., 2025; 2024), which contain hundreds of thousands of QApairs, derived from MIMIC-CXR. These datasets rely on templates but use radiology reports as their original

Table 1: Medical VQA dataset comparison. Our dataset is the largest to date and provides localized and tagged answers.

Dataset	#QA	Boxes	Tags	Answers
CXR-QBA (Ours) 4 fine-tuning grade	42.2M 7.5M	1	1	detailed
VQA-RAD	3.5K	Х	х	brief
SLAKE	14K	X	X	brief
ImageCLEF	15K	X	X	brief
PMC-VQA	227K	X	X	brief
MIMIC-CXR-VQA	377K	X	X	brief
Medical-CXR-VQA	780K	X	X	brief
CheXinstruct	8.5M	X	X	brief

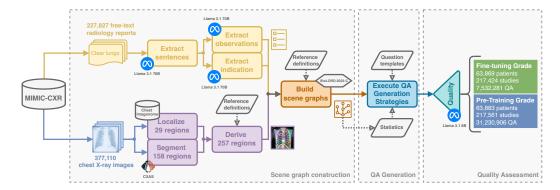


Figure 2: Overview of our dataset construction pipeline. First, we construct scene graphs based on information extracted from the radiology reports and regions localized in the images. Next, we generate question-answer pairs based on templates and the scene graphs. Finally, we automatically assess the quality of generated QA-pairs.

information source, where MIMIC-Ext-MIMIC-CXR-VQA leverages Chest ImaGenome's (Wu et al., 2021a) scene graphs and Medical-CXR-VQA employs an LLM-based extraction strategy similar to ours but without semantic entity mapping, localization, and extraction of textual descriptions. The largest chest X-ray VQA dataset to date, CheXinstruct (Chen et al., 2024), contains 8.5M QA-pairs with images from multiple data sources. However, compared to our dataset, its questions and answers are less diverse, being purely template-based and derived from dataset annotations instead of being directly conditioned on the reports. Additionally, none of the described datasets provide the level of detail and annotation richness found in our dataset, which includes bounding boxes, tags, and more detailed, multi-part answers that mirror radiology report sentences.

Grounded Report Generation While localization is not yet common for medical VQA tasks, grounded report generation, i.e. predicting radiology reports with bounding boxes is gaining popularity. Notable examples include MAIRA-2 (Bannur et al., 2024), trained on reports manually annotated with bounding boxes and MedTrinity-25M (Xie et al., 2025), a large-scale public dataset with automatically generated reports with bounding boxes. ChEX (Müller et al., 2025) is another model producing textual answers with bounding boxes but does not support VQA tasks.

Scene Graph Construction for Chest X-Rays During our VQA dataset construction, we automatically derive scene graphs from radiology reports. A similar approach is employed by Chest ImaGenome (Wu et al., 2021a;b), which uses rule-based information extraction, and RadGraph (Jain et al., 2021), which uses a relation extraction model. We do not rely on any of those scene graphs but instead leverage LLM-based extraction with semantic entity mapping to extract our own more comprehensive and reliable scene graphs. Notably, our method defines a larger set of (localized) regions (257) and findings (221) compared to Chest ImaGenome (29 regions, 53 findings), while additionally rewriting report sentences to focus on specific aspects related to individual graph nodes.

3 THE CXR-QBA DATASET

We present our dataset *CXR-QBA*, a large-scale *chest X-ray* (*CXR*) VQA dataset derived from *MIMIC-CXR* (Johnson et al., 2024b; 2019a), consisting of more than 42M QA-pairs. As shown in Fig. 1, each QA sample (for a given chest X-ray) consists of a *question* (*Q*), a *bounding box* (*B*) supplemented *answer* (*A*), and additional tags (regions, findings, certainties, and more). To build our dataset, we propose an automatic pipeline highlighted in Fig. 2. More specifically, we first construct (visually grounded) scene graphs based on the MIMIC-CXR radiology reports using LLM-based information extraction, semantic concept mapping, and localization models (Sec. 3.1). These scene graphs provide a structured description of the study, including sentences (derived from the report) for individual observations. They serve as a data source for our question-answer generation, where we utilize both template-based answers and answers derived from the rewritten report sentences (Sec. 3.2). Finally, we automatically assess the quality of question-answer pairs using LLM-based evaluations (Sec. 3.3).

3.1 Scene Graph Construction

Given a MIMIC-CXR study with a radiology report and accompanying CXRs, we construct a scene graph (Sec. D.1) consisting of sentence nodes, observation nodes, region nodes, an indication node, and edges between them. Sentence nodes are directly extracted from the reports, containing the raw sentences and their identified section names. Observation nodes represent individual aspects described in the report's FINDINGS or IMPRESSION section, containing (i) a textual description, (ii) bounding boxes for associated CXR images and (iii) additional tags, such as positivity, certainty, laterality, regions, and finding classes. Region nodes are created for mentioned anatomical structures and key regions. The indication node contains information from the INDICATION section, including a textual description and an individual observation node, derived from the FINDINGS and IMPRESSION sections, that can act as an answer to the indication. We construct these scene graphs in three steps: (a) region localization, (b) information extraction and (c) building the graphs using entity mapping. We refer to Sec. E.1 for details.

Region Localization The bounding boxes in our scene graphs (and the derived QA-pairs) are based on fine-grained anatomical structures, allowing us to localize associated findings very precisely. We use the CXAS (Seibold et al., 2022; 2023) model to predict segmentation masks of 158 anatomical structures on the 377 110 CXRs from MIMIC-CXR-JPG (Johnson et al., 2024a; 2019b). Additionally, we use the bounding boxes provided by the Chest ImaGenome (Wu et al., 2021a;b) dataset, which are provided for 29 anatomical structures in most frontal images of MIMIC-CXR. Next, we derive a total of 257 localized anatomical structures based on combinations (e.g. intersections, unions, super bounding boxes, etc.) of the available masks and bounding boxes. Finally, we discard any masks or boxes that are too small and derive bounding boxes from the segmentation masks. Note that we define 53 further regions/structures that are either non-localized (e.g. interstitial) or for which we do not have bounding boxes, leading to a total of 310 structures/regions.

Information Extraction We use the 227 827 free-text radiology reports provided by MIMIC-CXR as the main source of information for our scene graphs. Using the Llama 3.1 70B (Grattafiori et al., 2024) model with few-shot prompting, we extract the relevant information (tags and textual descriptions) in three steps. First, we extract individual sentences from the reports and detect their sections. Next, we extract information about the INDICATION section and detect which FINDINGS or IMPRESSION sentences may provide information related to the indication. Finally, we extract individual observations described in the FINDING/IMPRESSION sentences.

Building Scene Graphs using Entity Mapping Given the extracted information from the reports and the computed bounding boxes, we now construct the final scene graph. Therefore, we first map extracted tags to pre-defined sets of values, our reference definitions. This assures high quality and consistency of the scene graphs and enables mapping of observations to the extracted bounding boxes. The reference definitions are based on tags used in other datasets (including PadChest (Bustos et al., 2020) and Chest ImaGenome (Wu et al., 2021b)) as well as SNOMED-CT (SNOMED International, 2023)) and have been verified by clinical experts. They include synonym lists, hierarchies, and relationships. For more robust mapping, we utilize the BioLORD (Remy et al., 2024) model as a sentence transformer and identify the closest matching concept based on their semantic embeddings. Additionally, we try to fill in missing information where possible, such as inferring the region from an identified finding. Finally, we build a tree of region nodes (using the reference data) and attach the indication information extracted from the report.

3.2 QUESTION-ANSWER GENERATION

Based on the information available in the scene graphs, we generate question-answer pairs (Sec. D.2) using a template-based approach co-designed with radiologists. We incorporate the textual descriptions from the observation nodes – which have been derived directly from the report – to provide diverse and fine-grained answers. Each answer may consist of multiple answer parts (as shown in Fig. 3 and Sec. A), each describing an individual aspect of the answer with its own sentence, bounding boxes, and tags. We categorize answer parts into three types: (i) main-answers, (ii) details, and (iii) related-information, allowing for controlled answer granularity. Answer parts are generated either from templates using scene graph information or directly from observation nodes (Sec. E.2). Answer parts may also be structured hierarchically, where we use parent-child edges from the scene graph.

To generate the question-answer pairs, we employ different strategies for the four types of questions (shown in Fig. 1):

- 1. **Indication**: We use the paraphrased indication as the question and create the answer based on the indication node in the scene graph, answering the indication based on information in the FINDINGS and IMPRESSION sections.
- 2. **Study abnormality**: We generate study-level questions using 13 different templates, with answer parts (Fig. 3) based on (filtered) observation nodes.
- 3. **Region abnormality**: We generate questions about individual regions using 6 different templates, considering any region mentioned and additionally randomly sampling non-mentioned regions for balancing.
- 4. Finding: We generate questions about individual findings using 7 different templates, considering any finding mentioned and additionally randomly sampling non-mentioned findings for balancing.

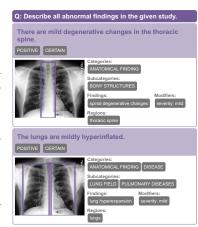


Figure 3: Answer with multiple parts for different aspects, each with a sentence, tags, and boxes.

3.3 QUALITY ASSESSMENT

The dataset construction procedure described so far allows us to automatically generate large amounts of QA-pairs. However, in each of the steps, errors may be introduced, affecting the overall quality of the dataset. For example, errors during information extraction could lead to incorrect tags, therefore leading to incorrectly filled answer templates or incorrectly selected observation nodes for answers.

To identify and filter such cases, we employ an automatic quality assessment strategy using an LLM-as-a-judge. More specifically, we use Llama 3.1 8B (Grattafiori et al., 2024) to rate questions and answers by the following five criteria: *entailment* (does the answer factually align with the original report?), *relevance* (is the answer relevant to the question?), *completeness* (is the answer missing something?), as well as *question and answer clarity* (is the question/answer clear and grammatically correct?). Additionally, we assess the quality of the used scene graphs by identifying missing information (e.g. missing tags or localization) or issues during the construction process. Finally, we combine these individual assessments to compute an overall quality rating as one of A++, A+, A+,

4 EVALUATION AND ANALYSIS

4.1 EVALUATION OF THE SCENE GRAPHS

We evaluate our scene graphs by comparing their tags and bounding boxes to hand-labeled expert annotations on MIMIC-CXR, using the scene graphs from Chest ImaGenome (Wu et al., 2021a) as a baseline. First, we evaluate the plausibility of finding tags by comparing study-level labels derived from our scene graphs to two reference annotation sets: the radiologist annotations in MIMIC-CXR-JPG v.2.1.0 (Johnson et al., 2024a) with 13 CheXpert (Irvin et al., 2019) classes and the CXR-LT 2024 (Holste et al., 2025; 2024) gold-standard dataset (task 2 test set) with 12 additional rare (long-tail) classes. As shown in Tab. 2a, our approach (slightly) outperforms Chest ImaGenome, with strong improvements (20%) on long-tail classes, demonstrating the value of our fine-grained finding tags (221 classes) in capturing nuanced study details. To evaluate the accuracy of finding bounding boxes, we compare them with annotations from MS-CXR (Boecking et al., 2024; 2022) (6 classes) and REFLACX (Bigolin Lanfredi et al., 2021; 2022) (18 classes). We compute study-level pixel masks for each finding as the union of all bounding boxes from positive observation nodes that contain

Table 2: Evaluation of our scene graphs, comparing finding tags (a) and associated bounding boxes (b) to expert annotations on MIMIC-CXR subsets, with 95% confidence intervals (bootstrapping, n = 1000). Compared to Chest ImaGenome's (Wu et al., 2021a) scene graphs, we achieve competitive or superior performance, showing that our construction process yields plausible scene graphs.

(a) Evaluation of finding tags against 13 CheXpert (CXP) classes from the MIMIC-CXR-JPG test set and 25 classes, 13 CXP and 12 long-tail (LT) classes, from the CXR-LT 2024 gold standard dataset (Sec. C.2). We report the Matthews Correlation Coefficient (MCC) macro-averaged over different finding subsets (CXP-5, CXP-7, CXP-13, LT) and micro-averaged. Compared to Chest ImaGenome, we produce slightly more accurate tags, performing especially well on long-tail classes, highlighting the importance of our fine-grained tags.

	MIMIC-CXR-JPG Test [MCC]					CXR-L	Γ 2024 Gold	[MCC]	
Classes	CXP-5	CXP-7	CXP-13	Micro	CXP-7	CXP-13	LT-only	CXR-LT	Micro
Ours (scene graphs)	0.8 [0.77, 0.82]	0.81 [0.79, 0.84]	0.69 [0.67, 0.71]	0.71 [0.69, 0.73]	0.65 [0.61, 0.69]	0.57 [0.54, 0.6]	0.71 [0.67, 0.74]	0.64 [0.61, 0.66]	0.67 [0.65, 0.69]
Chest ImaGenome	0.78 [0.75, 0.81]	0.8 [0.78, 0.83]	0.66 [0.64, 0.69]	0.67 [0.65, 0.68]	0.65 [0.61, 0.68]	0.56 [0.54, 0.59]	0.59 [0.55, 0.63]	0.58 [0.55, 0.6]	0.64 [0.62, 0.66]

(b) Evaluation of finding bounding boxes against 6 finding classes from MS-CXR and 18 classes from REFLACX (Sec. C.3). We report the pixel-level Intersection-over-Union (IoU), Intersection-over-Prediction (IoP), and Intersection-over-Target (IoT), each thresholded at 30%, and micro-averaged. Compared to Chest ImaGenome, our bounding boxes are better matching the hand-labeled boxes, especially leading to smaller and more precise boxes (larger IoP), which we assume is due to our more fine-grained region annotations.

	MS-CXR				REFLACX	
	[IoU@30]	[IoP@30]	[IoT@30]	[IoU@30]	[IoP@30]	[IoT@30]
Ours (scene graphs)	0.51 [0.47, 0.54]	0.56 [0.52, 0.6]	0.94 [0.92, 0.96]	0.45 [0.44, 0.47]	0.54 [0.53, 0.56]	0.87 [0.86, 0.88]
Chest ImaGenome	0.45 [0.42, 0.49]	0.48 [0.45, 0.52]	0.98 [0.97, 0.99]	0.42 [0.4, 0.43]	0.46 [0.44, 0.47]	0.95 [0.94, 0.96]

the specific finding tag. We calculate pixel-level Intersection-over-Union (IoU), Intersection-over-Prediction (IoP), and Intersection-over-Target (IoT) for each finding class, considering only image pairs with positive predictions and targets. Thresholding at 30% IoU/IoP/IoT, we micro-average the results, reporting the percentage of accurately localized finding-boxes in Tab. 2b. On the IoU metric, our scene graphs perform slightly better than the ones from Chest ImaGenome. The low IoP values indicate that bounding boxes are often too large, but high IoT values suggest that they generally cover the finding boxes well. This discrepancy arises because bounding boxes are derived from anatomical regions mentioned in reports, whereas hand-labeled annotations are more precise. Notably, our approach produces more precise boxes (higher IoP) than Chest ImaGenome, likely due to our large number of fine-grained region annotations (257 region classes).

Our analysis confirms that our scene graphs contain plausible finding tags and bounding boxes, with competitive or better quality than Chest ImaGenome. The bounding box quality, in turn, validates the plausibility of our region tags. Overall, our construction process yields high-quality scene graphs, making them a reliable foundation for generating QA samples.

4.2 QUALITY OF THE QA-SAMPLES

We assess the quality of our 42.2M QA pairs using an LLM-as-a-judge approach co-designed with trained radiologists (Sec. 3.3). For quality control, we conducted a human study across the five rating criteria (entailment, relevance, completeness, question clarity, and answer clarity), sampling 100 cases per criterion and rating them with the same information provided to the LLM. We found that at most 2% of cases were overrated by the LLM (e.g., assigning fine-tuning instead of pre-training grade), with high overall agreement (Cohen's kappa) and a slight tendency of the LLM to be more critical than human ratings. Details are provided in Sec. C.1.

Quality rating results for our dataset are shown in Fig. 4a. We found that 18.6% were fine-tuning grade, 58.8% were pre-training grade, and 22.6% were marked for exclusion. Notably, 85% of individual main answers were rated A or higher. We also analyzed the main causes of ratings (Fig. 4b) and found that A+ samples were limited by minor incompleteness (minor details missing), A samples

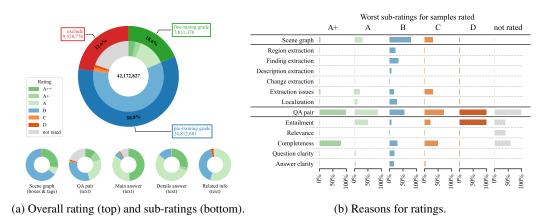


Figure 4: Results of quality assessment (Sec. 3.3). We identified a significant amount of fine-tuning grade samples, while even pre-training grade samples provide factually accurate answers, especially having high quality main answers.

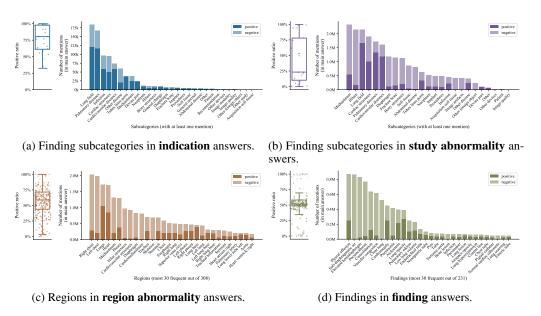


Figure 5: Distribution of tags (finding subcategories, regions, findings) mentioned in answers of different question types (indication, study abnormality, region abnormality, finding). We show their positive ratios, i.e. how often they are mentioned in positive versus in negative answers (**left**) and plot the number of positive and negative mentions of the most frequent tags (**right**). These fine-grained tags enable filtering and balancing the dataset or can be used as additional supervision.

by minor entailment aspects (facts not explicitly mentioned in the report), while B samples were restricted by issues with region/finding/localization extraction, completeness, and text clarity. Ratings C were caused by major incompleteness or extraction issues, ratings D by contradicting entailments, while non-rated samples where due to the LLM-judge not producing parsable outputs. Our analysis shows that even pre-training grade samples provide factually accurate answers with minor flaws, making them suitable for pre-training purposes.

4.3 FINDING- AND REGION-DISTRIBUTION IN QA-SAMPLES

Our answers include additional tags for findings (and their categories), regions, and answer positivity (positive or negative finding), enabling filtering and balancing for specific applications. For instance, undersampling negative answers can help mitigate model biases towards negative predictions. In

Fig. 5 we analyze the distribution of these tags. We observe that indication questions tend to have more positive mentions (Fig. 5a) – as there is a specific indication to check for – while study abnormality questions have more negative ones (Fig. 5b) – as many samples are negative overall. In region abnormality questions, most regions are mentioned slightly more often with positive than with negative findings (Fig. 5c), while for finding questions mentions are mostly balanced (Fig. 5d). This shows the success of our balanced region/finding sampling used for these two question types.

4.4 ANSWER CHARACTERISTICS

Our QA-samples provide detailed free-text answers consisting of one or even multiple sentences (i.e. answer parts). In Fig. 6, we analyze the distribution of lengths of these answers and study differences between types of answers or questions. The median answer length is 14 words, with similar lengths for most question types except for indication questions, where answers are much longer (46 words). We also observe that related information answers are much longer (22 words) than main answers (9 words) or details answers (7 words), which is expected as they can provide a lot additional context to the answers. Answers describing positive findings are typically very long (18 words), considerably longer than negative finding answers (10 words). This highlights that our dataset provides nuanced finding description in their answers, following the level of detail typically present in radiology reports.

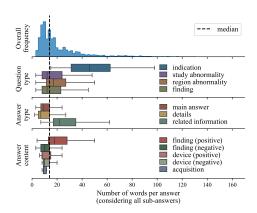


Figure 6: Distribution of answer lengths. We provide nuanced answers with detailed freetext finding descriptions.

5 STRUCTURED VQA TASK

To showcase the utility of our dataset, we introduce the *Structured Visual Question Answering (structured VQA)* task. Given a free-text question, the model is expected to generate a (multi-part) free-text answer accompanied by bounding boxes and tags (e.g., findings, regions). Unlike classical VQA, this task demands more structured answers, thus improving interpretability and clinical utility.

Baseline Model We implement a baseline model based on the Llava architecture (Liu et al., 2023), using Rad-DINO (Pérez-García et al., 2025) for image encoding and the Llama 3.2 3B (Grattafiori et al., 2024) language model. Our CXR-QBA dataset provides the necessary targets, which we format into sequences using XML-style structures and special tokens to represent tags and bounding boxes. We then fine-tune the model for one epoch on 1M QApairs (MIMIC-CXR train split). We provide some ablations in Sec. F.1. However, our primary goal is to demonstrate the utility of our dataset and establish a baseline. Detailed studies of model design choices and scaling are therefore considered out of scope and left for future work.

Table 3: Results on the structured VQA task, with 95% confidence intervals. Comparison of our baseline trained on this task and MAIRA-2 with adapted prompt. Our dataset enables training vision-language models to predict logically correct, visually grounded, and tagged answers.

			Ours	MAIRA-2
		Logical Prec.	0.76 [0.75, 0.76]	0.25
	RadFact	Logical Rec.	0.75	0.64 [0.63, 0.65]
	Rac	Grounding Prec.	0.87	0.69 [0.67, 0.71]
OA		Grounding Rec.	0.89 [0.88, 0.89]	0.12
RadStrucVQA		Finding Prec.	0.68 [0.67, 0.69]	-
adSt		Finding Rec.	0.66 [0.66, 0.67]	-
~	Tags	Finding-pos Prec.	0.41 [0.40, 0.43]	-
	-	Finding-pos Rec.	0.26 [0.25, 0.27]	-
		Region Prec.	0.67 [0.66, 0.68]	-
		Region Rec.	0.66 [0.65, 0.67]	-

RadStrucVQA Metric For evaluation, we introduce the *RadStrucVQA* metric, which closely follows the RadFact (Bannur et al., 2024) metric introduced for radiology report generation but is generalized to structured VQA. Like RadFact, we identify whether individual predicted answer parts

are entailed with target answer parts and vice-versa, in our casing using Llama 3.1 8B. For entailed pairs, we compute whether they are visually grounded, i.e. whether their bounding boxes are precise enough considering their references, and whether finding and region tags are correctly reported. This is conducted bi-directionally, using either the targets as references for the predictions or vice-versa, resulting in precision or recall scores, respectively. More details can be found in Sec. F.2.

Results We evaluate our baseline model on our fine-tuning grade test set. As the structured VQA task is newly introduced, no dedicated baselines exist. For comparison, we thus select MAIRA-2 (Bannur et al., 2024), the only publicly available CXR report generation model that supports bounding box prediction. To partially support the task, we adapt its prompt. However, tag prediction is not supported. Further details are provided in Sec. F.3.

Results are shown in Tab. 3. Our model achieves high scores in both textual content (logical) and grounding metrics. As expected, MAIRA-2 performed lower on all metrics, but achieved 85% of our model's logical recall, suggesting it captured most relevant information while also including extraneous details (lower precision). This aligns with its training objective of comprehensive reporting but also indicates that our dataset's answers do not contradict with MAIRA-2's predictions, further confirming the quality of our dataset. MAIRA-2's grounding precision significantly exceeded its recall, because it was trained to predict bounding boxes only for positive findings. Our model successfully predicts finding and region tags in most cases. However, performance drops when focusing solely on positive findings (finding-pos), indicating potential underprediction due to our training procedure or limitations in the pre-trained components. While further analysis would be required, such problems may also lead to flaws in textual answers. Importantly, our dataset's detailed tags enable fine-grained analysis of such issues while also enabling potential solutions like data filtering or balancing, making it well-suited for complex training scenarios.

6 DISCUSSION

6.1 USE CASES AND IMPACT

We introduced the *CXR-QBA* dataset, a large-scale, automatically constructed VQA dataset for CXRs offering structured, tagged, and localized answers. By design, it is particularly well-suited for the structured VQA task proposed in Sec. 5. Additionally, its large size and detailed answers make it a valuable resource for pre-training vision-language models. The accompanying tags further enable filtering and balancing of the dataset to suite specific needs. Furthermore, our fine-grained scene graphs with bounding boxes, textual descriptions, and tags can serve as a versatile data source for various purposes. For instance, they can be leveraged to create customized datasets for grounded report generation or VQA, or even as a direct training source for graph generation models to predict scene graphs on unseen chest X-rays, enabling the creation of even larger datasets. Moreover, the bounding boxes and tags provided with the scene graphs can be used for longitudinal analysis, including region-level examination. They can also be used to train models for pathology localization or classification, providing fine-grained and long-tail diagnosis targets that are often lacking in existing datasets. Finally, we release all source code for dataset construction, enabling future work to transfer our dataset construction framework to other modalities beyond chest X-rays.

6.2 LIMITATIONS

Our dataset was automatically constructed, relying on models and templates instead of human annotations. While this enables the generation of a large number of QA-pairs, it may also introduce potential errors and biases. We apply automatic quality assessments and extensive quality controls to mitigate these risks, but users should still be aware that inaccuracies may remain and exercise caution in critical applications. We strongly advise against using this dataset as the sole source for fine-tuning or evaluating models in clinical practice. For such purposes, we recommend using small-scale, gold-standard (fully human-annotated) datasets, which are, however, beyond the scope of this work. Furthermore, our template-based approach may also limit diversity and introduce grammatical errors, though we partially mitigate these issues by incorporating answers from actual report sentences and through quality assessment. Additionally, the dataset focuses on individual chest X-ray studies, excluding longitudinal and differential questions as well as other imaging modalities. Future work could extend our approach to broader question types and modalities.

ETHICS STATEMENT

This work does not directly involve human subjects and does not collect new human subject data. All human subject data used (such as chest X-rays or radiology reports) are derived from existing public datasets, which were collected independently of and prior to this work. We will release the dataset as a credentialed dataset on the Physionet platform (Goldberger et al., 2000). This ensures compliance with the license of MIMIC-CXR (Johnson et al., 2024b), from which our dataset is derived. It also serves to protect the included data, which is indirectly derived from patient health information. While enabling researcher access, it restricts use for other purposes and requires completion of a privacy and ethics course. Furthermore, to protect privacy and comply with the MIMIC-CXR license, all data processing was performed locally, and no data were sent to external services.

REPRODUCIBILITY STATEMENT

We provide full details to ensure reproducibility of our work. The construction of the dataset is described in Sec. 3, with additional information in Secs. D and E. The structured VQA task and associated metrics are presented in Sec. 5, with further details in Sec. F. All source code for dataset construction and the structured VQA task is available at https://anonymous.4open.science/r/mimic-ext-cxr-qba/ and will be made publicly accessible. The constructed dataset will be released through the PhysioNet platform (Goldberger et al., 2000), where it will be accessible to researchers after credentialing by PhysioNet (which requires a short online training). This hosting is necessary because the dataset is derived from MIMIC-CXR.

REFERENCES

Seongsu Bae, Daeun Kyung, Jaehee Ryu, Eunbyeol Cho, Gyubok Lee, Sunjun Kweon, Jungwoo Oh, Lei JI, Eric Chang, Tackeun Kim, and Edward Choi. MIMIC-Ext-MIMIC-CXR-VQA: A Complex, Diverse, And Large-Scale Visual Question Answering Dataset for Chest X-ray Images, 2024. URL https://physionet.org/content/mimic-ext-mimic-cxr-vqa/1.0.0/.

Shruthi Bannur, Kenza Bouzid, Daniel C. Castro, Anton Schwaighofer, Anja Thieme, Sam Bond-Taylor, Maximilian Ilse, Fernando Pérez-García, Valentina Salvatelli, Harshita Sharma, Felix Meissen, Mercy Ranjit, Shaury Srivastav, Julia Gong, Noel C. F. Codella, Fabian Falck, Ozan Oktay, Matthew P. Lungren, Maria Teodora Wetscherek, Javier Alvarez-Valle, and Stephanie L. Hyland. MAIRA-2: Grounded Radiology Report Generation, September 2024. URL http://arxiv.org/abs/2406.04449.arXiv:2406.04449 [cs].

Asma Ben Abacha, Sadid A. Hasan, Vivek V. Datla, Dina Demner-Fushman, and Henning Müller. VQA-Med: overview of the medical visual question answering task at ImageCLEF 2019. *Proceedings of CLEF (Conference and Labs of the Evaluation Forum) 2019 Working Notes*, September 2019. Meeting Name: CLEF (Conference and Labs of the Evaluation Forum) 2019 Working Num Pages: 11 Place: Lugano, Switzerland Publisher: 9-12 September 2019.

Ricardo Bigolin Lanfredi, Mingyuan Zhang, William Auffermann, Jessica Chan, Phuong-Anh Duong, Vivek Srikumar, Trafton Drew, Joyce Schroeder, and Tolga Tasdizen. REFLACX: Reports and eye-tracking data for localization of abnormalities in chest x-rays, 2021. URL https://physionet.org/content/reflacx-xray-localization/1.0.0/.

Ricardo Bigolin Lanfredi, Mingyuan Zhang, William F. Auffermann, Jessica Chan, Phuong-Anh T. Duong, Vivek Srikumar, Trafton Drew, Joyce D. Schroeder, and Tolga Tasdizen. REFLACX, a dataset of reports and eye-tracking data for localization of abnormalities in chest x-rays. *Sci Data*, 9(1):350, June 2022. ISSN 2052-4463. doi: 10.1038/s41597-022-01441-z. URL https://www.nature.com/articles/s41597-022-01441-z.

Benedikt Boecking, Naoto Usuyama, Shruthi Bannur, Daniel C. Castro, Anton Schwaighofer, Stephanie Hyland, Maria Wetscherek, Tristan Naumann, Aditya Nori, Javier Alvarez-Valle, Hoifung Poon, and Ozan Oktay. Making the Most of Text Semantics to Improve Biomedical Vision–Language Processing. In Shai Avidan, Gabriel Brostow, Moustapha Cissé, Giovanni Maria Farinella, and Tal Hassner (eds.), *Computer Vision – ECCV 2022*, volume 13696, pp. 1–21. Springer Nature Switzerland, Cham, 2022. doi: 10.1007/978-3-031-20059-5_1.

```
URL https://link.springer.com/10.1007/978-3-031-20059-5_1. Series Title: Lecture Notes in Computer Science.
```

- Benedikt Boecking, Naoto Usuyama, Shruthi Bannur, Daniel Coelho de Castro, Anton Schwaighofer, Stephanie Hyland, Harshita Sharma, Maria Teodora Wetscherek, Tristan Naumann, Aditya Nori, Javier Alvarez Valle, Hoifung Poon, and Ozan Oktay. MS-CXR: Making the Most of Text Semantics to Improve Biomedical Vision-Language Processing, 2024. URL https://physionet.org/content/ms-cxr/1.1.0/.
- Aurelia Bustos, Antonio Pertusa, Jose-Maria Salinas, and Maria de la Iglesia-Vayá. PadChest: A large chest x-ray image dataset with multi-label annotated reports. *Medical Image Analysis*, 66: 101797, December 2020. ISSN 1361-8415. doi: 10.1016/j.media.2020.101797. URL https://www.sciencedirect.com/science/article/pii/S1361841520301614.
- Zhihong Chen, Maya Varma, Justin Xu, Magdalini Paschali, Dave Van Veen, Andrew Johnston, Alaa Youssef, Louis Blankemeier, Christian Bluethgen, Stephan Altmayer, Jeya Maria Jose Valanarasu, Mohamed Siddig Eltayeb Muneer, Eduardo Pontes Reis, Joseph Paul Cohen, Cameron Olsen, Tanishq Mathew Abraham, Emily B. Tsai, Christopher F. Beaulieu, Jenia Jitsev, Sergios Gatidis, Jean-Benoit Delbrouck, Akshay S. Chaudhari, and Curtis P. Langlotz. A Vision-Language Foundation Model to Enhance Efficiency of Chest X-ray Interpretation, December 2024. URL http://arxiv.org/abs/2401.12208.arXiv:2401.12208 [cs].
- Ary L. Goldberger, Luis A. N. Amaral, Leon Glass, Jeffrey M. Hausdorff, Plamen Ch. Ivanov, Roger G. Mark, Joseph E. Mietus, George B. Moody, Chung-Kang Peng, and H. Eugene Stanley. PhysioBank, PhysioToolkit, and PhysioNet. *Circulation*, 101(23):e215–e220, June 2000. doi: 10.1161/01.CIR.101.23.e215. URL https://www.ahajournals.org/doi/full/10.1161/01.cir.101.23.e215. Publisher: American Heart Association.
- Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, et al. The Llama 3 Herd of Models, November 2024. URL http://arxiv.org/abs/2407.21783.arXiv:2407.21783 [cs].
- Gregory Holste, Yiliang Zhou, Song Wang, Ajay Jaiswal, Mingquan Lin, Sherry Zhuge, Yuzhe Yang, Dongkyun Kim, Trong-Hieu Nguyen-Mau, Minh-Triet Tran, Jaehyup Jeong, Wongi Park, Jongbin Ryu, Feng Hong, Arsh Verma, Yosuke Yamagishi, Changhyun Kim, Hyeryeong Seo, Myungjoo Kang, Leo Anthony Celi, Zhiyong Lu, Ronald M. Summers, George Shih, Zhangyang Wang, and Yifan Peng. Towards long-tailed, multi-label disease classification from chest X-ray: Overview of the CXR-LT challenge. *Medical Image Analysis*, 97:103224, October 2024. ISSN 1361-8415. doi: 10.1016/j.media.2024.103224. URL https://www.sciencedirect.com/science/article/pii/S136184152400149X.
- Gregory Holste, Mingquan Lin, Song Wang, Yiliang Zhou, Yishu Wei, Hao Chen, Atlas Wang, and Yifan Peng. CXR-LT: Multi-Label Long-Tailed Classification on Chest X-Rays, 2025. URL https://physionet.org/content/cxr-lt-iccv-workshop-cvamd/2.0.0/.
- Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. LoRA: Low-Rank Adaptation of Large Language Models. October 2021. URL https://openreview.net/forum?id=nZeVKeeFYf9.
- Xinyue Hu, Lin Gu, Kazuma Kobayashi, Liangchen Liu, Mengliang Zhang, Tatsuya Harada, Ronald M. Summers, and Yingying Zhu. Interpretable medical image Visual Question Answering via multi-modal relationship graph learning. *Medical Image Analysis*, 97:103279, October 2024. ISSN 1361-8415. doi: 10.1016/j.media.2024.103279. URL https://www.sciencedirect.com/science/article/pii/S1361841524002044.
- Xinyue Hu, Lin Gu, Kazuma Kobayashi, liangchen liu, Mengliang Zhang, Tatsuya Harada, Ronald Summers, and Yingying Zhu. Medical-CXR-VQA dataset: A Large-Scale LLM-Enhanced Medical Dataset for Visual Question Answering on Chest X-Ray Images, 2025. URL https://physionet.org/content/medical-cxr-vqa-dataset/1.0.0/.
- Jeremy Irvin, Pranav Rajpurkar, Michael Ko, Yifan Yu, Silviana Ciurea-Ilcus, Chris Chute, Henrik Marklund, Behzad Haghgoo, Robyn Ball, Katie Shpanskaya, Jayne Seekins, David A. Mong,

Safwan S. Halabi, Jesse K. Sandberg, Ricky Jones, David B. Larson, Curtis P. Langlotz, Bhavik N. Patel, Matthew P. Lungren, and Andrew Y. Ng. CheXpert: A Large Chest Radiograph Dataset with Uncertainty Labels and Expert Comparison. *Proceedings of the AAAI Conference on Artificial Intelligence*, 33(01):590–597, July 2019. ISSN 2374-3468. doi: 10.1609/aaai.v33i01.3301590. URL https://ojs.aaai.org/index.php/AAAI/article/view/3834. Number: 01.

- Saahil Jain, Ashwin Agrawal, Adriel Saporta, Steven Truong, Du Nguyen Duong N. Duong, Tan Bui, Pierre Chambon, Yuhao Zhang, Matthew Lungren, Andrew Ng, Curtis Langlotz, and Pranav Rajpurkar. RadGraph: Extracting Clinical Entities and Relations from Radiology Reports. Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks, 1, December 2021. URL https://datasets-benchmarks-proceedings.neurips.cc/paper/2021/hash/c8ffe9a587b126f152ed3d89a146b445-Abstract-round1.html.
- Alistair Johnson, Matthew Lungren, Yifan Peng, Zhiyong Lu, Roger Mark, Seth Berkowitz, and Steven Horng. MIMIC-CXR-JPG chest radiographs with structured labels, 2024a. URL https://physionet.org/content/mimic-cxr-jpg/2.1.0/.
- Alistair Johnson, Tom Pollard, Roger Mark, Seth Berkowitz, and Steven Horng. MIMIC-CXR Database, 2024b. URL https://physionet.org/content/mimic-cxr/2.1.0/.
- Alistair E. W. Johnson, Tom J. Pollard, Seth J. Berkowitz, Nathaniel R. Greenbaum, Matthew P. Lungren, Chih-ying Deng, Roger G. Mark, and Steven Horng. MIMIC-CXR, a de-identified publicly available database of chest radiographs with free-text reports. *Sci Data*, 6(1):317, December 2019a. ISSN 2052-4463. doi: 10.1038/s41597-019-0322-0. URL https://www.nature.com/articles/s41597-019-0322-0.
- Alistair E. W. Johnson, Tom J. Pollard, Nathaniel R. Greenbaum, Matthew P. Lungren, Chih-ying Deng, Yifan Peng, Zhiyong Lu, Roger G. Mark, Seth J. Berkowitz, and Steven Horng. MIMIC-CXR-JPG, a large publicly available database of labeled chest radiographs, November 2019b. URL http://arxiv.org/abs/1901.07042.arXiv:1901.07042 [cs].
- Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph Gonzalez, Hao Zhang, and Ion Stoica. Efficient Memory Management for Large Language Model Serving with PagedAttention. In *Proceedings of the 29th Symposium on Operating Systems Principles*, SOSP '23, pp. 611–626, New York, NY, USA, October 2023. Association for Computing Machinery. ISBN 979-8-4007-0229-7. doi: 10.1145/3600006.3613165. URL https://dl.acm.org/doi/10.1145/3600006.3613165.
- Jason J. Lau, Soumya Gayen, Asma Ben Abacha, and Dina Demner-Fushman. A dataset of clinically generated visual questions and answers about radiology images. *Sci Data*, 5(1):180251, November 2018. ISSN 2052-4463. doi: 10.1038/sdata.2018.251. URL https://www.nature.com/articles/sdata2018251.
- Ji Lin, Jiaming Tang, Haotian Tang, Shang Yang, Wei-Ming Chen, Wei-Chen Wang, Guangxuan Xiao, Xingyu Dang, Chuang Gan, and Song Han. AWQ: Activation-aware Weight Quantization for On-Device LLM Compression and Acceleration. *Proceedings of Machine Learning and Systems*, 6: 87–100, May 2024. URL https://proceedings.mlsys.org/paper_files/paper/2024/hash/42a452cbafa9dd64e9ba4aa95cclef21-Abstract-Conference.html.
- Bo Liu, Li-Ming Zhan, Li Xu, Lin Ma, Yan Yang, and Xiao-Ming Wu. Slake: A Semantically-Labeled Knowledge-Enhanced Dataset For Medical Visual Question Answering. In 2021 IEEE 18th International Symposium on Biomedical Imaging (ISBI), pp. 1650–1654, April 2021. doi: 10.1109/ISBI48211.2021.9434010. URL https://ieeexplore.ieee.org/abstract/document/9434010. ISSN: 1945-8452.
- Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual Instruction Tuning. Advances in Neural Information Processing Systems, 36:34892-34916, December 2023. URL https://papers.nips.cc/paper_files/paper/2023/hash/6dcf277ea32ce3288914faf369fe6de0-Abstract-Conference.html.

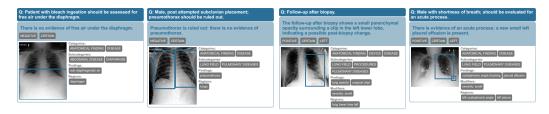
Philip Müller, Georgios Kaissis, and Daniel Rueckert. ChEX: Interactive Localization and Region Description in Chest X-Rays. In Aleš Leonardis, Elisa Ricci, Stefan Roth, Olga Russakovsky, Torsten Sattler, and Gül Varol (eds.), *Computer Vision – ECCV 2024*, pp. 92–111, Cham, 2025. Springer Nature Switzerland. ISBN 978-3-031-72664-4. doi: 10.1007/978-3-031-72664-4_6.

Fernando Pérez-García, Harshita Sharma, Sam Bond-Taylor, Kenza Bouzid, Valentina Salvatelli, Maximilian Ilse, Shruthi Bannur, Daniel C. Castro, Anton Schwaighofer, Matthew P. Lungren, Maria Teodora Wetscherek, Noel Codella, Stephanie L. Hyland, Javier Alvarez-Valle, and Ozan Oktay. Exploring scalable medical image encoders beyond text supervision. *Nat Mach Intell*, 7(1):119–130, January 2025. ISSN 2522-5839. doi: 10.1038/s42256-024-00965-w. URL https://www.nature.com/articles/s42256-024-00965-w. Publisher: Nature Publishing Group.

- Peng Qi, Yuhao Zhang, Yuhui Zhang, Jason Bolton, and Christopher D. Manning. Stanza: A Python Natural Language Processing Toolkit for Many Human Languages. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, 2020. URL https://nlp.stanford.edu/pubs/qi2020stanza.pdf.
- François Remy, Kris Demuynck, and Thomas Demeester. BioLORD-2023: semantic textual representations fusing large language models and clinical knowledge graph insights. *Journal of the American Medical Informatics Association*, 31(9):1844–1855, September 2024. ISSN 1527-974X. doi: 10.1093/jamia/ocae029. URL https://doi.org/10.1093/jamia/ocae029.
- Constantin Seibold, Simon Reiß, Saquib Sarfraz, Matthias A. Fink, Victoria Mayer, Jan Sellner, Moon Sung Kim, Klaus H. Maier-Hein, Jens Kleesiek, and Rainer Stiefelhagen. Detailed Annotations of Chest X-Rays via CT Projection for Report Understanding. arXiv, October 2022. doi: 10.48550/arXiv.2210.03416. URL https://bmvc2022.mpi-inf.mpg.de/0058.pdf. arXiv:2210.03416 [cs].
- Constantin Seibold, Alexander Jaus, Matthias A. Fink, Moon Kim, Simon Reiß, Ken Herrmann, Jens Kleesiek, and Rainer Stiefelhagen. Accurate Fine-Grained Segmentation of Human Anatomy in Radiographs via Volumetric Pseudo-Labeling, June 2023. URL http://arxiv.org/abs/2306.03934. arXiv:2306.03934 [eess].
- SNOMED International. SNOMED CT, 2023. URL https://www.snomed.org.
- Tao Tu, Shekoofeh Azizi, Danny Driess, Mike Schaekermann, Mohamed Amin, Pi-Chuan Chang, Andrew Carroll, Charles Lau, Ryutaro Tanno, Ira Ktena, Anil Palepu, Basil Mustafa, Aakanksha Chowdhery, Yun Liu, Simon Kornblith, David Fleet, Philip Mansfield, Sushant Prakash, Renee Wong, Sunny Virmani, Christopher Semturs, S. Sara Mahdavi, Bradley Green, Ewa Dominowska, Blaise Aguera y Arcas, Joelle Barral, Dale Webster, Greg S. Corrado, Yossi Matias, Karan Singhal, Pete Florence, Alan Karthikesalingam, and Vivek Natarajan. Towards Generalist Biomedical AI. NEJM AI, 1(3):Aloa2300138, February 2024. doi: 10.1056/Aloa2300138. URL https://ai.nejm.org/doi/abs/10.1056/Aloa2300138. Publisher: Massachusetts Medical Society.
- Joy Wu, Nkechinyere Agu, Ismini Lourentzou, Arjun Sharma, Joseph Paguio, Jasper Seth Yao, Edward Christopher Dee, William Mitchell, Satyananda Kashyap, Andrea Giovannini, Leo Anthony Celi, Tanveer Syeda-Mahmood, and Mehdi Moradi. Chest ImaGenome Dataset, 2021a. URL https://physionet.org/content/chest-imagenome/1.0.0/.
- 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, Leo A. Celi, and Mehdi Moradi. Chest ImaGenome Dataset for Clinical Reasoning, July 2021b. URL http://arxiv.org/abs/2108.00316. arXiv:2108.00316 [cs].
- Yunfei Xie, Ce Zhou, Lang Gao, Juncheng Wu, Xianhang Li, Hong-Yu Zhou, Sheng Liu, Lei Xing, James Zou, Cihang Xie, and Yuyin Zhou. MedTrinity-25M: A Large-scale Multimodal Dataset with Multigranular Annotations for Medicine, March 2025. URL http://arxiv.org/abs/2408.02900. arXiv:2408.02900 [cs].

Xiaoman Zhang, Chaoyi Wu, Ziheng Zhao, Weixiong Lin, Ya Zhang, Yanfeng Wang, and Weidi Xie. Development of a large-scale medical visual question-answering dataset. *Commun Med*, 4 (1):1–13, December 2024. ISSN 2730-664X. doi: 10.1038/s43856-024-00709-2. URL https://www.nature.com/articles/s43856-024-00709-2. Publisher: Nature Publishing Group.

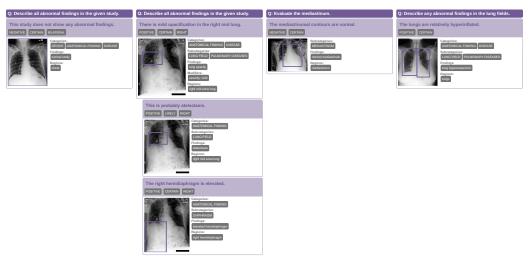
A EXAMPLE QA-PAIRS FROM OUR DATASET



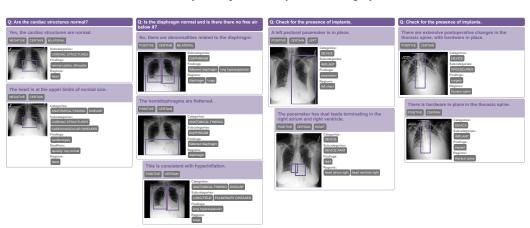
(a) Negative.

(b) Positive.

Figure 7: Examples of **indication** questions. Questions are based on the paraphrased INDICATION section while each main answer is generated based on the indication node from the scene graph.



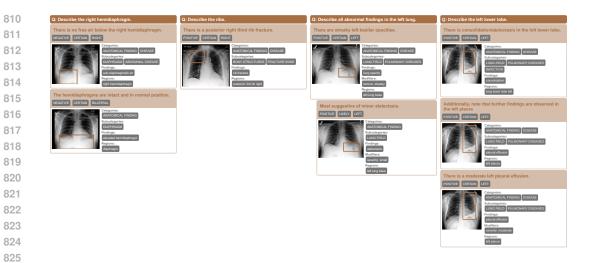
(a) Abnormality descriptions (study-level or category-level).



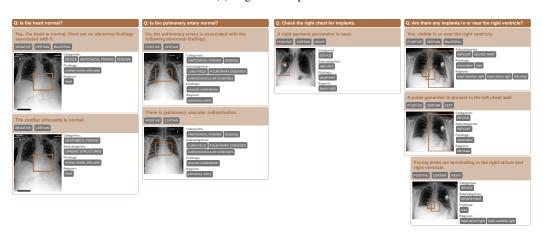
(b) Abnormality assessment (category-level).

(c) Device descriptions.

Figure 8: Examples of **study abnormality** questions. Questions are based on one of 13 templates. Answers may consist of several answer parts, where each describes an individual aspect (about the overall study or a finding category). Individual answer parts are constructed based on observation nodes, filtered based on finding categories relevant to the question, where individual answer parts may be organized hierarchically (indicated by indentations) based on parent-child edges in the scene graph. Additionally assessment answers (b) start with a template-based yes/no answer.



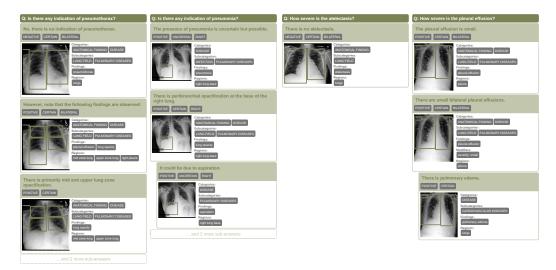
(a) Region description.



(b) Region assessment.

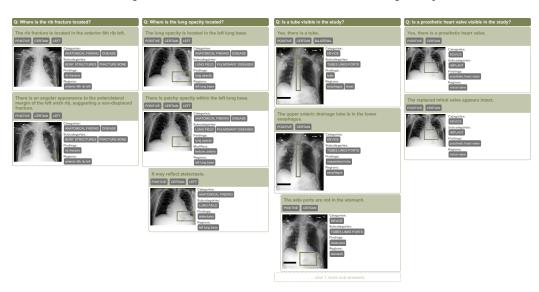
(c) Region devices.

Figure 9: Examples of **region abnormality** questions. Questions are based on one of 6 templates. Answers may consist of several answer parts, where each describes an individual aspect (about the region). Individual answer parts are constructed based on observation nodes relevant to the region, where individual answer parts may be organized hierarchically (indicated by indentations) based on parent-child edges in the scene graph. Additionally assessment answers (b) start with a template-based yes/no answer. Some templates also ask specifically about devices in the region (c).



(a) Finding assessment.

(b) Finding description.

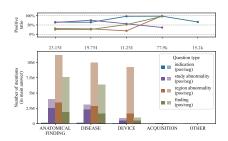


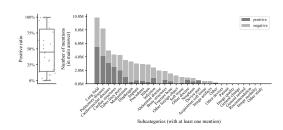
(c) Finding location.

(d) Device.

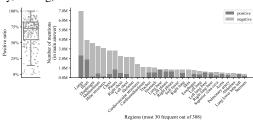
Figure 10: Examples of **finding** questions. Questions are based on one of 7 templates. Answers start with a template-based answer part to identify the finding presence (a), provide a severity summary (b), describe the location (c), or presence of a device (d). Additional details may be provided in answer parts based on observation nodes relevant to the finding, where individual answer parts may be organized hierarchically (indicated by indentations) based on parent-child edges in the scene graph.

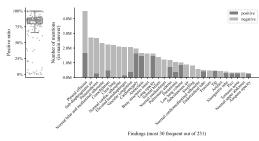
B FINDING- AND REGION-DISTRIBUTION IN QA-SAMPLES





- (a) Finding (main-)categories for different question types (indication, study abnormality, region abnormality, finding).
- (b) Finding sub-categories (all samples).





(c) Regions (all samples).

(d) Findings (all samples).

Figure 11: Tags (finding main- and sub-categories, regions, findings) mentioned in answers. We show their positive ratios (**top/left**), i.e. how often they are mentioned in positive versus in negative answer parts and plot the number of positive and negative mentions of the most frequent tags (**bottom/right**).

C EVALUATION DETAILS

C.1 QA EVALUATION: ASSESSMENT OF THE LLM-AS-A-JUDGE

Table 4: Systematic study to reliably quantify the effectiveness of the LLM-as-a-judge. For each of the five quality criteria (entailment, relevance, completeness, question clarity, answer clarity), we independently sampled 100 cases (without replacement). Each of these cases was rated by a human annotator provided with the same information available to the LLM-as-a-judge. We show two metrics: (i) the percentage of cases where LLM judgment lead to an overrated grade (e.g. fine-tuning instead of pre-training grade) compared to human judgment and (ii) Cohen's kappa, a standard metric for inter-rater reliability. In a maximum of 2% of the cases the LLM judgment lead to an overrated grade while we found overall high agreement with a tendency of the LLM to be more critical than human ratings (especially for completeness and clarity).

Metric	Entailment	Relevance	Completeness	Question clarity	Answer clarity
% overrated grade Cohen's kappa	(7 classes) 1% 0.49	(4 classes) 0% 0.65	(5 classes) 2% 0.41	(6 classes) 0% 0.32	(5 classes) 0% 0.32

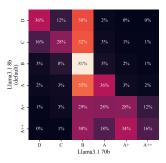


Figure 12: Confusion matrix comparing the assigned quality ratings between using Llama3.1 8b (default) and Llama3.1 70b as an LLM-judge (see Secs. 3.3 and 4.2). In most cases, ratings differ only slightly. Most importantly, low-quality samples (as rated by Llama3.1 70b) are almost never assigned to fine-tuning grades (A or higher) by Llama3.1 8b. We thus decided to use Llama3.1 8b as our default rater, as it is much more computationally efficient.

C.2 Scene Graph Evaluation: Finding Tags

Table 5: Evaluation of finding tags against the 13 CheXpert (CXP) classes from the MIMIC-CXR-JPG test set (Sec. 4.1). We show finding-level scores, macro-averages over subsets and the micro-average, with 95% confidence intervals (bootstrapping, n = 1000).

			MIMIC-0	CXR-JPG Joh	nson et al.	(2024a) Test		
	[Pre	ecision]	[R	[Recall]		[F1]		ICC]
	Ours	Chest ImaG.	Ours	Chest ImaG.	Ours	Chest ImaG.	Ours	Chest ImaG
Findings in CXP-5, C	XP-7, and	CXP-13						
Atelectasis	0.82 [0.78, 0.87]	0.78 [0.73, 0.83]	0.99 [0.97, 1.0]	0.99 [0.98, 1.0]	0.9 [0.87, 0.93]	0.88 [0.84, 0.9]	0.84 [0.8, 0.88]	0.81 [0.77, 0.85]
Cardiomegaly	0.64 [0.58, 0.7]	0.67 [0.61, 0.74]	0.85 [0.78, 0.9]	0.82 [0.75, 0.87]	0.73 [0.67, 0.78]	0.74	0.61 [0.54, 0.68]	0.63 [0.56, 0.69]
Consolidation	0.83 [0.73, 0.91]	0.77	0.87 [0.79, 0.94]	0.93	0.85 [0.78, 0.91]	0.84	0.83 [0.75, 0.9]	0.83 [0.76, 0.89]
Edema	0.94 [0.9, 0.98]	0.9 [0.85, 0.95]	0.8 [0.73, 0.86]	0.8	0.86 [0.82, 0.9]	0.85 [0.8, 0.89]	0.83 [0.77, 0.87]	0.8
Pleural Effusion	0.9 [0.86, 0.93]	0.86 [0.82, 0.9]	0.98 [0.96, 1.0]	0.97 [0.94, 0.99]	0.94 [0.92, 0.96]	0.91 [0.89, 0.94]	0.89 [0.85, 0.92]	0.85 [0.8, 0.89]
Findings in CXP-7 and	d CXP-13							
Pneumonia	0.92 [0.87, 0.96]	0.89 [0.84, 0.94]	0.94 [0.89, 0.97]	0.95 [0.91, 0.98]	0.93 [0.89, 0.96]	0.92 [0.89, 0.95]	0.91 [0.87, 0.95]	0.9 [0.86, 0.94]
Pneumothorax	0.78 [0.64, 0.89]	0.79	0.84 [0.71, 0.95]	0.89	0.8 [0.69, 0.89]	0.84	0.79 [0.68, 0.88]	0.83 [0.72, 0.91]
Findings in CXP-13								
Enlarged Cardiom.	0.51 [0.37, 0.65]	0.61 [0.41, 0.8]	0.39 [0.29, 0.51]	0.23 [0.13, 0.33]	0.44 [0.34, 0.55]	0.33 [0.2, 0.45]	0.39 [0.28, 0.51]	0.33 [0.19, 0.45]
Lung Lesion	0.17 [0.12, 0.22]	0.68	0.81 [0.69, 0.91]	0.87	0.28 [0.21, 0.35]	0.76	0.25 [0.17, 0.32]	0.74 [0.64, 0.83]
Lung Opacity	0.62 [0.56, 0.69]	0.28	0.83	1.0 [1.0, 1.0]	0.71 [0.65, 0.76]	0.43	0.61 [0.54, 0.68]	0.2
Pleural Other	0.54 [0.36, 0.71]	0.3	0.87 [0.7, 1.0]	0.92 [0.76, 1.0]	0.67 [0.49, 0.8]	0.45 [0.31, 0.58]	0.67 [0.51, 0.79]	0.5 [0.37, 0.6]
Fracture	0.67 [0.54, 0.79]	0.6	0.92 [0.82, 1.0]	0.82 [0.68, 0.93]	0.77 [0.67, 0.86]	0.69	0.77 [0.67, 0.85]	0.68 [0.56, 0.78]
Support Devices	0.61 [0.56, 0.66]	0.62 [0.56, 0.67]	0.98 [0.96, 1.0]	0.83 [0.77, 0.88]	0.75 [0.71, 0.79]	0.71 [0.66, 0.75]	0.63 [0.59, 0.68]	0.54 [0.48, 0.61]
Macro-averages								
CheXpert-5 (CXP-5)	0.83 [0.79, 0.85]	0.8 [0.77, 0.83]	0.9 [0.87, 0.92]	0.9 [0.88, 0.92]	0.85 [0.83, 0.87]	0.84 [0.82, 0.86]	0.8 [0.77, 0.82]	0.78 [0.75, 0.81]
CheXpert-7 (CXP-7)	0.83 [0.8, 0.86]	0.81	0.89 [0.87, 0.92]	0.91	0.86	0.85	0.81 [0.79, 0.84]	0.8
CheXpert-13 (CXP-13)		0.67 [0.65, 0.7]	0.85 [0.83, 0.87]	0.85 [0.82, 0.87]	0.74 [0.72, 0.76]	0.72 [0.7, 0.74]	0.69 [0.67, 0.71]	0.66 [0.64, 0.69]
Micro	0.68 [0.66, 0.7]	0.63	0.89 [0.87, 0.9]	0.88	0.77 [0.75, 0.79]	0.73 [0.72, 0.75]	0.71 [0.69, 0.73]	0.67 [0.65, 0.68]

Table 6: Evaluation of finding tags against the 13 CXP and 12 long-tail (LT) classes from the CXR-LT 2024 gold standard dataset (Sec. 4.1). We show finding-level scores, macro-averages over different subsets and the micro-average, with 95% confidence intervals (bootstrapping, n=1000).

			CXR-	LT 2024 Hol	ste et al. (20	025) Gold		
	[Pre	cision]	[Re	ecall]	[F1]	[N	ICC]
	Ours	Chest ImaG.	Ours	Chest ImaG.	Ours	Chest ImaG.	Ours	Chest ImaG
Findings in CXP-5, C	XP-7, CXP	-13, and CX	R-LT					
Atelectasis	0.55 [0.47, 0.62]	0.56 [0.49, 0.61]	0.82 [0.75, 0.88]	0.99 [0.97, 1.0]	0.66 [0.59, 0.71]	0.71 [0.65, 0.76]	0.48 [0.4, 0.57]	0.59 [0.54, 0.65]
Cardiomegaly	0.82 [0.76, 0.88]	0.85	0.85 [0.79, 0.91]	0.8	0.84	0.82	0.73 [0.66, 0.79]	0.72 [0.65, 0.79]
Consolidation	0.82 [0.73, 0.91]	0.74 [0.63, 0.83]	0.86 [0.76, 0.93]	0.89	0.84 [0.77, 0.9]	0.8 [0.73, 0.87]	0.8 [0.72, 0.87]	0.76 [0.67, 0.83]
Edema	0.73 [0.63, 0.83]	0.69 [0.6, 0.79]	0.64 [0.55, 0.74]	0.71 [0.61, 0.8]	0.68 [0.6, 0.76]	0.7 [0.62, 0.77]	0.59 [0.5, 0.69]	0.6 [0.5, 0.69]
Pleural Effusion	0.82	0.78 [0.73, 0.84]	0.93 [0.9, 0.97]	0.97 [0.94, 0.99]	0.87 [0.83, 0.91]	0.87 [0.83, 0.9]	0.77 [0.7, 0.82]	0.76 [0.7, 0.82]
Findings in CXP-7, C			[0.5, 0.57]	[0.54, 0.55]	[0.05, 0.71]	[0.05, 0.7]	[0.7, 0.02]	[0.7, 0.02]
Pneumonia	0.38	0.13	0.45	0.76	0.41	0.23	0.38	0.24
Pneumothorax	[0.19, 0.58] 0.85	[0.07, 0.21] 0.8	[0.25, 0.67] 0.85	[0.55, 0.94] 0.96	[0.22, 0.58] 0.85	0.87	[0.19, 0.55] 0.83	[0.13, 0.34] 0.86
	[0.73, 0.95]	[0.7, 0.9]	[0.72, 0.94]	[0.89, 1.0]	[0.75, 0.92]	[0.8, 0.94]	[0.72, 0.91]	[0.78, 0.93]
Findings in CXP-13 an Enlarged Cardiom.	nd CXR-L7 0.78	Ր 1.0	0.13	0.1	0.22	0.18	0.24	0.27
	[0.57, 0.95]	[1.0, 1.0] 0.05	[0.07, 0.2]	[0.05, 0.17]	[0.13, 0.32]	[0.1, 0.29]	[0.13, 0.35]	[0.19, 0.36]
Lung Lesion	0.01	[0.0, 0.12]	0.5 [0.0, 1.0]	0.75 [0.0, 1.0]	0.02	0.1 [0.0, 0.22]	0.0 [-0.09, 0.1]	0.18 [-0.02, 0.3]
Lung Opacity	0.92 [0.87, 0.96]		0.77 [0.72, 0.84]	1.0 [1.0, 1.0]	0.84 [0.8, 0.88]	0.7 [0.66, 0.74]	0.73 [0.66, 0.79]	0.35 [0.29, 0.4]
Pleural Other	0.11 [0.0, 0.24]	0.23 [0.13, 0.34]	0.19 [0.0, 0.43]	1.0 [1.0, 1.0]	0.14 [0.0, 0.29]	0.38 [0.23, 0.51]	0.1 [-0.05, 0.26]	0.45 [0.34, 0.55]
Fracture	0.89 [0.78, 0.98]	0.84 [0.72, 0.94]	0.91 [0.81, 0.98]	0.82 [0.69, 0.93]	0.9 [0.82, 0.95]	0.83 [0.73, 0.9]	0.89 [0.8, 0.95]	0.81 [0.69, 0.89]
Support Devices	0.92 [0.88, 0.96]	0.93 [0.9, 0.97]	0.93 [0.89, 0.96]	0.83 [0.77, 0.88]	0.92 [0.9, 0.95]	0.88 [0.84, 0.91]	0.83 [0.77, 0.88]	0.75 [0.68, 0.82]
Findings in LT-only, a	nd CXR-L	Γ						
Calcification of Aorta	0.95 [0.83, 1.0]	0.95 [0.88, 1.0]	0.43 [0.28, 0.58]	0.93 [0.85, 1.0]	0.6 [0.43, 0.73]	0.94 [0.89, 0.99]	0.61 [0.48, 0.73]	0.94 [0.87, 0.99]
Emphysema	0.58 [0.41, 0.74]	0.54	0.81 [0.63, 0.95]	0.81	0.68 [0.52, 0.8]	0.65 [0.49, 0.77]	0.66 [0.5, 0.79]	0.63 [0.47, 0.75]
Fibrosis	0.27 [0.15, 0.43]	0.0 [0.0, 0.0]	0.52 [0.31, 0.74]	0.0 [0.0, 0.0]	0.36 [0.21, 0.52]	0.0 [0.0, 0.0]	0.33 [0.17, 0.5]	0.0 [0.0, 0.0]
Hernia	1.0	0.86	0.9	0.9	0.95	0.88	0.95	0.87
Infiltration	0.15	[0.68, 1.0] 0.38	0.33	[0.73, 1.0] 0.55	0.21	[0.74, 0.97] 0.44	[0.85, 1.0] 0.19	[0.73, 0.97] 0.44
Mass	[0.03, 0.3] 0.54	0.16, 0.67]	[0.08, 0.7]	[0.21, 0.88] 0.89	[0.05, 0.39] 0.64	0.42	[0.01, 0.39] 0.63	[0.17, 0.69] 0.46
Nodule	[0.34, 0.74] 0.92	[0.17, 0.4] 0.5	[0.56, 0.94] 0.74	0.91	[0.44, 0.78] 0.82	0.64	[0.45, 0.77] 0.81	0.64
Pleural Thickening	[0.79, 1.0] 0.78	[0.36, 0.63] 0.33	[0.57, 0.88] 1.0	[0.78, 1.0] 1.0	[0.68, 0.91] 0.88	[0.51, 0.75] 0.5	[0.68, 0.91] 0.88	[0.51, 0.74] 0.54
Pneumomediastinum	[0.62, 0.92] 0.94	[0.22, 0.45] 0.88	[1.0, 1.0] 0.84	[1.0, 1.0] 0.84	[0.77, 0.96] 0.89	[0.36, 0.62] 0.86	[0.78, 0.96] 0.88	[0.43, 0.64] 0.85
Pneumoperitoneum	[0.83, 1.0] 0.88	[0.73, 0.97]	[0.7, 0.96]	[0.7, 0.96] 1.0	[0.78, 0.96] 0.92	[0.73, 0.94]	[0.77, 0.95] 0.91	[0.72, 0.93]
	[0.72, 1.0] 0.97	[0.65, 0.96]	[0.85, 1.0] 0.8	[1.0, 1.0] 0.0	[0.81, 1.0]	[0.79, 0.98]	[0.8, 1.0] 0.87	[0.8, 0.98]
Subcut. Emphysema	[0.9, 1.0]	[0.0, 0.0]	[0.68, 0.92]	[0.0, 0.0]	0.88 [0.79, 0.95]	[0.0, 0.0]	[0.78, 0.94]	[0.0, 0.0]
Tortuous Aorta	0.83 [0.69, 0.95]	0.79 [0.64, 0.91]	0.88 [0.75, 0.97]	0.94 [0.84, 1.0]	0.85 [0.75, 0.93]	0.86 [0.75, 0.94]	0.84 [0.73, 0.93]	0.85 [0.74, 0.93]
Macro-averages								
CheXpert-5 (CXP-5)	0.75 [0.71, 0.78]	0.72 [0.69, 0.76]	0.82 [0.79, 0.85]	0.87 [0.84, 0.9]	0.78 [0.75, 0.8]	0.78 [0.75, 0.81]	0.67 [0.64, 0.71]	0.69 [0.65, 0.72]
CheXpert-7 (CXP-7)	0.71 [0.66, 0.75]	0.65 [0.62, 0.68]	0.77 [0.73, 0.81]	0.87 [0.83, 0.9]	0.73 [0.7, 0.77]	0.71 [0.69, 0.74]	0.65 [0.61, 0.69]	0.65 [0.61, 0.68]
CheXpert-13 (CXP-13)		0.63 [0.61, 0.65]	0.68 [0.63, 0.73]	0.81	0.63 [0.6, 0.65]	0.62 [0.6, 0.64]	0.57 [0.54, 0.6]	0.56 [0.54, 0.59]
LT-only	0.73 [0.69, 0.77]	0.53 [0.48, 0.57]	0.75 [0.7, 0.8]	0.73 [0.69, 0.77]	0.72 [0.68, 0.75]	0.59 [0.55, 0.62]	0.71 [0.67, 0.74]	0.59 [0.55, 0.63]
CXR-LT	0.7 [0.67, 0.72]	0.58 [0.56, 0.6]	0.71 [0.68, 0.75]	0.77 [0.74, 0.8]	0.67 [0.65, 0.69]	0.61 [0.58, 0.62]	0.64 [0.61, 0.66]	0.58 [0.55, 0.6]
Micro	0.69	0.62	0.76	0.8	0.72	0.7	0.67	0.64

C.3 Scene Graph Evaluation: Finding Boxes

Table 7: Evaluation of finding bounding boxes against 6 finding classes from MS-CXR (see Sec. 4.1). We show finding-level scores, macro-averages over different subsets and the micro-average, with 95% confidence intervals (bootstrapping, n=1000). We excluded 2 of the 8 finding classes, because there are no samples that have positive annotations from MS-CXR, Chest ImaGenome and our dataset.

		MS-CXR Boecking et al. (2024)						
	[IoI]	J@30]	[Iol	P@30]	[IoT@30]			
	Ours	Chest ImaG.	Ours	Chest ImaG.	Ours	Chest ImaG.		
Atelectasis	0.28 [0.12, 0.42]	0.1	0.5 [0.34, 0.66]	0.14 [0.05, 0.28]	0.83 [0.71, 0.93]	0.85 [0.74, 0.94]		
Cardiomegaly	0.96	0.97 [0.95, 0.99]	1.0	1.0	0.96	0.99		
Consolidation	0.31	0.2 [0.1, 0.31]	0.41	0.24 [0.12, 0.35]	0.91	0.98		
Edema	0.52 [0.32, 0.71]	0.52	0.52	0.52 [0.32, 0.71]	1.0	1.0		
Pneumonia	0.48 [0.41, 0.57]	0.28	0.58	0.34 [0.26, 0.42]	0.93	1.0		
Pneumothorax		0.15	0.14 [0.1, 0.19]	0.15	0.96 [0.93, 0.98]	0.98		
Macro	0.45 [0.4, 0.5]	0.37 [0.33, 0.41]	0.53 [0.48, 0.58]	0.4 [0.36, 0.44]	0.93	0.97 [0.95, 0.98]		
Micro	0.51 [0.47, 0.54]	0.45	0.56 [0.52, 0.6]	0.48	0.94	0.98		

Table 8: Evaluation of finding bounding boxes against 18 finding classes from REFLACX (see Sec. 4.1). We show finding-level scores, macro-averages over different subsets and the micro-average, with 95% confidence intervals (bootstrapping, n=1000). Note that we excluded 11 of the 29 finding classes, because there are no samples that have positive annotations from REFLACX, Chest ImaGenome and our dataset.

		REFLACX B	igolin Lant	fredi et al. (20	021) all pha	ises
	[IoI]	J@30]	[Io]	P@30]	[IoT@30]	
	Ours	Chest ImaG.	Ours	Chest ImaG.	Ours	Chest ImaG
Abnormal mediastinal contour	0.08 [0.0, 0.31]	0.25 [0.0, 0.57]	0.08 [0.0, 0.31]	0.25 [0.0, 0.57]	1.0 [1.0, 1.0]	1.0 [1.0, 1.0]
Acute fracture	0.0 [0.0, 0.0]	0.0 [0.0, 0.0]	0.0 [0.0, 0.0]	0.0 [0.0, 0.0]	1.0 [1.0, 1.0]	0.0
Atelectasis	0.29 [0.26, 0.33]	0.15	0.47 [0.44, 0.51]	0.2	0.76 [0.73, 0.78]	0.93
Consolidation	0.39 [0.33, 0.45]	0.27	0.51 [0.45, 0.57]	0.34	0.8 [0.74, 0.85]	0.95
Emphysema	1.0 [1.0, 1.0]	1.0 [1.0, 1.0]	1.0 [1.0, 1.0]	1.0 [1.0, 1.0]	1.0 [1.0, 1.0]	1.0 [1.0, 1.0]
Enlarged cardiac silhouette	0.96 [0.94, 0.97]	0.96	1.0 [0.99, 1.0]	0.99 [0.99, 1.0]	0.96 [0.95, 0.97]	0.98
Enlarged hilum	0.5 [0.0, 1.0]	0.5 [0.0, 1.0]	0.5 [0.0, 1.0]	0.8 [0.23, 1.0]	0.5 [0.0, 1.0]	0.5 [0.0, 1.0]
Fracture	0.0 [0.0, 0.0]	0.0 [0.0, 0.0]	0.0 [0.0, 0.0]	0.0 [0.0, 0.0]	0.0 [0.0, 0.0]	0.0 [0.0, 0.0]
Groundglass opacity	0.28 [0.22, 0.34]	0.31	0.48 [0.4, 0.54]	0.38	0.77 [0.71, 0.83]	0.96
Hiatal hernia	0.19 [0.0, 0.43]	0.4 [0.17, 0.67]	0.27	0.4 [0.17, 0.67]	0.94 [0.78, 1.0]	1.0 [1.0, 1.0]
High lung volume / emphysema		0.58 [0.35, 0.79]	0.58 [0.35, 0.79]	0.58	0.9 [0.75, 1.0]	1.0 [1.0, 1.0]
Interstitial lung disease	0.5 [0.0, 1.0]	0.8 [0.0, 1.0]	0.8 [0.0, 1.0]	0.8 [0.0, 1.0]	0.8 [0.12, 1.0]	1.0 [1.0, 1.0]
Lung nodule or mass	0.18 [0.08, 0.31]	0.09	0.21 [0.1, 0.35]	0.09 [0.02, 0.2]	0.89 [0.77, 0.97]	0.91
Pleural abnormality	0.16 [0.13, 0.19]	0.14	0.2 [0.17, 0.23]	0.17	0.91 [0.88, 0.93]	0.92
Pleural effusion	0.42 [0.2, 0.65]	0.37 [0.17, 0.6]	0.53 [0.3, 0.75]	0.53 [0.31, 0.75]	1.0 [1.0, 1.0]	0.95 [0.82, 1.0]
Pleural thickening	0.0 [0.0, 0.0]	0.0 [0.0, 0.0]	0.0 [0.0, 0.0]	0.0 [0.0, 0.0]	1.0 [1.0, 1.0]	1.0 [1.0, 1.0]
Pneumothorax	0.04 [0.01, 0.08]	0.13	0.04 [0.01, 0.08]	0.13	0.9 [0.84, 0.96]	0.96
Pulmonary edema	0.51 [0.45, 0.56]	0.58	0.55 [0.5, 0.6]	0.58 [0.53, 0.63]	0.95 [0.93, 0.97]	1.0
Macro	0.34 [0.27, 0.42]	0.37 [0.3, 0.45]	0.41 [0.34, 0.49]	0.41	0.84 [0.78, 0.91]	0.87 [0.8, 0.95]
Micro	0.45 [0.44, 0.47]	0.42 [0.4, 0.43]	0.54 [0.53, 0.56]	0.46	0.87 [0.86, 0.88]	0.95 [0.94, 0.96]

D DATASET STRUCTURE

D.1 Scene Graph Structure

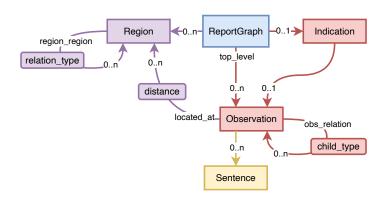


Figure 13: Scene graph structure overview.

Sentence Nodes Sentence nodes are directly associated with raw sentences in the report, i.e. there is exactly one sentence node per identified report sentence. They contain the following attributes:

• **sent_id**: Identifier, unique per study. Example: S01.

• **section**: Name of the section that the sentence belongs to, as specified in the report. If the sentence is not part of a section, FINAL_REPORT_NO_SECTION or PRE_FINAL_REPORT_NO_SECTION are used.

Examples: FINDINGS, IMPRESSION, REASON_FOR_EXAM.

• **section_type**: The identified type of section used for classifying the type of content of the sentence. IGNORE is used for irrelevant sections.

Examples: FINDINGS, IMPRESSION, INDICATION.

• **sentence**: The raw sentence as written in the report.

Observation Nodes Observation nodes are created for each individually described aspect (i.e. observation) in the report's FINDINGS or IMPRESSION section. Hereby, a single sentence may be related to several observation nodes and a single observation may be derived from several sentences (if they describe related aspects). Observation nodes are structured hierarchically, i.e. they may have other observation nodes as parents. An observation node contains the following attributes:

• **obs_id**: Identifier, unique per study. Example: 001.

For child nodes this also contains the parent id, e.g. 001.02.

- **summary_sentence**: Textual description of the observation, directly derived from the associated report sentences. In some cases, this may be an exact copy of the report sentences but it may also paraphrase parts of it.
- name: Abbreviated version of the summary_sentence.
- child_level: Hierarchy level, 0 for top-level, larger numbers for deeper hierarchy levels.
- child_type: Type of parent-child relation.

 Possible options: regional_distinction, related_region, associated_with, device_part, recommendation, comparison_only.
- **regions**: List of associated regions, each paired with an optional list of distance annotations.

Example: [("heart", ["1 cm above"])]

 non_resolved_regions: Similar to regions but with regions that could not be semantically mapped to reference definitions.

1237 1238

1239

1240

1241

1188 1189	• laterality: Laterality of the region. Possible options: left, right, likely bilateral, bilateral, unknown.
1190	• default_regions : List of regions that have been added because they are defaults for
1191	the identified findings (obs_entities).
1192	• obs_entities: List of (directly) associated findings.
1193	Example: ["pleural effusion"].
1194	• • • •
1195	• obs_entities_parents : List of findings that are considered parents of findings in obs_entities.
1196	
1197 1198	• non_resolved_obs_entities: Similar to obs_entities but with findings that could not be semantically mapped to reference definitions.
1199 1200	 obs_categories: List of associated finding super-categories. Example: ["ANATOMICAL_FINDING"].
1201 1202	• obs_subcategories : List of associated finding sub-categories. Example: ["LUNG_FIELD"].
1203	•
1204 1205	• probability : Likelihood of the observation being positive. Short term, derived from what is mentioned in the report.
1205	• certainty: How certain is the observation. Derived from probability.
1207	Possible options: certain, likely, uncertain, comparison_only,
1208	recommendation.
1209	• positiveness: Whether the observation is positive or negative. Derived from
1210	probability.
1211	Possible options: pos, neg, comparison_only, recommendation.
1212	• modifiers: Modifiers of the finding. Dictionary with keys for each type of modifier and
1213	lists of the individual modifier values.
1214	Possible modifier type: severity, texture, spread, temporal. Example: {"severity": ["mild"], "spread": ["focal"]}.
1215	
1216 1217	• change_sentence : Optional textual description of any changes to the prior study of the same patient, if it was mentioned in the report.
1218	• changes: List of change types mentioned in the change_sentence.
1219	Example: worsening.
1220	• from_report : Whether this observation was explicitly mentioned in the report (true) or automatically added (false).
1221	
1222	• obs_quality : Extraction quality of the observation, consisting of several individual
1223 1224	aspects. See Tab. 9.
1224	• localization: Bounding boxes for this observation, for each associated image. Dictio-
1225	nary with keys equaling image ids (each study may correspond to several images). Values
1227	contain:
1228	- image_id
1229	- bboxes : List of bounding boxes in the (x_1, y_1, x_2, y_2) format in original image-pixel
1230	coordinates.
1231	- localization_reference_ids: List of region names from which the bounding
1232	boxes are derived.
1233	 missing_localization: List of associated region names for which no localization is available for this image.
1234	
1235	 is_fallback: Whether this localization is a fallback, i.e. the original region localization was not available but a more coarse localization was used instead.
1236	- localization guality: Quality of the localization. See Tab. 9

Region Nodes Region nodes are created for each anatomical structure mentioned in any observation and for key regions. They contain the following attributes:

- localization_quality: Quality of the localization. See Tab. 9.

• region: Name of the region and unique identifier within each study. Example: left lung.

Under review as a conference paper at ICLR 2026 1242 • **laterality**: Laterality of the region. 1243 Possible options: left, right, bilateral, unknown (i.e. not clearly definable). 1244 localization: Bounding boxes for this region, for each associated image. Same format 1245 as for observation nodes. 1246 • region_localization_quality: Quality of the localization attribute. See Tab. 9. 1247 1248 **Indication Node** Each study contains an optional indication node with information extracted from 1249 the INDICATION section. If present, it contains the following attributes: 1250 1251 • indication_summary: Summary of the indication, directly derived from the 1252 INDICATION section of the report, but typically paraphrased. 1253 • patient_info: Any information about the patient, if mentioned in the INDICATION section. A subset of the content in indication_summary. 1255 • indication: Indication for the study, if mentioned in the INDICATION section. A subset of the content in indication_summary. 1257 • evaluation: Any required evaluation of the patient (i.e. what should be evaluated with this study), if mentioned in the INDICATION section. A subset of the content in 1259 indication_summary. 1261 • associated_sentence_ids: List of sent_ids from sentence nodes that are related to the indication. 1262 1263 • associated_obs_ids: List of obs_ids from observation nodes that are related to the 1264 indication. 1265 • answer_for_indication: A single observation node containing the answer to the 1266 question (implicitly) asked by the provided indication. This is a special observation node 1267 with obs_id = OIND. Its textual description is directly derived from the FINDINGS and 1268 IMPRESSION sections but conditioned on the INDICATION section. 1270 1271

Root Node and Relations Each study contains a single root node called the ReportGraph. It contains general metadata about the study and its scene graphs:

- patient_id: Unique patient ID, the subject_id from MIMIC-CXR.
- **study_id**: Unique study ID, from MIMIC-CXR. Each patient may have several studies.
- study_quality: The overall extraction quality of the scene graph for this study, consisting of several individual aspects. See Tab. 9.
- study_img_localization_quality: Dictionary of localization qualities for each image with keys corresponding to image IDs. See Tab. 9.

Additionally, it is connected to all other nodes and links to the top-level (root) observations. Thus, it contains the following:

• **sentences**: List of all sentence nodes.

1272

1274

1278

1279 1280

1281

1282

1283 1284

1285

1286

1290 1291

1293

1294

1295

- **observations**: Dictionary of all observation nodes, indexed by their obs_id.
- top_level_obs_ids: List of all top-level (root) observation node IDs, i.e. their obs_ids.
- **regions**: Dictionary of all region nodes, indexed by their region attribute.
- indication: The indication node, if it exists.

Nodes can also be connected by the following relations:

- located_at_relations (observation \rightarrow region): Specifies where an observation is located with the following additional attributes:
 - **distances**: List of distance annotations, e.g. ["3cm above"].
 - where_specified: How this relation was derived. Possible options: direct, bilateral, sub_region.

1296	• obs_relations (observation↔observation): Specifies a parent-child relation
1297	between two observations, with the following additional attribute:
1298	
1299	- child_type: Type of parent-child relation.
1300	Possible options: regional_distinction, related_region, associated_with.device_part, recommendation, comparison_only.
1301	associated_with, device_part, recommendation, comparison_only.
1302	• obs_sent_relations (observation↔sentence): Specifies from which sen-
1303	tences an observation was derived.
1304	• region_region_relations (region \rightarrow region): Specifies a relation between two

regions with the following additional attribute: - relation_type: Type of relation.

Possible options: sub_region, bilateral (the bilateral version of a region), left (the left version of a region), right (the right version of a region).

D.2 QUESTION-ANSWER STRUCTURE

1305 1306

1309 1310

1311

1312

1313

1314 1315

1316

1317

1318 1319

1320

1321

1322

1324

1326

1328

1332

1333

1334

1335 1336

1337

1338

1339

1340

1341

1347

1348 1349 **QA-Pair** Each question-answer pair consists of a free-text question (attribute question), an answer consisting of structured answer parts (attribute answers). Additionally, it contains the following metadata:

- question_id: Identifier, unique within the associated study.
- question_type: The QA-template used to generate this QA-pair.
- question_strategy: The strategy used to generate QA-pair. See Secs. 3.2 and E.2.2.
- **variables**: Key-value pairs of variables (and their values) used during generation, e.g. to fill the template. See Sec. E.2.1.
- obs_ids: List of obs_idss of observation nodes (in the scene graph) from which the
 answer is derived.
- contains_report_answers: Whether any of the answer parts was derived from the report, i.e. from observation nodes.
- contains_template_answers: Whether any of the answer parts was generated based on a template.
- **extraction_quality**: The overall extraction quality of the associated observations in the scene graph, consisting of several individual aspects. See Tab. 9.
- question_img_localization_quality: Quality of the localizations per image. See Tab. 9.
- question_quality: The overall question-answer text quality, consisting of several individual aspects. See Tab. 10.
- rating: The overall rating of the QA-pair. See Sec. D.3

Answers are structured hierarchically, consisting of a list of answer parts (attribute answers) and sub-answers (children) of these answers, where there can be several hierarchy levels. The hierarchy levels are derived from the parent-child structure of associated observation nodes (based on obs_relations, Sec. D.1). Additionally, there are different types of answer parts:

- main_answer: Required to answer the question. There is always at least one main-answer per question.
- **details**: Providing additional details for the main answer, which are however not mandatory to answer the question.
- related_information: Not directly answering the question, but may be related and provides context.

Each individual answer part contains the following attributes:

• answer_id: Identifier, unique within each study. Contains the question_id.

hierarchy levels (sub-answers).

• answer_type: Type of answer part.

The answer text.

summary_sentence in the observation node (Sec. D.1).

• answer_level: Hierarchy level, 0 for top-level answer part, larger numbers for deeper

Possible options: main_answer, details, related_information.

1350

1351

1352

1353

1354

1355

1356

1401

1402

1403

the report.

1357 1358		• name_tag: Abbreviated version of the text. Either generated from a template or based on name in the observation node (Sec. D.1).
1359		• laterality: Laterality of the region. See observation node (Sec. D.1).
1360		Possible options: left, right, likely bilateral, bilateral, unknown.
1361		• regions: List of associated regions. See observation node (Sec. D.1). Distances are not
1362		provided here.
1363		Example: ["heart"]
1364 1365		• obs_entities : List of (directly) associated findings. See observation node (Sec. D.1). Example: ["pleural effusion"].
1366		• obs_entities_parents: List of findings that are considered parents of findings in
1367		obs_entities. See observation node (Sec. D.1).
1368		• obs_categories: List of associated finding super-categories. See observation node
1369		(Sec. D.1).
1370		Example: ["ANATOMICAL_FINDING"].
1371		• obs_subcategories: List of associated finding sub-categories. See observation node
1372		(Sec. D.1).
1373		Example: ["LUNG_FIELD"].
1374		• certainty : How certain is the observation. See observation node (Sec. D.1).
1375		Possible options: certain, likely, uncertain, comparison_only,
1376		recommendation.
1377		• positiveness: Whether the observation is positive or negative. See observation node
1378 1379		(Sec. D.1).
1380		Possible options: pos, neg, comparison_only, recommendation.
1381		• modifiers: Modifiers of the finding. List of pairs of modifier type and value. See
1382		observation node (Sec. D.1).
1383		Possible modifier type: severity, texture, spread, temporal.
1384		<pre>Example: [("severity", "mild"), ("spread", "focal")].</pre>
1385		• localization: Bounding boxes for this answer part, for each associated image. Dic-
1386		tionary with keys equaling image ids (each study may correspond to several images). See
1387		observation node (Sec. D.1).
1388		• sub_answers: List of child answers (deeper in the hierarchy). Each sub-answer is another
1389		answer-part with all attributes and potentially further sub-answers.
1390		• from_report: Whether this answer part is derived from the report, i.e. an observation
1391		node (true), or from a template (false).
1392		• extraction_quality: The overall extraction quality of the associated observations in
1393		the scene graph, consisting of several individual aspects. See Tab. 9.
1394		• answer_quality: The overall answer text quality, consisting of several individual as-
1395		pects. See Tab. 10.
1396		
1397	D.3	QUALITY
1398	Dat!	We distinguish between the fellowing consult actions for each OA action
1399	Katii	ngs We distinguish between the following overall ratings for each QA-pair:
1400		• A++: Perfect and complete content; all information in the answer is explicitly mentioned in

Either generated from a template or based on

tioned in the report, but some minor details may be missing or irrelevant.

• A+: Perfect and mostly complete content; all information in the answer is explicitly men-

- A: Very good content with minor issues not affecting the overall quality; some tags or boxes may be inferred or minor issues (e.g. grammatical) may be present in the text.
- **B**: Good content; factually correct answers, which may however be not fully complete or slightly unclear.
- C: Poor content; answers may be misleading or contain completely unclear information.
- **D**: Incorrect content; answers may be contradicting the report and are not usable.
- not rated: Quality could not be assessed, e.g. due to invalid LLM-rater outputs.

These ratings are derived based on individual aspects that will be described in the following paragraphs. Possible quality levels for each aspects and the resulting rating are presented in Tabs. 9 and 10. The final rating is computed as the minimum (worst) rating over all individual aspects.

Scene Graph Extraction Quality For each scene graph, we provide a quality rating based on how well it could be constructed/extracted. Tab. 9 shows the considered aspects with their potential quality levels and resulting ratings.

Table 9: Quality levels for the 6 scene graph quality aspects, with their resulting ratings.

	Quality level	Value	Resulting rating
How w	ell are region tags identified? (attribute regio	ns)	
Region extraction	NO_REGIONS	0	В
	DEFAULT_REGIONS_ONLY	1	В
	CONTAINS_DEFAULT_REGIONS	2	A
	CONTAINS_NON_RESOLVED_REGIONS	3	A
	RESOLVED_REGIONS_ONLY	4	A++
How w	ell are finding tags identified? (attribute obs_e	ntitie	s)
gr ion	NO_ENTITIES	0	В
Finding extraction	CONTAINS_NON_RESOLVED_ENTITIES	1	A
ex E	RESOLVED_ENTITIES_ONLY	2	A++
How w	ell are textual descriptions extracted? (attribut	es summa	ary_sentence and name)
tion ion	CHANGE_IN_SENTENCE_OR_NAME	0	В
Description extraction	UNDERSCORES_IN_SENTENCE_OR_NAME	1	A
Des	NO_ISSUES	2	A++
How w	ell are mentions of change extracted? (attribute	es chanç	ge_sentence and change
	CHANGE_SENTENCE_REMOVED	0	В
Change extraction	UNDERSCORES_IN_CHANGE_SENTENCE	1	A
Che extra	CONTAINS_NON_RESOLVED_CHANGES	2	A
	NO_ISSUES	3	A++
Have tl	nere been any issues in the extraction and scen	e graph	construction pipeline?
	DISCARDED	-1	D
п	NON_INTERPRETABLE	0	C
Extraction issues	MOSTLY_INTERPRETABLE	1	В
Extra	IGNORABLE	2	A
_	FIXABLE	3	A+
	NOLISSUES	4	A++
How w	ell could observations/regions be localized? (at	ttribute 1	ocalization))
_	NO_LOCALIZATION	0	В
ation	FALLBACK_LOCALIZATION	1	В
Localization	INCOMPLETE_LOCALIZATION	2	A
	BBOX_LOCALIZATION	3	A++
	BBOX_AND_MASK_LOCALIZATION	4	A++

Finding extraction is also referred to as entity extraction, description extraction as sentence/name quality.

QA Text Quality For each QA-pair, we provide quality rating for its text, i.e. the question text and the textual descriptions in its answer parts. Tab. 10 shows the considered aspects with their potential quality levels and resulting ratings.

Table 10: Quality levels for the 5 QA-pair text quality aspects, with their resulting ratings.

	Quality level	Value	Resulting rating
Does	s the answer factually align with the original	report?	
(rate	d per answer-part, given the question and the re-	port)	
	NON_ALIGNED_CONTRADICTING	-3	D
	NON_ALIGNED_MISLEADING	-2	C
ent	NON_ALIGNED_NON_INFERABLE	-1	В
Entailment	ALIGNED_GENERAL_STATEMENT	0	A
En	ALIGNED_NEGATIVE_NOT_MENTIONED	1	A+
	ALIGNED_INFERABLE	2	A++
	ALIGNED_MENTIONED	3	A++
Is th	e answer relevant for the given question?		
(rate	d per answer-part, given the question but indepe	endent of t	he report)
v	IRRELEVANT_INFO	-2	A
Relevance	REDUNDANT_INFO	-1	A
Rele	RELATED_INFO	0	$A \hbox{++} (A \hbox{+++} for \verb related_information answer)$
_	RELEVANT_MAIN_ANSWER	1	A++ (A for related_information answer)
Does	the answer cover all aspects in the report th	at are rel	evant to the question?
(rate	d for the full answer, given the question and the	report)	
×	INCOMPLETE_MISLEADING	-2	С
Completeness	INCOMPLETE_NON_MISLEADING	-1	В
ıplet	NOT_ANSWERED	0	В
Com	DETAILS_MISSING	1	A+
	FULLY_COMPLETE	2	A++
Is th	e generated question clear and grammaticall	y correct	?
(rate	d for the question, given nothing else)		
	UNANSWERABLE	-3	С
arity	UNRELATED_TO_CHEST_XRAY	-2	В
Question clarity	UNCLEAR_QUESTION	-1	В
estic	GRAMMATICAL_ERRORS	0	A
Õ	UNUSUAL_SENTENCE_STRUCTURE	1	A
	OPTIMAL	2	A++
Is th	e answer clear and grammatically correct?		
(rate	d per answer-part, given nothing else)		
5.	NOT_UNDERSTANDABLE	-2	С
larit	UNCLEAR_ANSWER	-1	В
ver c	GRAMMATICAL_ERRORS	0	A
Answer clarity	UNUSUAL_SENTENCE_STRUCTURE	1	A
7	OPTIMAL	2	A++

E DATASET CONSTRUCTION DETAILS

E.1 Scene Graph Construction

E.1.1 REGION LOCALIZATION

 We use the CXAS (Seibold et al., 2022; 2023) model to predict segmentation masks of 158 anatomical structures on the 377,110 CXRs from MIMIC-CXR-JPG (Johnson et al., 2024a; 2019b; Goldberger et al., 2000). Additionally, we use the bounding boxes provided by the Chest ImaGenome (Wu et al., 2021a;b; Goldberger et al., 2000) dataset, which are provided for 29 anatomical structures in most frontal images of MIMIC-CXR. The masks predicted by CXAS are post-processed with morphological operations to filter out outlier pixels.

We specify 257 localized regions in our reference definitions. For each of these regions, we define how the bounding boxes are derived. We consider the following options:

- CXAS masks: Some regions are directly associated with one of the 158 anatomical structures for which the CXAS model predicts segmentation masks. In these cases, we compute the bounding box around the predicted segmentation mask.
- Chest ImaGenome boxes: Some regions are directly associated with one of the 29 anatomical structures for which Chest ImaGenome provides bounding boxes. In such cases, we use these provided bounding boxes if no CXAS masks are associated.
- **Bilateral regions**: Some regions refer to a pair of bilateral regions (e.g. *lungs* refers to *left lung* and *right lung*). In these cases, we simply use the two bounding boxes of the left and right versions, but do not fuse them.
- Parent regions: For some regions we do not have exact correspondences to available masks
 or boxes but we have available sub-regions. In these cases, we compute the super bounding
 box, a single box, around all specified child regions.
- Fusions: In some rare cases, we combine multiple individual masks or bounding boxes. We compute intersections or unions of boxes or masks, before inferring the final bounding box.

After computing all regions, we filter out regions with a too small bounding box area. For images where a specific region is not available, we try to use alternative regions as fallbacks instead, e.g. using a more coarse parent regions as an alternative. Note that this is often the case for lateral images as there no Chest ImaGenome boxes are available.

E.1.2 Information Extraction

Extracting the Sentences First, we extract individual sentences from the reports, detect their sections (e.g. FINDINGS, IMPRESSION, INDICATION, ...), discard sentences without relevant information, and merge sentences containing similar information (e.g. if findings are described in both the FINDINGS and IMPRESSION section). Therefore, each full report is passed in a single step to the LLM, which predicts the individually separated sentences as well as their sections and related sentences. We use the prompt shown in Listing 1 (with few-shot examples similar to Listing 2) and apply it to the full radiology report. After parsing the LLM outputs, we apply the Stanza (Qi et al., 2020) tokenizer to each identified sentence and try to further split it. The LLM also identified potentially related sentences. We use this information to identify sentence clusters containing related information. Such sentence clusters are the basis for the next step, i.e. observation extraction. We successfully extracted sentence from 227 626 studies (reports) while having parse errors for 209 studies.

```
1566
                          Listing 1: LLM prompt used for sentence extraction.
1567
       Extract all sentences from the given textual report.
1568
       You will be given a (free-text) medical radiology report describing one
1569
           \hookrightarrow or more chest X-rays of a single patient.
1570
       # Rules:
1571
       - Split the report into sentences and extract all sentences in the report
1572
1573
       - Do not rewrite the sentences!
1574
       - For each sentences, identify the its section name (written in the
1575
           \hookrightarrow report).
       If a sentence is not part of a section but is part of the "FINAL REPORT",
1576
           \hookrightarrow then use "FINAL\_REPORT\_NO\_SECTION". If a sentence is not part
1577
           \hookrightarrow of a section but the sentence is before the "FINAL REPORT", use the
1578
           \hookrightarrow section name "PRE\_FINAL\_REPORT\_NO\_SECTION".
1579
       - For each sentence, classify the content written therein into one of the
1580
              following types: [EXAM\_TECHNIQUE, INDICATION, FINDINGS,
           \hookrightarrow IMPRESSION, PRE\_FINAL\_REPORT, IGNORE]. This is typically inferred
1581
           \hookrightarrow from the section name but may also be influenced by the content of
1582
           \hookrightarrow the sentence. Some example sections names for each type are given
1583
           \hookrightarrow below:
1584
           EXAM\_TECHNIQUE: EXAMINATION, EXAM, TECHNIQUE
1585
           INDICATION: INDICATION, INDICATIONS, HISTORY, CLINICAL HISTORY,
                \hookrightarrow CLINICAL, REASON, REASON FOR EXAM
1586
           FINDINGS: FINDING, FINDINGS
1587
            IMPRESSION: IMPRESSION, IMPRESSIONS, RECOMMENDATION
1588
           PRE\_FINAL\_REPORT: WET\_READ, WET\_READ\_VERSION\_#1, PRE\_FINAL\
1589
                \hookrightarrow _REPORT\_NO\_SECTION
1590
           IGNORE: COMPARISON, COMPARISONS, REFERENCE EXAM, NOTIFICATION
       - Split the report into individual sentences and report each sentence in
1591
           \hookrightarrow its own line, removing any newlines present in the sentence.
1592
       - For enumerations: each point is considered an independent sentence!
1593
           \hookrightarrow Remove the numbering.
1594
       - Specify sentence IDs of similar, previous sentences that each sentence
1595
           \hookrightarrow could be merged with. A sentence should be merged with all previous
           \hookrightarrow sentences that either describe the same aspect or that refer to
1596
           \hookrightarrow each other (e.g. if a sentence provides further details to a
1597
           \hookrightarrow previous one). A bullet point may also be associated with a
           \hookrightarrow sentence, even if the other sentence has a different bullet number
1599
           \hookrightarrow or none at all.
1600
       - Follow the examples given below!
1601
       # Examples:
1602
       <FEWSHOT>
1603
1604
       # Input Report (extract data from this report):
1605
       --- START OF REPORT -
       <REPORT>
1606
       --- END OF REPORT ---
1608
       # Hints:
1609
         Infer the output format from the examples!
1610
       - Do not add any explanations or text BEFORE or AFTER the extracted
           \hookrightarrow sentences, i.e. start with the first sentence!
1611
1612
       # Proceed:
1613
1614
```

1637

1638

1639

1640

1641

1642

1643

1644

1645

1646

1647 1648

1672 1673

Listing 2: Few-shot example for sentence extraction.

```
1621
       ***Example: Report***
1622
       --- START OF REPORT ---
1623
                                        FINAL REPORT
1624
      PORTABLE CHEST OF \_\_\_
1625
      COMPARISON: \_\_\ radiograph.
1626
1627
      FINDINGS: No pleural effusion or pneumothorax.
1628
      --- END OF REPORT ---
1629
       ***Example 5: Output***
1630
       [S01] FINAL\_REPORT\_NO\_SECTION(EXAM\_TECHNIQUE) - merge with []:
1631
          → PORTABLE CHEST OF \_\_\_
1632
       [S02] COMPARISON(IGNORE) - merge with []: \_\_\ radiograph.
1633
       [S03] FINDINGS(FINDINGS) - merge with []: No pleural effusion or
1634

→ pneumothorax.

1635
```

Extracting the Observations In this step, we consider each sentence cluster (as identified during sentence extraction), in the FINDINGS and IMPRESSION sections. A sentence cluster contains one or more sentences that describe related aspects and may stretch over one of both of these sections. From each of these clusters, we now extract mentioned observations using the prompt shown in Listing 3 with few-shot examples similar to Listing 4. We apply this prompt to each sentence cluster individually and extract zero, one, or multiple observations each. The output is provided in the json-format and follows a similar structure as the final observation node, but we optimized it to be easy to fill by the LLM. The LLM is allowed to freely assign values to each of the json-fields. For name and summary_sentence, we prompt the model to stay close to the original sentence, but it must remove any mentions of change and only keep the part relevant to the individual observation (if several observations are mentioned in one sentence). We successfully extracted observations from 227 266 studies (reports) while having parse errors for 360 studies.

Listing 3: LLM prompt used for observation extraction.

```
1649
       Extract structured information from the given textual report.
1650
       You will be given sentences from a (free-text) medical radiology report
1651
           \hookrightarrow describing one or more chest X-rays of a single patient.
1652
1653
       # Guidelines:
       <GUIDE>
1654
1655
       # Rules:
1656
       - Follow the examples given below!
1657
1658
       # Examples:
       <FEWSHOT>
1659
1661
         Check for any "change" modifiers (see guidelines).
1662
       - If there is a "change" modifier, rewrite the "summary\_sentence" such
1663
           \hookrightarrow that it describes only what is visible in the current image,
           \hookrightarrow without any mentions of change or comparisons! Describe the change
1664
           \hookrightarrow in the "change\_sentence". Do this for all top-level AND child
1665
           \hookrightarrow observations.
       - Make sure to include all children of observations, even if they repeat
           \hookrightarrow information from the parent!
1668
1669
       # Proceed with the Input Sentence:
       Sentence(s): <SENT>
       Output JSON-List:
1671
```

1695

1696

1698

1700

1701

1702

1703

Listing 4: Few-shot example for observation extraction.

```
1675
       Sentence(s): Left more than right basilar atelectasis.
1676
       Output JSON-List:
1677
1678
                "name": "bibasilar atelectasis", "entity": "atelectasis",
1679
                "probability": "positive", "change": null,
                "summary\_sentence": "Bibasilar atelectasis.",
1681
                "change\_sentence": null,
1682
                "regions": ["bibasilar"],
1683
                "children": [
1684
                         "child\_type": "regional\_distinction", "name": "left
1685
                            \hookrightarrow basilar atelectasis", "entity": "atelectasis",
1686
                         "probability": "positive", "change": null,
1687
                         "summary\_sentence": "Left more than right basilar
1688
                            \hookrightarrow atelectasis.",
                         "change\_sentence": null,
1689
                         "regions": ["left basilar"]
1690
1691
                ]
           }
1693
1694
```

Extracting the Indication Next, we extract information about the INDICATION section and detect which FINDINGS or IMPRESSION sentences may provide information related to the indication. Therefore, the extracted INDICATION sentences and a list of all FINDINGS and IMPRESSION sentences are passed to the LLM using the prompt shown in Listing 5 with few-shot examples similar to Listing 6. The LLM predicts a json-structure containing several text fields for summaries of aspects in the indication, an answer_for_indication derived from the FINDINGS and IMPRESSION section, as well as relevant sentence IDs. We successfully extracted indictions from 227 596 studies (reports) while having parse errors for 30 studies.

Listing 5: LLM prompt used for indication extraction.

```
1704
1705
       Extract structured information from the given (free-text) medical report.
1706
       You will be given the indication sentence from a report and additionally
1707
           \hookrightarrow the sentences from the findings section.
1708
       # Rules:
1709
       - Extract / summarize the given indication information. Use only the
1710
           \hookrightarrow provided indication sentence.
1711
       - Additionally, identify the finding sentences associated with the
1712
           \hookrightarrow indication, i.e. the sentence that answer the quesiton of the
           \hookrightarrow indication or are highly relevant to it. Based on these finding
1713
           \hookrightarrow sentences, provide an answer to the question asked in the
1714
           \hookrightarrow indication.
1715
       - Follow the examples given below!
1716
       # Examples:
1717
       <FEWSHOT>
1718
1719
       # Hints:
1720
       - For each attribute, write full sentences instead of single terms or
1721
           \hookrightarrow bullet points.
       - In the "answer\_for\_indication", describe in YOUR OWN WORDS how the
1722
           \hookrightarrow question asked in the evaluation can be answered based on the
1723
           \hookrightarrow findings. Only include the key information.
1724
       - Use the JSON structure from the examples!
1725
1726
       # Proceed with the Input:
1727
       **Input: **
       INDICATION: <IND>
```

```
1728

1729

1730

1731

**Output JSON:**
```

Listing 6: Few-shot example for indication extraction.

E.1.3 BUILDING SCENE GRAPHS

Entity Mapping We apply semantic entity mapping to modifiers (used to fill the attributes probability, certainty, positiveness, and modifiers), regions (attribute regions), finding entities (attribute obs_entities), and changes (attribute changes).

For each of these we consider the associated tags extracted by the LLM during observation extraction and encode them into text embeddings using the BioLORD (Remy et al., 2024) model. We also encode all potential tags and their synonyms, defined for each type of tag in our reference definitions. Then we compute the cosine similarities of each tag with all reference tags of the same type. We pick the reference tag with the highest cosine similarity but threshold it at 0.5. If no reference tag was identified with cosine similarity ≥ 0.5 , then we mark the tag as non-resolved. For finding entities, we follow a slightly more complicated matching approach. Instead of only considering the finding entity tags extracted by the LLM, we also consider pairs of these entities and extracted region tags as well as the extracted summary sentences and names for matching. We then try to match each of those with the reference finding tags and pick the ones with the highest cosine similarities.

The matched reference finding tags are stored in the <code>obs_entities</code> attribute (non-resolved ones are kept in <code>non_resolved_obs_entities</code>), matched reference regions are stored in the <code>regions</code> attribute, where we also store the distance as identified by the LLM (non-resolved regions are kept in <code>non_resolved_regions</code>). The matched changes are store in the <code>changes</code> attribute (non-resolved changes are discarded). For all modifiers, we use the modifier type defined for the matched reference tag. We matched all modifiers against all types of modifiers, which means that the modifier type identified by the LLM can be overwritten during matching. Finally, we extract the probability from the modifiers (this is a special modifier type), store it in the <code>probability</code> attribute and infer the <code>certainty</code> and <code>positiveness</code> attributes from it (using the reference definitions). The remaining modifiers are stored in the <code>modifiers</code> attribute (non-matched ones are discarded).

We additionally try to identify the laterality of the observations. Here, we do not use semantic entity mapping but rely on keywords instead. We consider the raw finding entities, regions, as well as the summary sentences, and search for any laterality-related mentions such as *left*, *right*, *bilateral*, and related terms. From this we infer the laterality and store it into the laterality attribute.

Reference Data and Standardization Using the reference definitions, we infer all obs_entities_parents, obs_categories, obs_subcategories, and default_regions from the matched obs_entities.

Next, we inspect the <code>summary_sentence</code> and <code>name</code> attributes (extracted by the LLM) for underscores or mentions of changes. We track such issues (which are used for quality assessment) but do not apply any cleanup. Similarly, we check the <code>change_sentence</code> for underscores and assert that it contains mentions of changes.

We further inspect the structure of observations and their children. If an observation mentions multiple different findings and has one child for each of these findings, then we lift these children to the top-level and discard the parent. Similarly, we merge multiple duplicate observations into one.

Finally, we try to resolve missing regions or improve their precision. If no regions could be extracted, we rely on the default_regions derived from the obs_entities instead, but consider the identified laterality. We also check whether these default_regions are more precise than extracted ones. Then we check whether any identified region contradicts the identified laterality and remove them. We then either split or merge bilateral versions of the same region.

Graph Construction Based on the matched regions, we associate bounding boxes with the observations if available. Additionally, we build a tree of all mentioned regions and fill missing intermediate regions based on the reference data. This allows us to build a graph of region nodes relevant to the study.

We construct region_region_relations based on the reference data alone. located_at_relations are constructed based on the regions attribute of observations (direct specified). Additionally, we infer located_at_relations relations for sub regions (sub_region) and bilateral versions of regions (bilateral). obs_relations are constructed based on the parent-child structure of observations and their child type, as predicted by the LLM. obs_sent_relations are constructed based on the sentences each observation was derived from.

Finally, we attach the indication information extracted from the report. Therefore, we build an additional observation node based on the LLM-extracted answer_for_indication and the LLM-extracted associated sentences, from which we can infer the associated observations and can infer all relevant tags.

E.2 QUESTION-ANSWER GENERATION

E.2.1 TEMPLATE ENGINE

To construct QA-pairs, we develop a template-engine that considers the information in the scene graphs to construct the answers. The template engine generates a QA-pair by running the following steps:

- 1. Filter observations and studies based on the template configuration, e.g. only keeping observations of specific sub-categories.
- 2. Run a QA-strategy (indication, study abnormality, region abnormality, or finding) on the remaining scene graph. The strategy provides multiple named subsets of observations, variables to fill the template, as well as an overall state consisting of multiple tags (e.g. is the study positive, are there any devices, ...).
- 3. Construct the template-based main answer by selecting and filling the answer-template based on the state returned by the QA-strategy and the returned variables. Tags and bounding boxes can be inferred from defined observation subsets. (Not all templates provide such main answers)
- 4. Pick observation subsets identified by the QA-strategy and convert the observations into answer parts. The template configuration defines which subsets are picked and how they are ordered. Additionally, template-based prefix- or fallback-answers can be defined for each subset. Some subsets can also be excluded based on the QA-strategy state. These answers can be main-answers, details, or related information as defined in the configuration.

Additionally, the template engine supports variables, i.e. each template can be used to generate multiple QA-pairs. Variables can either be defined as lists (configured in the template) or can be

provided by the QA-strategy (which might infer variables from the current scene graph, e.g. all mentioned regions). The question may then also contain such template variables.

E.2.2 STRATEGIES AND TEMPLATES

Indication In this strategy, we use the extracted indication (if available) as the question. More precisely, we use the indication_summary attribute from the indication node as the question text. The main-answer is constructed from the indication observation (i.e. the answer to the indication based on the finding sentences), while detail answer parts are constructed based on all associated finding observations. We include this question, if an indication observation is present in the scene graph.

Study abnormality In this strategy, we generate questions about abnormalities. This includes descriptions of the full study or specific categories of observations (e.g. devices), description of only abnormal findings, and yes/no questions of whether there are positive findings (overall or of specific categories) present in the study. We use the templates defined in Tab. 11.

The strategy identifies five types of observations: (i) finding (positive), (ii) finding (negative), (iii) device (positive), (iv) device (negative), (v) acquisition. Based on the specific template, these are selected as main-answers, details, or related information. Additionally, a template answer can be included, which is selected based on whether the study is abnormal or not. Some templates use different subcategories as variables, i.e. one question is generated for each of the defined subcategories, where observations are filtered based on this subcategory.

Table 11: Study abnormality templates.

Template (ID)	Question Example	Variables	Main answer	Details	Related Inf.
describe_all B01_describe_all	Describe the given study.	-	finding (positive) device (positive) device (negative) finding (negative) acquisition	-	-
describe_abnormal B02_describe_abnormal	Describe all abnormal findings in the given study.	-	finding (positive)	-	device (positive)
is_abnormal B03_is_abnormal	Are there any abnormal findings?	-	template finding (positive)	finding (negative)	device (positive) finding (negative)
is_normal B04_is_normal	Is the study normal?	-	template	finding (positive) finding (negative)	finding (negative) device (positive)
describe_subcat B08_describe_subcat	Evaluate the cardiac structures.	subcategory	finding (positive) finding (negative)	-	-
describe_abnormal_subcat B09_describe_abnormal_subcat	Describe any pulmonary diseases and disorders suggested by the study.	subcategory	finding (positive)	-	_
is_abnormal_subcat B10_is_abnormal_subcat	Are there any fractures or bone diseases apparent from the study?	subcategory	template finding (positive)	finding (negative)	finding (negative)
is_normal_subcat B11_is_normal_subcat	Are the mediastinal and hilar contours normal?	subcategory	template finding (positive)	finding (negative)	finding (negative)
describe_device B12_describe_device	Check the presence and position of devices, tubes, lines, and other foreign objects.	subcategory	device (positive) device (negative)	_	-
has_devices B13_has_devices	Are there any signs of prior surgical procedures?	subcategory	template device (positive)	device (negative)	device (negative)
describe_acquisition B14_describe_acquisition	Assess the image quality and describe aspects related to image acquisition.	-	acquisition	-	-
describe_imaging_artifacts B15_describe_imaging_artifacts	Describe any apparent imag- ing artifacts and imaging- related shadows.	_	acquisition	-	-
has_imaging_artifacts B16_has_imaging_artifacts	Are there any imaging artifacts or imaging-related shadows?	-	template acquisition	-	-

Region abnormality In this strategy, we generate question about anatomical regions. This includes describing regions, answering yes/no questions about the abnormality of regions, or describing specific aspects of regions (e.g. devices). We use the templates defined in Tab. 12.

For a given region, the strategy first identifies observations associated with that region and classifies them into the five types defined in the study abnormality strategy. Additionally, it identifies observations in related regions. This includes positive findings in parent regions or the opposite laterality. Additionally, a template answer can be included, which is selected based on whether the region is abnormal or not.

Before generating QA-pairs, the strategy first identifies a set of regions. For each of these regions an individual QA-pair is generated. The set of regions is computed as follows: We always include a set of pre-defined default regions (the lungs, the heart, ...) and include all regions explicitly mentioned in observations, as well as their parent regions. Additionally, we randomly sample regions. Their sampling probabilities are computed based on how often they are associated with positive vs. negative findings, i.e. the more often a region is associated with positive findings and the less often it is associated with negative findings, the more often we sample it as a question. This assures that we generate additional negative questions for regions that are only/mostly mentioned with positive findings.

Table 12: Region templates.

Template (ID)	Question Example	Variables	Main answer	Details	Related Inf.
describe_region C01_describe_region	Describe the left lung.	region	finding (positive) device (positive) finding (negative) device (negative)	-	related regions
describe_abnormal_region C02_describe_abnormal_region	Describe all abnormal findings in the lung bases.	region	finding (positive)	-	device (positive) related regions
is_abnormal_region C03_is_abnormal_region	Are there any abnormal findings in the mediastinum?	region	template finding (positive)	device (positive) finding (negative)	related regions
is_normal_region C04_is_normal_region	Is the heart normal?	region	template finding (positive)	region (positive)	finding (negative) related regions
describe_region_device C07_describe_region_device	Check the right chest for implants.	region subcategory	device (positive) device (negative)	-	related regions
has_region_device C08_has_region_device	Are there any tubes, lines, or ports in or near the left lung?	region subcategory	template device (positive)	device (negative)	device (negative) related regions

Finding In this strategy, we generate question about specific findings (radiological findings, diseaes, devices, ...). This includes descriptions of findings, yes/no questions about the presence of findings, location of findings, and severity of findings. We use the templates defined in Tab. 13.

For a given finding/device entity, the strategy first identifies observations associated with it and classifies them into the five types defined in the study abnormality strategy. Additionally, it identifies observations that contain related finding/device entities. This includes parent findings (i.e. findings that are parents of the current one), same subcat findings (i.e. findings having the same sub-category), correlated findings (based on statistics computed over the whole scene graph dataset), indications of the current finding, and findings that are indicative of the current finding. The observation subset can be selected based on the template configuration. Additionally, a template answer can be included, which is selected based on whether the finding is present or not and based on severity levels. This template may also be filled with information about the localization of the finding.

Before generating QA-pairs, the strategy first identifies a set of finding/device entities. For each of these entities an individual QA-pair is generated. The set of entities is computed as follows: We always include a set of pre-defined default entities and include all entities explicitly mentioned in observations, as well as their parent entities. Additionally, we randomly sample entities. Their sampling probabilities are computed based on how often they are mentioned positively vs. negatively (over all scene graphs), i.e. the more often a finding is mentioned positively and the less often it is mentioned negatively, the more often we sample it as a question. This assures that we generate additional negative questions for findings that are only/mostly mentioned positively.

Table 13: Finding templates.

Template (ID)	Question Example	Variables	Main answer	Details	Related Inf.
describe_finding D01_describe_finding	Describe the pleural effusion.	finding	finding (positive) finding (negative)	-	parent findings indications indicative of same subcat correlated
has_finding D02_has_finding	Is there any indication of pneumonia?	finding	template	finding (positive) finding (negative)	parent findings indications same subcat correlated
where_is_finding D03_where_is_finding	Where is the lung nodule located?	finding	template	finding (positive) finding (negative)	parent findings indications
how_severe_is_finding D04_how_severe_is_finding	How severe is the car- diomegaly?	finding	template	finding (positive) finding (negative)	parent findings indications
describe_device D05_describe_device	Describe the endotracheal tube.	device	device (positive) device (negative)	-	parent findings same subcat
has_device D06_has_device	Is a pacemaker visible in the study?	device	template	device (positive) device (negative)	same subcat
where_is_device D07_where_is_device	Where are the surgical clips located?	device	template	device (positive) device (negative)	same subcat

E.3 QUALITY ASSESSMENT

Scene Graph Quality The scene graph quality aspects are computed by simply inspecting the observations nodes and checking which fields are set or empty. Additionally, we track issues during the graph construction procedure and derive quality aspects from them.

QA Quality We automatically assess the quality of the textual content of QA-pairs using Llama 3.1 8B (Grattafiori et al., 2024) as a judge for the five criteria presented in Tab. 10.

For rating *entailment* (Listing 7), we condition the model on the report, the question, as well as the answer parts and we rate each answer part individually.

Listing 7: LLM prompt used for entailment evaluation of generated answers.

```
You will be given a Report (medical report of a chest X-ray study), a
    \hookrightarrow Question (about the study), and an Answer (to the question)
   \hookrightarrow consiting of several (numbered) sentences.
Your task is to assess/rate whether each of the answer sentences is true
    \hookrightarrow or not, given a the reference report about the chest X-ray. This
   \hookrightarrow task is known as entailment verification.
Assess the quality of each answer sentence independently and use one of
    \hookrightarrow the rating options provided below to assess how well the facts in
   \hookrightarrow each sentences align with the report.
# Guidelines:
 Rate each sentence in the Answer individually; do NOT use any prior
    \hookrightarrow answer sentences as context or source
 Provide the rating for each answer sentences in its own line starting
   \hookrightarrow with the sentence number followed by your rating
- For each sentence, use ONE of the rating options provided below, do NOT
   \hookrightarrow use any other options
- An example format will be provided
- DO NOT REPEAT the question or answer sentences in your response!
## Rating Options -- ONLY USE ONE OF THE FOLLOWING OPTIONS
```

Example: The same finding is described in the answer and the report

- ALIGNED_MENTIONED: Answer aligns with the report (is factually correct)

 \hookrightarrow and all facts are explicitly stated in the report.

```
1998
       - ALIGNED_INFERABLE: Answer aligns with the report (is factually correct)
1999
           \hookrightarrow but some facts are NOT explicitly stated in the report, can
2000
           \hookrightarrow however be derived from what is written there.
2001
           Example: The answer provides a more general description of what is
                \hookrightarrow written in the report.
2002
       - ALIGNED_NEGATIVE_NOT_MENTIONED: Answer does NOT contradict the report (
2003
           \hookrightarrow may factually correct) but some facts (negative findings) cannot be
2004
           \hookrightarrow derived from the report, are however likely correct because they
           \hookrightarrow are negative findings and nothing contradictory is mentioned in the
2006
           \hookrightarrow report.
           Example: The answer mentions that a finding is not present but this
2007
                \hookrightarrow is not explicitly mentioned and does not contradict anything in
2008
                \hookrightarrow the report.
2009
       - ALIGNED_GENERAL_STATEMENT: Answer is a more general statement or
2010
           \hookrightarrow summary that is not explicitly mentioned but aligns roughly with
           \hookrightarrow the overall report.
2011
           Example: Summaries of whether the study is positive or negative.
2012
       - NON_ALIGNED_NON_INFERABLE: Answer does NOT contradict the report but
2013
           \hookrightarrow the correctneaa of some facts cannot be validated using the report.
2014
           Example: The answer mentions that a finding is present but this is
2015
               \hookrightarrow never mentioned in the report and cannot be concluded from it.
2016
       - NON_ALIGNED_MISLEADING: Answer does NOT directly contradict the report
           \hookrightarrow but the description is highly misleading considering the report.
2017
           Example: The answer mentions that a finding is not present, which is
2018
                \hookrightarrow never mentioned in the report but could likely be present
2019
                \hookrightarrow considering the report.
2020
       - NON_ALIGNED_CONTRADICTING: Answer contradict with the report.
2021
           Example: The answer describes that a finding is not present, which is
                \hookrightarrow however mentioned as present in the report or vice versa.
2022
2023
       # Example Format:
2024
       Report:
2025
       --- START OF REPORT ---
2026
       --- END OF REPORT ---
2027
2028
       Question: ...
2029
2030
       Answer (2 sentences to rate):
2031
       [01] First answer sentence.
       [02] Second answer (last sentence in this example).
2032
2033
       Rating (provide 2 ratings):
2034
       [01] ALIGNED_MENTIONED
2035
       [02] NON_ALIGNED_NON_INFERABLE
2036
2037
       # Proceed with the following Report, Question, and Answer sentences:
2038
       Report:
2039
       --- START OF REPORT ---
2040
       <REPORT>
       --- END OF REPORT ---
2041
2042
       Question: <QUEST>
2043
2044
       Answer (NUMANS) sentences to rate):
2045
       <ANSWERS>
2046
       Rating (provide < NUMANS > ratings):
2047
```

For rating *relevance* (Listing 8), we condition the model on the question as well as the answer parts (but not on the report) and we rate each answer part individually.

20482049

2050

2051

Listing 8: LLM prompt used for relevance evaluation of generated answers.

```
2052
2053
       You will be given a question (about a chest X-ray study), and an answer (
           \hookrightarrow to the question) consiting of several (numbered) sentences.
2054
2055
       Your task is to assess/rate whether each of the answer sentences relevant
2056
          \hookrightarrow to answer the question or is redundant.
2057
       Assess the quality of each answer sentence and use one of the rating
2058
           \hookrightarrow options provided below.
2059
       # Guidelines:
2060
       - Rate each sentence in the answer individually; but check for redundancy
2061
           \hookrightarrow with previous sentences.
2062
       - Provide the rating for each answer sentences with the sentence number
2063
           \hookrightarrow followed by your rating.
       - For each sentence, use ONE of the rating options provided below, do NOT
2064
          \hookrightarrow use any other options.
2065
       - An example format will be provided.
2066
       - DO NOT REPEAT the question or answer sentences in your response!
2067
2068
       ## Rating Options -- ONLY USE ONE OF THE FOLLOWING OPTIONS
       - RELEVANT_MAIN_ANSWER: Fullfill ALL of the following
2069
           a) Are relevant to the question
2070
           b) Are needed to answer the question or provide details
2071
           c) Are not redundant to previous RELEVANT_MAIN_ANSWER sentences
2072
       - RELATED_INFO: Fullfill ALL of the following
2073
           a) Are NOT needed to answer the question
           b) Provide additional context related to the question or other answer
2074
               \rightarrow sentences
2075
           c) Are not redundant to any previous sentences
2076
       - REDUNDANT_INFO(...): Fullfill ALL of the following
2077
           a) Would fullfill criteria a-b) for RELEVANT_MAIN_ANSWER, or
2078
               \hookrightarrow RELATED_INFO
           b) Contains exactly the same information that was already provided in
2079
               \hookrightarrow a previous sentence of the same type (ONLY consider previous
2080
               2081
           c) Does not provide any additional details or related information
2082
           d) Could be removed without changing the content of the answer
2083
           Note: replace ... with the sentences IDs OF PREVIOUS SENTENCE with

→ which the current sentence is redundant

2084
       - IRRELEVANT_INFO: Fullfill ALL of the following
2085
           a) Does not classify as any of the above
2086
           b) No information in the sentence is relevant or related to the
2087
               \hookrightarrow question
2088
       # Example Format:
2089
       Question: ...
2090
2091
       Answer (4 sentences to rate):
2092
       [01] First answer sentence.
2093
       [02] Second answer.
       [03] Third answer, containint no additional information, everything was
2094
           \hookrightarrow already mentioned in 01 and 02.
2095
       [04] Fourth sentence.
2096
2097
       Rating (provide 4 ratings):
2098
       [01] RELEVANT MAIN ANSWER
       [02] IRRELEVANT_INFO
2099
       [03] REDUNDANT_INFO(01,02)
2100
       [04] RELATED_INFO
2101
2102
2103
       # Proceed with the following Question and Answer sentences:
       Question: <QUEST>
2104
2105
      Answer (NUMANS) sentences to rate):
```

2153

2154

2155

21562157

2158 2159 Report:

<REPORT>

<ANSWERS>

--- START OF REPORT ---

--- END OF REPORT ---

Question: <QUEST>

Answer (to rate):

```
2106
       <ANSWERS>
2107
2108
       Rating (provide < NUMANS > ratings):
2109
2110
       For rating completeness (Listing 9), we condition the model on the report, the question, as well as the
2111
       full answer and we rate the full answer as a whole.
2112
               Listing 9: LLM prompt used for completeness evaluation of generated answers.
2113
2114
       You will be given a Report (medical report of a chest X-ray study), a
           \hookrightarrow Question (about the study), and an Answer (to the question) .
2115
2116
       Your task is to assess/rate whether the provided Answer contains all the
2117
            \hookrightarrow necessary information to answer the Question, considering the
2118
           \hookrightarrow Report as the source of truth.
2119
       # Guidelines:
2120
       - Do not assess whether the answer is correct but whether it is contains
2121
           \hookrightarrow all relevant information from the Report to answer the Question.
2122
       - Use ONE of the rating options provided below, do NOT use any other
2123
           \hookrightarrow options.
2124
       - Answer with a short explanation (a few words) followed by "->" and the
           \hookrightarrow rating option.
2125
       - An example format will be provided.
2126
       - DO NOT REPEAT the report, question, or answer sentences in your
2127
           \hookrightarrow response!
2128
       ## Rating Options -- ONLY USE ONE OF THE FOLLOWING OPTIONS
2129
       - FULLY_COMPLETE: All facts from the report that are relevant to the
2130
           \hookrightarrow question are included and the question is answered.
2131
       - DETAILS_MISSING: The main facts from the report that are relevant to
2132
           \hookrightarrow the question are included BUT some details are missing.
2133
       - NOT_ANSWERED: While facts from the report may be contained, the answer
2134
           \hookrightarrow does not relate to the question at all.
       - INCOMPLETE_NON_MISLEADING: Main facts are missing, but these should not
2135
           \hookrightarrow lead to a misrepresentation of the facts (e.g. only some negative
2136
           \hookrightarrow findings are not mentioned).
2137
       - INCOMPLETE_MISLEADING: Important facts are missing, such that the
2138
           \hookrightarrow answer may mislead the reader.
2139
       # Example Format:
2140
       Report:
2141
       --- START OF REPORT ---
2142
       . . .
2143
       --- END OF REPORT ---
2144
       Question: ...
2145
2146
       Answer (to rate):
2147
2148
2149
       Rating (your task):
       severity is missing -> DETAILS_MISSING
2150
```

Proceed with the following Report, Question, and Answer:

```
2160
2161
       Rating (your task):
2162
2163
       For rating question clarity (Listing 10), we condition the model on the question only and rate it.
2164
2165
             Listing 10: LLM prompt used for question clarity evaluation of generated questions.
2166
       You will be given a medical Question about a radiological chest X-ray
2167
           \hookrightarrow study (which is not provided).
2168
       Your task is to assess/rate the clarity of the Question, i.e. whether its
2169
           \hookrightarrow wording is clear and unambiguous, and whether it is easy to
2170
           \hookrightarrow understand and answer.
2171
       # Guidelines:
2172
       - Use ONE of the rating options provided below, do NOT use any other
2173
           → options
2174
       - Answer with a short explanation (a few words) followed by "->" and the
2175
           \hookrightarrow rating option
       - An example format will be provided
2176
       - DO NOT REPEAT any part of the question or answer sentences in your
2177
           \hookrightarrow response!
2178
2179
       ## Rating Options -- ONLY USE ONE OF THE FOLLOWING OPTIONS
2180
       - OPTIMAL: The question is mostly clear, unambiguous, and can be answered
            \hookrightarrow . It is well-structured and concise without grammatical errors.
2181
       - UNUSUAL_SENTENCE_STRUCTURE: The question is mostly clear, unambiquous,
2182
            \hookrightarrow and can be answered. However, the sentence structure is unusual or
2183
           \hookrightarrow complex. There are no grammatical errors.
2184
       - GRAMMATICAL_ERRORS: The question is mostly clear, unambiguous, and can
2185
           \hookrightarrow be answered. However, there are grammatical errors that may affect
           \hookrightarrow the clarity. The sentence may or may not be well-structured.
2186
       - UNRELATED_TO_CHEST_XRAY: The question is mostly clear and unambiguous.
2187
           \hookrightarrow However, it does not make sense to ask this question about a chest
2188
           \hookrightarrow X-ray study, because it does not relate to the content that can be
2189
           \hookrightarrow observed in a chest X-ray. There may or may not be grammatical
2190
           \hookrightarrow errors or unusual sentence structure.
       - UNCLEAR_QUESTION: The question may be misunderstood, is ambiguous, or
2191
           \hookrightarrow otherwise unclear. Any answer could be misleading or incorrect,
2192
           \hookrightarrow even with proper medical knowledge and context. There may or may
2193
           \hookrightarrow not be grammatical errors or unusual sentence structure.
2194
2195
       Note that simply stating the indication/history motivating the study is
            \hookrightarrow considered a valid question (and should not be rated as
2196
           → UNCLEAR_QUESTION solely for not being an explicit question)!
2197
2198
       # Example Format:
2199
       Ouestion (to rate): ...
2200
       Rating (your task):
2201
       The question is unrelated to chest X-rays -> UNRELATED_TO_CHEST_XRAY
2202
2203
2204
       # Proceed with the following Question:
2205
       Question (to rate): < QUEST>
2206
       Rating (your task):
2207
2208
```

For rating *answer clarity* (Listing 11), we condition the answer parts only (but not on the report or question) and we rate each answer part individually.

2209

22102211

2212 2213

Listing 11: LLM prompt used for answer clarity evaluation of generated answers.

```
You will be given a medical Answer to an unknown question about a \hookrightarrow radiological chest X-ray study (which is not provided).
```

```
2214
       Your task is to assess/rate the clarity of each sentence of the Answer, i
2215
           \hookrightarrow .e. whether its wording is clear and unambiguous, and whether it is
2216
           \hookrightarrow easy to understand.
2217
       # Guidelines:
2218
       - Rate each sentence in the Answer individually; do NOT use any prior
2219
           \hookrightarrow answer sentences as context or source
2220
       - Provide the rating for each answer sentences in its own line starting
2221
            \hookrightarrow with the sentence number followed by your rating
2222
       - For each sentence, use ONE of the rating options provided below, do NOT
           \hookrightarrow use any other options
2223
       - An example format will be provided
2224
       - DO NOT REPEAT the question or answer sentences in your response!
2225
2226
       ## Rating Options -- ONLY USE ONE OF THE FOLLOWING OPTIONS
       - OPTIMAL: The answer sentence is mostly clear and unambiguous. It is
2227
           \hookrightarrow well-structured and concise without grammatical errors.
2228
       - UNUSUAL_SENTENCE_STRUCTURE: The answer sentence is mostly clear and
2229
           \hookrightarrow unambiguous. However, the sentence structure is unusual or complex.
2230
           \hookrightarrow There are no grammatical errors.
2231
       - GRAMMATICAL ERRORS: The answer sentence is mostly clear and unambiguous
           \hookrightarrow . However, there are (severe) grammatical errors that affect the
2232
           \hookrightarrow clarity. The sentence may or may not be well-structured.
2233
       - UNCLEAR_ANSWER: The answer sentence may be misunderstood, is ambiguous,
2234
           \,\hookrightarrow\, or otherwise unclear. There may or may not be grammatical errors
2235
           \hookrightarrow or unusual sentence structure.
2236
       - NOT_UNDERSTANDABLE: The answer sentence cannot be understood at all. It
2237
           \hookrightarrow is completely unclear, nonsensical, gibberish, or contradictory in
           \,\hookrightarrow\, itself. There may or may not be grammatical errors or unusual
2238
           \hookrightarrow sentence structure.
2239
2240
       # Example Format:
2241
       Answer (4 sentences to rate):
       [01] This first sentence.
2242
       [02]
            This is the second answer sentence.
2243
       [03] Some text where it is unclear what is meant.
2244
       [04] This is the last answer sentence.
2245
2246
       Rating (provide 4 ratings):
       [01] GRAMMATICAL_ERRORS
2247
       [02] OPTIMAL
2248
       [03] UNCLEAR_ANSWER
2249
       [04] OPTIMAL
2250
2251
       # Proceed with the following Answer sentences:
2252
       Answer (NUMANS) sentences to rate):
2253
       <ANSWERS>
2254
2255
       Rating (provide < NUMANS > ratings):
2256
```

E.4 RESOURCES FOR DATASET CONSTRUCTION AND EVALUATION

E.4.1 SOURCE DATASETS

2257

2258 2259

2260 2261

2262

2263

2264

2265 2266

2267

MIMIC-CXR (Johnson et al., 2024b; 2019a; Goldberger et al., 2000) We use the MIMIC-CXR dataset version 2.1.0 (https://physionet.org/content/mimic-cxr/2.1.0/ as the source of radiology reports from which we extract the scene graphs. It contains 227 835 radiographic (chest X-ray) studies performed at the Beth Israel Deaconess Medical Center in Boston, MA, USA. It is licensed under the PhysioNet Credentialed Health Data License 1.5.0.

MIMIC-CXR-JPG (Johnson et al., 2024a; 2019b; Goldberger et al., 2000) We use the MIMIC-CXR-JPG dataset version 2.1.0 (https://physionet.org/content/mimic-cxr-jpg/

2.1.0/) as the source of images for localization (the CXAS segmntation model is applied on these images). Additionally, we use the provided radiologist annotations (mimic-cxr-2.1.0-test-set-labeled.csv) as targets to evaluate the quality of extracted finding tags (Tabs. 2a and 5). The dataset is derived from MIMIC-CXR and is licensed under the PhysioNet Credentialed Health Data License 1.5.0.

Chest ImaGenome (Wu et al., 2021a;b; Goldberger et al., 2000) We use the Chest ImaGenome Dataset version 1.0.0 (https://physionet.org/content/chest-imagenome/1.0.0/) as a source of anatomical region bounding boxes for localization. Additionally, we use their provided scene graphs as a baseline for the evaluations of our scene graphs (Tabs. 2 and 5 to 8). It contains scene graphs for 242 072 frontal images from MIMIC-CXR that have been created using rule-based natural language processing and CXR atlas-based bounding box detection. The dataset is derived from MIMIC-CXR and is licensed under the PhysioNet Credentialed Health Data License 1.5.0.

CXR-LT 2024 (Holste et al., 2025; 2024; Goldberger et al., 2000) We use the CXR-LT 2024 dataset version 2.0.0 (https://physionet.org/content/cxr-lt-iccv-workshop-cvamd/2.0.0/) as targets to evaluate the quality of extracted finding tags (Tabs. 2a and 6). More precisely, we use the gold standard dataset provided for Task 2 in the CXR-LT 2024 challenge tasks (406 reports, 26 classes). The dataset is derived from a small subset of MIMIC-CXR and was hand-labeled by radiologists. It is licensed under the PhysioNet Credentialed Health Data License 1.5.0.

MS-CXR (Boecking et al., 2024; 2022; Goldberger et al., 2000) We use the MS-CXR dataset version 1.1.0 (https://physionet.org/content/ms-cxr/1.1.0/) as targets to evaluate the quality of extracted finding boxes (Tabs. 2b and 7). The dataset contains 1162 image-sentence pairs of bounding boxes and corresponding phrases (and their finding classes) for 8 different findings. It is derived from a small subset of MIMIC-CXR and was hand-labeled by radiologists. It is licensed under the PhysioNet Credentialed Health Data License 1.5.0.

REFLACX (Bigolin Lanfredi et al., 2021; 2022; Goldberger et al., 2000) We use the REFLACX dataset version 1.0.0 (https://physionet.org/content/reflacx-xray-localization/1.0.0/) as targets to evaluate the quality of extracted finding boxes (Tabs. 2b and 8). The dataset provides eye-tracking data collected for 3032 frontal chest x-rays from the MIMIC-CXR dataset. Additionally, it provides hand-labeled ellipses localizing for several anomalies present in the images. We only use the ellipses but do not use the eye-tracking data. It is licensed under the PhysioNet Credentialed Health Data License 1.5.0.

E.4.2 Models

LLM Selection We choose the information extraction model (Llama 3.1 70B) based on the following criteria:

- Open weights and locally hostable (required due to the MIMIC dataset licenses).
- Fast and efficient to run on large amounts of data.
- Sufficient extraction quality and instruction following. Sophisticated reasoning capabilities are not required.
- No vision component is required.

While using other models would be possible, we choose Llama 3.1 70B based on preliminary experiments (at the start of dataset construction). Those experiments indicated that it provides good trade-offs for the points above.

The study of alternative models may be interesting. However, we do not expect major improvements in data quality, as our semantic entity mapping and our automatic quality assessment already provide a robust framework to assure data quality even with small data extraction errors. Therefore, we leave this to future work.

Llama 3.1 70B (Grattafiori et al., 2024) We use the AWQ-INT4 (Lin et al., 2024) quantized version of Llama 3.1 70B Instruct provided by the Huggingface hub at https://huggingface.co/hugging-quants/Meta-Llama-3.1-70B-Instruct-AWQ-INT4. The model is derived from the https://huggingface.co/meta-llama/Llama-3.1-70B-Instruct and is licensed under the LLAMA 3.1 COMMUNITY LICENSE AGREEMENT. We limit the maximum number of tokens to 6144.

Llama 3.1 8B (Grattafiori et al., 2024) We use the AWQ-INT4 (Lin et al., 2024) quantized version of Llama 3.1 70B provided by the Huggingface hub at https://huggingface.co/hugging-quants/Meta-Llama-3.1-8B-Instruct-AWQ-INT4. The model is derived from the https://huggingface.co/meta-llama/Llama-3.1-8B-Instruct and is licensed under the LLAMA 3.1 COMMUNITY LICENSE AGREEMENT. We limit the maximum number of tokens to 8192.

CXAS (Seibold et al., 2022; 2023) We use the model provided by the CXAS Python library https://pypi.org/project/cxas/. See also https://github.com/ConstantinSeibold/ChestXRayAnatomySegmentation. It is licensed under the Attribution-NonCommercial-ShareAlike 4.0 International license. We run segmentation of all anatomical structures on half the original image resolution (half original image width and height).

BioLORD (Remy et al., 2024) We use the BioLORD-2023-C variant provided by the Huggingface model hub at https://huggingface.co/FremyCompany/BioLORD-2023-C and licensed under the MIT license. To apply the model, we use the Sentence Transformers library (https://github.com/UKPLab/sentence-transformers), which is licensed under the Apache-2.0 license.

Model Inference Details For all LLM-based information extraction steps, we rely on the vLLM library (Kwon et al., 2023) (https://github.com/vllm-project/vllm, Apache-2.0 license) for inference. We run all models with temperature = 0.0. All json-outputs are parsed using the Pydantic libary (https://docs.pydantic.dev).

E.4.3 COMPUTATIONAL COSTS

Each dataset construction step can run on an individual Nvidia A100 GPU, but we use multiple GPUs in parallel, with each GPU responsible for a different subset of the dataset. Semantic segmentation of all 158 anatomical structures using the CXAS models takes about 6 seconds per image, leading to a total of about 628 GPU hours. Sentence extraction takes about 1 second per study (report), leading to a total of about 65 GPU hours (for 227 835 studies). Observation extraction takes about 1.7 seconds per study, leading to a total of about 108 GPU hours. Indication extraction takes about 0.3 seconds per study, leading to a total of about 19 GPU hours. Scene graph construction (including entity matching) takes about 0.6 seconds per study, leading to a tool of about 38 hours. Question-answer generation does not require a GPU but takes about 9 seconds per study (including all question templates and strategies), leading to a total of about 24 days. However, multiple processes can be run in parallel on a single machine, leading to an effective time of only about a day for all 42M QA-pairs. Quality assessment of QA texts again requires a GPU and consists of 5 individual steps that can be run in parallel. Overall the assessment takes about 6 GPU days for all 42M QA-pairs.

Societal Impact As a large vision-language dataset for medical imaging, this dataset has significant potential for societal impact. However, its use as a training source for models employed in clinical or medical applications also poses a substantial risk of misdiagnosis, highlighting the need for caution. Therefore, we strongly advise against relying solely on this dataset for fine-tuning or evaluating such models. On the other hand, this dataset can facilitate the development of large and interactive VQA models, which can provide supplemental information for patients, serve as a training tool for healthcare professionals, or optimize clinical workflows. The provided annotations, including bounding boxes and tags, further enhance its utility by providing a level of transparency and explainability in model predictions, allowing for more informed interpretation and analysis. By sparking research in this direction, this dataset can contribute to the advancement of the field and ultimately lead to positive long-term societal impacts. Nevertheless, it is essential to approach this

dataset with caution, recognizing its limitations and potential risks if used improperly. As such, we consider this dataset a valuable research asset, but not yet suitable as a (sole) training source for real-world medical applications, emphasizing the need for careful evaluation and validation.

F STRUCTURED VQA TASK

F.1 FURTHER STRUCTURED VQA RESULTS

Table 14: Further results of our structured VQA task from Sec. 5. We show all the metrics from Tab. 3 with additional sub-metrics of our RadStrucVQA metric, each with 95% confidence intervals (bootstrapping, n=1000). Besides our default VQA model and the MAIRA-2 baseline, we also show alternative settings of our VQA model, namely training without bounding boxes and/or tags and predicting bounding boxes and tags before or after the text. Apart from these adaptions, the experimental setup was the same as in Sec. 5. We found that none of these adaptions has major influences on the results (apart from being capable of predicting boxes/tags), indicating that text, boxes, and tags in our dataset do not contradict each other. However, we observed minor improvements in text quality by adding bounding boxes

	Model	Ours (ablations)			Ours (default)	MAIRA-2	
	Boxes	× / / /					
	Tags	Х	X	X	✓	1	/
	Text	only text	after text	before text	after text	before text	after text
RadFact*	Logical Prec	0.75	0.76	0.76	0.76	0.76	0.25
	Logical Rec	0.74	0.75	0.75	0.75	0.75	0.64
	Logical F1.	0.74	0.74	0.75	0.75	0.75	0.27
	Grounding Prec	-	0.88	0.87	0.88	0.87	0.69
	Grounding Rec	-	0.88	0.90	0.88	0.89	0.12
Ra	Grounding F1.	_	0.83	0.83	0.83	0.83	0.32
	Spatial Prec	_	0.68	0.67	0.68	0.67	0.12
	Spatial Rec	_	0.67	0.68	0.68	0.68	0.07
	Spatial F1.	-	0.63	0.63	0.64	0.63	0.06
	Finding Prec.	-	-	-	0.68	0.68	-
	Finding Rec.	_	_	_	0.67	0.66	_
	Finding F1	_	_	_	0.68	0.67	_
	Finding-pos Prec.	_	_	_	0.40	0.41	_
	Finding-pos Rec.	_	_	_	0.29	0.26	_
	Finding-pos F1	_	_	_	0.39	0.39	_
	Region Prec.	_	_	_	0.67	0.67	_
	Region Rec.	_	_	_	0.66	0.66	_
	Region F1	_	_	_	0.66	0.67	_
	Region-pos Prec.	_	_	_	0.29	0.34	_
(S)	Region-pos Rec.	_	_	_	0.21	0.21	_
Tag	Region-pos F1	_	_	_	0.29	0.32	_
RadStrucVQA (Tags)	Main-category Prec.	_	_	_	0.73	0.73	_
8	Main-category Rec.	_	_	_	0.70	0.70	_
ruc	Main-category F1	_	_	_	0.72	0.72	_
ıgŞı	Main-category-pos Prec.	_	_	_	0.49	0.52	_
ጼ	Main-category-pos Rec.	_	_	_	0.36	0.34	_
	Main-category-pos F1	_	_	-	0.47	0.49	_
	Sub-category Prec.	_	_	-	0.71	0.71	_
	Sub-category Rec.	_	_	_	0.67	0.67	_
	Sub-category F1	_	_	-	0.69	0.69	_
	Sub-category-pos Prec.	_	_	_	0.47	0.50	_
	Sub-category-pos Rec.	_	_	-	0.34	0.32	_
	Sub-category-pos F1	_	_	_	0.45	0.46	_
	Bbox-pos-entity Prec.	_	_	_	0.31	0.32	_
	Bbox-pos-entity Rec.	_	_	_	0.22	0.20	_
	Bbox-pos-entity F1	_	_	-	0.26	0.26	_

*Our RadStrucVQA implementation.

F.2 RADSTRUCVQA METRIC

Following the RadFact (Bannur et al., 2024) metric, we split the predictions into individual elements. In our case, we treat each answer part as its own element, ignoring the hierarchy level and order. For each QA-sample, this results in a set of prediction elements $\hat{\mathcal{Y}}$, where $|\hat{\mathcal{Y}}|$ is the number of answer parts in the predicted answer, and a set of target elements \mathcal{Y} , where $|\mathcal{Y}|$ is the number of answer parts in the target answer.

For each RadStrucVQ sub-metric (sub \in {logical, grounding, finding, . . . }), we compute a sample-level precision p_{sub} and recall r_{sub} score individually:

$$p_{\text{sub}}\left(\hat{\mathcal{Y}}, \mathcal{Y}\right) = s_{\text{sub}}\left(\hat{\mathcal{Y}}, \mathcal{Y}\right),$$
 (1)

$$r_{\text{sub}}(\hat{\mathcal{Y}}, \mathcal{Y}) = s_{\text{sub}}(\mathcal{Y}, \hat{\mathcal{Y}}),$$
 (2)

where $s_{\text{sub}}(\mathcal{H}, \mathcal{C}) \in [0, 1]$ is a sub-metric specific scoring function considering the hypothesis set \mathcal{H} given the context set \mathcal{C} . For precision $\mathcal{H} = \hat{\mathcal{Y}}$ is the prediction set and $\mathcal{C} = \mathcal{Y}$ is the target set, while for recall $\mathcal{H} = \mathcal{Y}$ and $\mathcal{C} = \hat{\mathcal{Y}}$.

The score s_{sub} is computed as the fraction of relevant hypothesis elements $h \in \mathcal{H}$ that are entailed, using a sub-metric specific entailment definition, given the context \mathcal{C} . More precisely:

$$s_{\text{sub}}(\mathcal{H}, \mathcal{C}) = \frac{\left| \left\{ h \in \mathcal{H} \mid \text{entailed}_{\text{sub}}(h, \mathcal{C}[h]) \land \text{relevant}_{\text{sub}}(h) \right\} \right|}{\left| \left\{ h \in \mathcal{H} \mid \text{relevant}_{\text{sub}}(h) \right\} \right|}, \tag{3}$$

where C[h] is the evidence from C for h defined as

$$C[h] = \left\{ c \in C \middle| h \text{ is logically entailed with } C \land c \text{ provides evidence for } h \right\}. \tag{4}$$

We compute C[h] by prompting an LLM to (i) identify entailment of h given all context elements in C, where h can be ENTAILED or NOT_ENTAILED (neutral or contradicting); and (ii) provide the relevant evidence for entailment, i.e. the context units $c \in C$ that support h. The LLM is given only the textual descriptions of each element (answer part), i.e. the entailment classification is purely logical and does not consider localization or any tags. Note that $C[h] = \{\}$ if h is not entailed.

Given the hypothesis h and its evidence C[h], the sub-metric entailment is computed individually by

entailed_{sub}
$$(h, C[h]) \in \{\text{true}, \text{false}\},$$
 (5)

while the relevant subset of hypothesis elements is identified using the sub-metric specific

$$relevant_{sub}(h) \in \{true, false\}.$$
 (6)

The definitions for each sub-metric can be found in Tab. 15.

Table 15: RadStrucVQA sub metric definitions. The logical, grounding, and spatial sub-metrics follow the same principles as the corresponding sub-metrics in RadFact (Bannur et al., 2024).

Sub-metric	$\mathrm{entailed}_{\mathrm{sub}}\big(h,\mathcal{C}[h]\big)$	$\mathrm{relevant_{sub}}(h)$	
logical	${\cal C}[h]$ is not empty i.e. there is positive evidence for h in ${\cal C}$ and h does not contradict ${\cal C}$	always true	
grounding	$\begin{split} & \text{entailed}_{\text{logical}}\big(h, \mathcal{C}[h]\big) \wedge \text{IoH}(h, \mathcal{C}[h]) \geq 0.5 \\ & \text{where IoH is the Intersection between boxes in h and boxes in $\mathcal{C}[h]$} \\ & \text{over the total box area in h}, \\ & \text{with intersection/area computed based on box-masks (unions of boxes)} \end{split}$	h has bounding boxes \wedge entailed $_{ ext{logical}}ig(h,\mathcal{C}[h]ig)$	
spatial	$\mathrm{entailed}_{\mathtt{grounding}}\big(h,\mathcal{C}[h]\big)$	h has bounding boxes	
finding	$ \begin{array}{c} \operatorname{entailed_{logical}}\left(h,\mathcal{C}[h]\right) \wedge \\ \operatorname{each} \operatorname{of} \operatorname{the} \operatorname{finding} \operatorname{tags} \operatorname{in} h \operatorname{is} \operatorname{present} \operatorname{in} \operatorname{any} \operatorname{of} \mathcal{C}[h], \\ \operatorname{only} \operatorname{considering} \operatorname{the} \operatorname{subset} \operatorname{of} \mathcal{C}[h] \operatorname{ with the same positivity} \end{array} $	h has finding tags	
finding-pos	$\mathrm{entailed}_{\mathrm{finding}}\big(h,\mathcal{C}[h]\big)$	$ ext{relevant}_{ ext{finding}}(h) \land h ext{ is positive}$	
region	$ \begin{array}{c} \operatorname{entailed_{logical}}\left(h,\mathcal{C}[h]\right) \wedge \\ \operatorname{each} \ \text{of the region tags in} \ h \ \text{is present in any of} \ \mathcal{C}[h], \\ \operatorname{only considering the subset of} \ \mathcal{C}[h] \ \text{with the same positivity} \end{array} $	h has region tags	
region-pos	$\mathrm{entailed}_{\mathrm{region}}\big(h,\mathcal{C}[h]\big)$	$ ext{relevant}_{ ext{region}}(h) \land h ext{ is positive}$	
main-category	$ \begin{array}{c} \operatorname{entailed_{logical}}\left(h,\mathcal{C}[h]\right) \wedge \\ \operatorname{each} \text{ of the finding main category tags in } h \text{ is present in any of } \mathcal{C}[h] \\ \operatorname{only considering the subset of } \mathcal{C}[h] \text{ with the same positivity} \end{array} $	h has finding main category tags	
main-category-pos	$\mathrm{entailed_{main-category}} \left(h, \mathcal{C}[h] \right)$	h is positive relevant _{main-category} $h \land h$	
sub-category	$ \begin{array}{c} \operatorname{entailed_{logical}}\left(h,\mathcal{C}[h]\right) \wedge \\ \operatorname{each} \operatorname{of} \operatorname{the finding sub category tags in } h \operatorname{is present in any of } \mathcal{C}[h], \\ \operatorname{only considering the subset of } \mathcal{C}[h] \operatorname{ with the same positivity} \end{array} $	h has finding sub category tags	
sub-category-pos	$\operatorname{entailed_{sub-category}} ig(h, \mathcal{C}[h]ig)$	$_{\text{relevant}_{\text{sub-category}}(h) \land h \text{ is positive}}$	
bbox-pos-entity	$\mathrm{entailed}_{\mathrm{finding}}\big(h,\mathcal{C}[h]\big) \wedge \mathrm{entailed}_{\mathrm{grounding}}\big(h,\mathcal{C}[h]\big)$		

Implementation Details The final precision/recall scores are computed by averaging the sample-level scores. F1 scores can also be computed by first taking the per-sample harmonic mean of precision and recall before averaging the sample-level F1 scores. Invalid answers, samples with LLM parse errors during entailment computation, as well as samples without relevant hypotheses are ignored during averaging. We use the same entailment prompts and few-shot examples as in RadFact (Bannur et al., 2024) but use the Llama 3.1 8B (Grattafiori et al., 2024) model, allowing us to compute the metric locally.

F.3 EXPERIMENTAL SETUP

Vision-Language Model Training Our vision-language model follows the Llava architecture (Liu et al., 2023), using Rad-DINO (Pérez-García et al., 2025) (microsoft/rad-dino) for image encoding and the 3B Llama 3.2 language model (https://huggingface.co/meta-llama/Llama-3.2-3B-Instruct) connected via an MLP projection layer. We freeze the image encoder and all existing language model parameters but add new special tokens (with trainable embeddings) and apply LoRA (Hu et al., 2021) to the language model. Therefore, we only train the projection layer, the LoRA parameters, and the newly added token embeddings (keeping the existing token embeddings frozen). We train for one epoch on 1M samples from our CXR-QBA fine-tuning grade dataset (MIMIC-CXR's train split), where we use autoregressive training but only apply the loss to answer tokens. For image encoding and projection, we adopt the hyperparameters of MAIRA-2 (Bannur et al., 2024): We square-crop the images and resize them to 518×518 , leading to $37 \times 37 = 1369$ image patches (i.e. image tokens), then we use the features of the last image encoder layer, and project the image tokens using 4 projection layers with GeLU activations. For LoRA, we use r = 64, $\alpha = 16$, and dropout 0.05. The maximum number of tokens for the language model is restricted to 2048. We use the AdamW optimizer with cosine annealing scheduling with 500 warmup

steps, maximum learning rate 1e-3, no weight decay, a batch size of 4 with 16 accumulation steps, gradient norm clipping at 1.0, and bf16 precision. The model is evaluated on the test split (following MIMIC-CXR) of our CXR-QBA fine-tuning grade set.

Prompt and Special Tokens Our question prompt follows the template shown in Listing 12, where
<boi> (begin of image), <eoi> (end of image), and <imgref1> (first image reference) are newly added special tokens, tokens are replaced by image token features, and {QUESTION} is replaced by the specific question.

Listing 12: Question prompt.

The answers are formatted into sequences using XML-style structures and special tokens to represent tags and bounding boxes. An example is given in Listing 13.

Listing 13: Answer prompt.

We use special start and end tokens for answer parts (<answer> / </answer>), bounding boxes (<box>/</box>), and groups of tags (<regions>/</regions>, <probability> /</probability>, <categories> / </categories>, <entities> / </entities>, <modifiers>/</modifiers>). For some tags we use individual special tokens, namely for laterality (e.g. <bilateral>), regions (e.g. <lungs>), certainty (e.g. <certain>), positivity (e.g. <neg>), and main categories (e.g. <ANATOMICAL_FINDING>). For others we use start/end tokens and normal text, namely for sub-categories (<subcat> / </subcat>) and finding entities (<entity>/</entity>). Bounding boxes are listed after all other tags, where we use <box> / </box> tokens and refer back to the image using <imgref1>. Inside the box-tokens we use special relative coordinate tokens (following MAIRA-2 (Bannur et al., 2024)) that represent the normalized (x_1, y_1, x_2, y_2) coordinates of the bounding box, each quantized to 100 different tokens per dimension. We use different tokens for the x- and y-dimensions but share them for both corners (e.g. x_1 and x_2 share the same token set). The textual description is the last part of each answer part and consists of plain text without special tokens. If an answer consists of multiple answer parts, then each answer part uses an individual block as in Listing 13. All new token embeddings are initialized close to the existing token embeddings, where we try to initialize them based on keywords defined for each token. More precisely, given a set of keywords for a new token, we tokenize the keywords using the old vocabulary and compute the average embedding of all these tokens. This is then used as the initialization for the new token.

MAIRA-2 Baseline We use the MIARA-2 (Bannur et al., 2024) checkpoint available at https://huggingface.co/microsoft/maira-2. We freeze the full model but modify the prompt. More precisely, we use their original prompt for grounded report generation but slightly modify it, asking the model to answer to the question (included in the modified prompt) instead of reporting all findings in the image. The rest of the prompt is kept unchanged. This model is then evaluated on the same test set as our vision-language model. It is capable of generating individual answer parts, each with bounding boxes, but does not generate bounding boxes for negative answers and cannot generate any tags.