

Fact-Aware Multimodal Retrieval Augmentation for Accurate Medical Radiology Report Generation

Anonymous ACL submission

Abstract

Multimodal foundation models hold significant potential for automating radiology report generation, thereby assisting clinicians in diagnosing cardiac diseases. However, generated reports often suffer from serious factual inaccuracy. In this paper, we introduce a fact-aware multimodal retrieval-augmented pipeline in generating accurate radiology reports (FactMM-RAG). We first leverage RadGraph to mine factual report pairs, then integrate factual knowledge to train a universal multimodal retriever. Given a radiology image, our retriever can identify high-quality reference reports to augment multimodal foundation models, thus enhancing the factual completeness and correctness of report generation. Experiments on two benchmark datasets show that our multimodal retriever outperforms state-of-the-art retrievers on both language generation and radiology-specific metrics, up to 6.5% and 2% score in F1CheXbert and F1RadGraph. Further analysis indicates that employing our factually-informed training strategy imposes an effective supervision signal, without relying on explicit diagnostic label guidance, and successfully propagates fact-aware capabilities from the multimodal retriever to the multimodal foundation model in radiology report generation.

1 Introduction

Within hospitals worldwide, chest radiology serves as a critical technique in identifying cardiac diseases and abnormalities. Results of a chest radiograph are typically consolidated in a radiology report, including the source X-ray and a radiologist-produced findings section detailing clinical observations. Manually generating these reports, however, can be both time-consuming and potentially inaccessible in under-resourced hospitals (Speets et al., 2006; Iyeke et al., 2022). Recent multimodal foundation models have exhibited remarkable capabilities in challenging

healthcare tasks, motivating an automation of this process to enhance physicians' efficiency on clinical decision-making and improve patient health outcomes (Çallı et al., 2021; Li et al., 2023; Moor et al., 2023; Tu et al., 2023; Sun et al., 2024).

Although prior medical multimodal foundation models have demonstrated promising capabilities on report generation given the radiology image, they still suffer from serious hallucinations by generating factually inaccurate reports (Pal et al., 2023; Ahmad et al., 2023; Pal and Sankarasubbu, 2024). Factual correctness is especially critical in chest radiology domains, as minute textual differences can drastically invert radiology report meaning and downstream prescribed treatments (Delbrouck et al., 2022; Xie et al., 2023; Liu et al., 2024). Retrieval-Augmented Generation (RAG) has emerged as a popular paradigm to address this issue by grounding text generation with retrieved relevant knowledge given a query (Lewis et al., 2021; Chen et al., 2022; Gao et al., 2024). However, developing medical multimodal retrievers remains challenging, requiring retrievers to bridge the gap between symptomatic image semantics and factually-equivalent report text.

To capture fine-grained details in chest radiographs and improve the factual completeness of generated reports, we introduce FactMM-RAG, a fact-aware multimodal retrieval-augmented pipeline for generating accurate radiology reports given a radiology image. By designing a novel report pair-mining procedure incorporating factual knowledge, we develop a fact-aware retriever to augment multimodal foundation models in generating accurate chest X-ray radiology reports. Specifically, we first leverage RadGraph (Jain et al., 2021) to mine factually-oriented report pairs by annotating consistent radiology entities and relations between query and reference reports with

certain abnormalities. Next, we train a universal multimodal encoding architecture through mined report pairs to conduct multimodal dense retrieval. Given an unseen patient’s radiology image, our retriever encodes it and searches for the most similar factually-informed reference report from an available report corpus. Passing them together into a multimodal foundation model unlocks its fact-aware potential to generate more accurate radiology reports.

Our experiments reveal that our retriever outperforms all state-of-the-art retrievers in both language generation and clinically relevant metrics on the MIMIC-CXR and CheXpert datasets, achieving up to 6.5% and 2% score in F1CheXbert and F1RadGraph for final RAG evaluation. We also investigate our retriever’s fact-aware capability controlled by factual similarity thresholds and confirm that our factually-informed training strategy can impose a useful supervision signal without relying on explicit diagnostic label guidance. Further analysis through retrieval evaluation metrics shows that the fact-aware capability of our retriever can be effectively propagated to the multimodal foundation models. Lastly, our case study highlights that among reports describing the same symptom from different retrievers, those generated by our model are more accurate and achieve greater factual correctness.

Our main contributions can be summarized as follows:

- We propose a fact-aware medical multimodal retriever to augment multimodal foundation models in generating accurate chest X-ray radiology reports.
- We design a method for mining factually-informed radiology report pairs that trains multimodal encoders to retrieve high-quality reference reports.
- We demonstrate that on two benchmark datasets, our medical multimodal retriever outperforms state-of-the-art medical multimodal retrievers on both language generation and clinically relevant metrics.

The rest of this paper is organized as follows. We review related work in Section 2. We discuss the pipeline of FactMM-RAG in Sections 3. Section 4 and 5 discuss our experimental setup and results.

2 Related Work

Retrieval Augmented Generation. Retrieval Augmented Generation, utilizing external knowledge to enhance language models, has shown great promise in text-generation performance on factual accuracy especially for Open-Domain QA. (Borgeaud et al., 2022; Izacard et al., 2022). Guu et al. (2020); Lewis et al. (2021) involve end-to-end training through both generators and retrievers; Shi et al. (2023); Yu et al. (2023b) adapt the end-to-end pattern by employing black-box LLM training signal propagation for retriever tuning. Further works have expanded RAG to multiple modalities, employing unified image-text encoders (Radford et al., 2021) or separate pretrained encoders (Dosovitskiy et al., 2021; Raffel et al., 2023) and plugging retrieved documents into multimodal foundation models (Chen et al., 2022; Hu et al., 2023). Yasunaga et al. (2023) similarly integrates multimodal retrieval with both text and image generation capabilities.

Medical Multimodal Retriever. Joint training of image-text pairs in a shared embedding space, as exemplified by CLIP (Radford et al., 2021), facilitates visual and textual modality interactions, providing flexible representations for general-domain downstream tasks. Adapting general-domain multimodal retrievers to medical domains, however, is non-trivial due to the necessity of specialized knowledge. Zhang et al. (2022) introduces an unsupervised approach for radiology image representation learning from paired text descriptions. Huang et al. (2021) leverages global image-report and local sub-region features for multimodal retrieval and classification. Wang et al. (2022); You et al. (2023) propose medical knowledge extraction for constructing contrastive learning image-text pairs. Zhang et al. (2024) addresses the limited diversity within medical datasets, curating a large biomedical image-text collection towards a biomedical multimodal foundation model. Nevertheless, these existing medical multimodal retrievers neglect specific image information and do not adequately emphasize factual accuracy, resulting in imprecision when retrieving radiology reports.

Medical Multimodal Foundation Model. Significant efforts have been made in applying multimodal foundation models to the medical

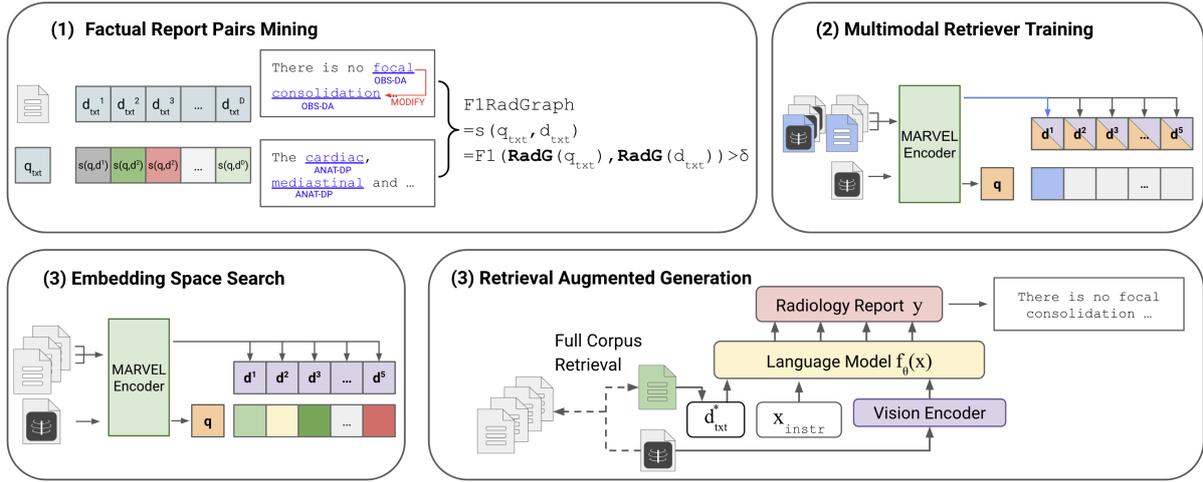


Figure 1: An overview of the FactMM-RAG system. It mainly contains three stages: (1) Leveraging RadGraph to characterize each radiology report and mine factually-informed report pairs; (2) Integrating factual knowledge into the training of the universal multimodal retriever; (3) Given the radiology image, employing the fact-aware multimodal retriever to search for factually-informed reference reports and augmenting the multimodal foundation model in generating accurate radiology reports.

imaging domain (Li et al., 2023; Moor et al., 2023; Tu et al., 2023; Sun et al., 2024). As chest X-ray radiology is the most commonly performed imaging examination, tailored medical multimodal foundation models for this critical area has gathered much attention (Chambon et al., 2022; Chen et al., 2021; Omkar Thawkar and Khan, 2023; Wu et al., 2023; Chen et al., 2024). Jain et al. (2021) advances this area by designing a novel information extraction schema to structure radiology reports from chest radiographs; Miura et al. (2021); Delbrouck et al. (2022) take a step forward, using reinforcement learning from semantic rewards to improve the factual quality of generated radiology reports; Chen et al. (2024) recently has also developed an instruction-tuned multimodal foundation model capable of sophisticated interpretation and analysis of chest X-rays.

One closely related line of work to ours is retrieval-based radiology report generation given only radiology images. For instance, Li et al. (2018) proposes a retrieval policy module to update radiology reports via hierarchical reinforcement learning; Endo et al. (2021) employs image-text embeddings from contrastive learning for retrieval-augmented radiology report generation; Ramesh et al. (2022) proposes synthesizing additional reports and reducing hallucinations from reference report priors to improve report generation.

3 Methodology

In this section, we present the overall methodology of FactMM-RAG. We first detail the training procedure of our fact-aware medical multimodal retriever in Section 3.1. We then provide the pipeline for retrieval-augmented radiology report generation with our multimodal retriever in section 3.2. The overview is illustrated in Figure 1.

3.1 Fact-aware Multimodal Retrieval

This section discusses the training process of the multimodal retriever with factual knowledge. Each patient in the corpus has a chest X-ray radiology image along with its corresponding report. We begin by annotating each report using RadGraph (Jain et al., 2021), then constructing factual report pairs to train our multimodal retriever. We describe these steps as follows.

Chest Radiograph Annotation. Since radiology reports are free-text, we utilize the RadGraph information extraction tool to extract structured knowledge graphs from them. Specifically, RadGraph employs named entity recognition and relation extraction models to identify radiological entities (e.g. carina, lungs, abnormalities) and the clinical relations between them (e.g. modify, located at, suggestive of). Each radiology report is then segmented into distinct regions and stored as $[(entity_1, entity label_1, relation_1), (entity_2, entity label_2, relation_2), \dots]$. After charac-

terizing the chest radiograph for each report in the training corpus, we construct factual report pairs.

Factual Report Pairs Mining. Each report has an associated medical label describing the symptom. We first utilize the query report to search for other reports with the same symptom, aiming to eliminate false negatives when constructing report pairs. Rather than solely relying on the diagnostic labels, we further capture the factually-oriented pathology semantics between different reports. Following F1RadGraph (Jain et al., 2021), we calculate the factual similarity $s(q_{txt}, d_{txt})$ between query report q_{txt} and other reports d_{txt} in the annotated format as follows,

$$s(q_{txt}, d_{txt}) = \frac{2 \cdot (\hat{q}_{txt} \cap \hat{d}_{txt})}{\text{length}(\hat{q}_{txt}) \cdot \text{length}(\hat{d}_{txt})}, \quad (1)$$

where $\hat{q}_{txt}, \hat{d}_{txt}$ denotes reports with only annotated entities and relations in RadGraph structured form. We then set a strict threshold δ to filter out searched reports with low similarity score:

$$N_{q_{txt}} = \{d_{txt} \in D | s(q_{txt}, d_{txt}) > \delta\}. \quad (2)$$

where $N_{q_{txt}}$ denotes factual positive report pairs for q_{txt} and D is the total training corpus. Since each query report is associated with a corresponding radiology image, these factual report pairs can also be applied to the query report’s radiology image. Next, we train our multimodal retriever with mined factual report pairs.

Multimodal Dense Retrieval. Following previous work (Zhou et al., 2024), we universally encode each query image q_{img} and other image-text pairs (d_{txt}, d_{img}) in the training corpus, using one encoder, MARVEL:

$$\mathbf{q} = \text{MARVEL}(q_{img}); \quad (3)$$

$$\mathbf{d} = \text{MARVEL}(d_{txt}, d_{img}), \quad (4)$$

where each image-text pair is represented as a single embedding. We then model the relevance score $f(q, d)$ between the query image and other image-text pairs by cosine similarity:

$$f(q, d) = \cos(\mathbf{q}, \mathbf{d}). \quad (5)$$

To inject factually-oriented medical knowledge into multimodal retrieval, we train the encoder to minimize the following loss,

$$\mathcal{L} = - \sum_{q_{img} \in D} \sum_{d^+ \in N_{q_{img}}} \log \frac{e^{f(\mathbf{q}, \mathbf{d}^+)/\tau}}{e^{f(\mathbf{q}, \mathbf{d}^+)/\tau} + \sum_{\mathbf{d}^-} e^{f(\mathbf{q}, \mathbf{d}^-)/\tau}}, \quad (6)$$

where d^+ are obtained through factual report pair mining and d^- are in-batch negative samples (Karpukhin et al., 2020). Then, we use our multimodal retriever and foundation model to perform retrieval-augmented radiology finding generation.

3.2 Retrieval Augmentation for Accurate Radiology Report Generation

Given our trained fact-aware multimodal retriever, we encode the query image and each report in the training corpus. Then, we retrieve the report with the highest relevance score to the query image as the factually-informed relevant report. Subsequently, we pass the query image along with the relevant report into a multimodal foundation model to perform retrieval-augmented generation training. The multimodal foundation model is finetuned by standard autogressive loss,

$$\mathcal{L} = -\frac{1}{n} \log \prod_i p_\theta(y_i | q_{img}, d_{txt}^*, x_{instr}, y_{<i}), \quad (7)$$

where q_{img} is the query image, d_{txt}^* is the retrieved factually-informed relevant report, x_{instr} is the prompt instruction, and y is the ground-truth report. During inference, we retrieve a relevant report from the training corpus using an unseen patient X-ray image, and pass them into the multimodal foundation model to generate findings with higher factual accuracy.

4 Experimental Setup

Dataset. Following Delbrouck et al. (2023), we use the processed MIMIC-CXR (Johnson et al., 2019) to train both retriever and foundation model. This dataset contains 125,417 training radiology image-report pairs, 991 validation pairs, and 1,624 test pairs. They are sourced from the Beth Israel Deaconess Medical Center. CheXpert (Irvin et al., 2019) is another chest X-ray dataset from Stanford Health Care. Since it contains complete finding reports only for a testing dataset containing 1000 pairs, we use it as zero-shot evaluation.

Evaluation Metrics. We evaluate our proposed system using both natural language generation and medically-tailored evaluation metrics.

For language fluency measures, we use ROUGE-L (Lin, 2004) to evaluate the longest common subsequence overlap between the generated and

reference findings, and BERTScore (Zhang et al., 2020a) to evaluate non-clinical semantic sentence similarity.

For clinical accuracy measures, we employ CheXbert (Smit et al., 2020) to generate the ground-truth diagnostic labels for finding reports, identifying 14 different types of observations. Following Delbrouck et al. (2023), we then calculate the F1CheXbert (Zhang et al., 2020b), which is the F1-score for 5 observations (Cardiomegaly, Edema, Consolidation, Atelectasis, Pleural Effusion) by comparing the generated report with the reference report’s classifications. Beyond using the limited diagnostic labels for evaluation, we also adopt F1RadGraph (Jain et al., 2021) to measure factual correctness by calculating the overlap in radiological entities and clinical relations between the generated report and the reference report.

Baselines. We mainly compare our retriever with other baselines under multimodal RAG setting. We include the following baselines, CLIP (Radford et al., 2021) is a multimodal retriever pretrained from general-domain image-text pairs; GLoRIA (Huang et al., 2021) leverages attention-weighted image regions with contextual words to learn localized and global representations for radiology images and reports; MedCLIP (Wang et al., 2022) and CXR-CLIP (You et al., 2023) build upon CLIP and utilize diagnostic labels as training signals for learning radiology image and text representations; BiomedCLIP (Zhang et al., 2024) extends the radiology-specific dataset and pretrains on a larger magnitude of biomedical data to learn multimodal representations; Med-MARVEL utilizes universal encoder MARVEL (Zhou et al., 2024) to conduct contrastive learning on each patient’s self image-report pair without further training on factual image-report pairs.

Implementation Details. In our experiments, we use MARVEL (Zhou et al., 2024) as our multimodal retriever backbone. MARVEL is a language model based on T5-ANCE (Yu et al., 2023a), trained with modality-balanced hard negatives. We use LLaVA (Liu et al., 2023) as our multimodal foundation model backbone. Since each radiology study contains multiple image views for each patient, we select the frontal view. We also concatenate the finding and impression sections to form the X-ray report. To reduce

training costs and address factual report pair imbalances, we rerank the retrieved reports by factual similarity and use the top 2 factual report pairs for each query to train our multimodal retriever. We leave the training details in Appendix A.

5 Evaluation Results

In this section, we present our experimental results. We first evaluate the overall performance between different retrievers under two settings in section 5.1. Next, we discuss the ablation studies in section 5.2. We then explore the fact-aware capability of our retriever in section 5.3 and section 5.4. Lastly, we show the superiority of our retriever through a case study in section 5.5.

5.1 Overall Performance

The results of our fact-aware RAG system are shown in Table 1. In MIMIC-CXR, FactMM-RAG outperforms state-of-the-art retrievers by a significant margin, up to 6.5% in F1CheXbert and 2% in F1RadGraph. In the CheXpert zero-shot evaluation, FactMM-RAG outperforms state-of-the-art retrievers by 2% and 1.2% in these two metrics, indicating our retriever’s generalization capability compared to other models.

Besides, we can observe that adopting the baseline retrievers on top of multimodal foundation models only yields marginal gains compared to the finetuning of foundation model generation without retrieval-augmentation. This shows that reports retrieved by baseline retrievers are factually-inferior to those from our retriever, potentially passing misleading information that prevents the foundation model from generating factual reports.

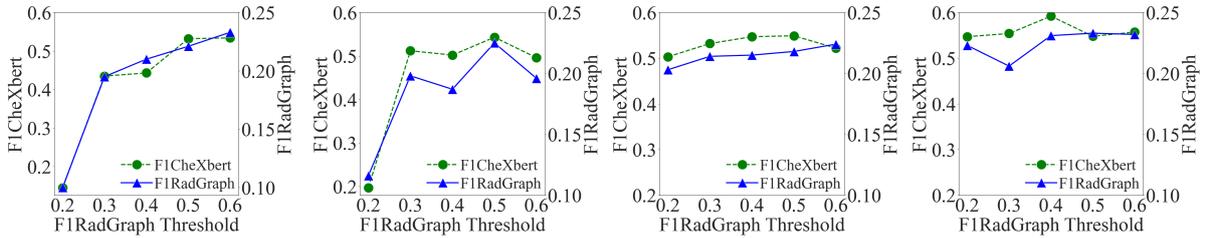
Specifically, compared to the retriever Med-MARVEL, we also observe factual-correctness performance gain based on two clinical metrics. Both use the same universal encoder backbone, but FactMM-RAG benefits from the injected factual medical knowledge, allowing it to search for the most similar and factually correct reports, thereby assisting the multimodal foundation model in generating more accurate reports.

5.2 Ablation Study

Multimodal Retrieval. Instead of relying on the multimodal foundation model to generate

Model	MIMIC-CXR				CheXpert			
	Factual Similarity		Textual Similarity		Factual Similarity		Textual Similarity	
	F1CheXbert	F1RadGraph	ROUGE-L	BERTScore	F1CheXbert	F1RadGraph	ROUGE-L	BERTScore
No Retriever	0.496	0.234	0.294	0.549	0.371	0.173	0.231	0.469
CLIP (Radford et al., 2021)	0.507	0.241	0.300	0.552	0.381	0.172	0.231	0.468
GLoRIA (Huang et al., 2021)	0.476	0.232	0.294	0.543	0.397	0.173	0.231	0.468
MedCLIP (Wang et al., 2022)	0.517	0.238	0.298	0.549	0.408	0.182	0.238	0.471
CXR-CLIP (You et al., 2023)	0.501	0.243	0.302	0.553	0.406	0.183	0.241	0.471
BiomedCLIP (Zhang et al., 2024)	0.502	0.233	0.293	0.546	0.380	0.173	0.232	0.469
Med-MARVEL (Zhou et al., 2024)	0.537	0.237	0.306	0.549	0.454	0.185	0.243	0.472
FactMM-RAG	0.602	0.257	0.307	0.561	0.475	0.185	0.236	0.475

Table 1: Overall performance of FactMM-RAG and baselines under the multimodal retrieval-augmentation setting. Models are evaluated by textual similarity and factual similarity between generated and reference reports. FactMM-RAG outperforms the best baseline with p-value < 0.05.



(a) F1CheXbert Threshold: 0.0 (b) F1CheXbert Threshold: 0.4 (c) F1CheXbert Threshold: 0.8 (d) F1CheXbert Threshold: 1.0

Figure 2: Factual performance of FactMM-RAG controlled by different F1CheXbert and F1RadGraph thresholds.

reports, we also evaluate the performance of the multimodal retrievers by directly encoding radiology images from the testing corpus and searching for the closest report from the training corpus for comparison with ground-truth reports. Table 2 shows that our retriever also achieves the best factual retrieval performance compared to other baselines under this setting across two datasets. This demonstrates that training the multimodal retriever with mined factually-informed report pairs can enhance its radiology image understanding capabilities and directly align it with precise reports.

Backbone Variation. We also investigate the impact of different retriever and foundation model backbones on radiology report generation in Table 2. We initialize our retriever model from two checkpoints: WebQA and ClueWeb in (Zhou et al., 2024). We observe that the ClueWeb checkpoint provides a marginal gain compared to the WebQA checkpoint. This can be attributed to the larger scale of the ClueWeb dataset used for pretraining. We also utilize Med-MARVEL as our retriever backbone, which exhibits similar performance to other backbones after training. This implies that even if our retriever is initialized

with a backbone from a general domain, our factually-informed training strategy enables it to fully leverage medical knowledge and quickly adapt to the radiology-specific domain without degrading performance.

5.3 Fact-aware Capability Control

The factual similarity threshold in Equation 1 plays a critical role in controlling the fact-awareness of our multimodal retriever. We examine the performance of FactMM-RAG under different thresholds, as shown in Figure 2. Not only utilizing F1RadGraph thresholds, we also employ F1CheXbert to curate additional thresholds from the report’s diagnostic labels to mine report pairs.

Under the same F1CheXbert threshold for mining report pairs, we observe that an increase in the F1RadGraph threshold correlates with an improvement in factual performance. However, adopting stricter thresholds for identifying report pairs does not yield further improvements and reaches saturation. After calculating the average number of report pairs per query, we find that high thresholds can exclude many relevant report pairs, as shown in Figure 3. This exclusion results in the potential loss of factually useful pairs, thereby

Model	MIMIC-CXR				CheXpert			
	Factual Similarity		Textual Similarity		Factual Similarity		Textual Similarity	
	F1CheXbert	F1RadGraph	ROUGE-L	BERTScore	F1CheXbert	F1RadGraph	ROUGE-L	BERTScore
Setting: Multimodal Retrieval								
CLIP (Radford et al., 2021)	0.341	0.160	0.238	0.489	0.285	0.130	0.207	0.439
GLoRIA (Huang et al., 2021)	0.346	0.137	0.211	0.453	0.359	0.135	0.216	0.447
MedCLIP (Wang et al., 2022)	0.539	0.198	0.261	0.508	0.478	0.161	0.225	0.454
CXR-CLIP (You et al., 2023)	0.516	0.215	0.277	0.524	0.444	0.167	0.230	0.458
BiomedCLIP (Zhang et al., 2024)	0.502	0.233	0.293	0.546	0.386	0.142	0.216	0.441
Med-MARVEL (Zhou et al., 2024)	0.550	0.212	0.279	0.525	0.479	0.160	0.222	0.454
FactMM-RAG	0.605	0.249	0.297	0.547	0.491	0.174	0.237	0.467
Setting: Multimodal Retrieval Augmented Generation								
ClueWeb-LLaVA _{1.5}	0.602	0.257	0.307	0.561	0.495	0.180	0.239	0.473
WebQA-LLaVA _{1.5}	0.572	0.262	0.304	0.562	0.456	0.184	0.237	0.474
Med-MARVEL-LLaVA _{1.5}	0.581	0.260	0.311	0.563	0.475	0.185	0.236	0.474
ClueWeb-LLaVA _{1.6}	0.601	0.252	0.303	0.558	0.492	0.178	0.237	0.471

Table 2: Ablation study of FactMM-RAG including multimodal retrieval and backbone variation.

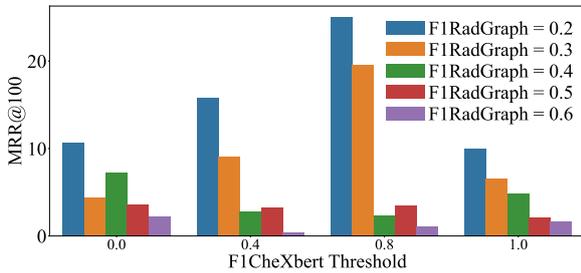
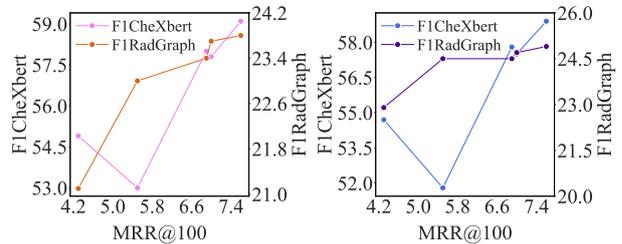


Figure 3: Retrieval evaluation of FactMM-RAG with different F1CheXbert and F1RadGraph thresholds.



(a) Multimodal Retrieval

(b) Multimodal RAG

Figure 4: Analysis of fact-aware capability propagation evaluated by MRR retrieval metric.

hindering the training of our multimodal retriever driven by additional factual medical knowledge.

Rather than relying on diagnostic labels from CheXbert to identify high-quality report pairs, Figure 2a demonstrates that the F1RadGraph threshold alone can also effectively mine factual report pairs for training our multimodal retriever. As the F1RadGraph threshold increases, FactMM-RAG even matches the performance under high threshold settings in Figure 2d. This signifies that employing our training strategy with curated factual query-report pairs still imposes useful supervision signals without relying on explicit diagnostic label guidance.

5.4 Fact-aware Capability Propagation

To further understand the benefits of our retriever for the foundation model, we explore the effective propagation of fact-aware capabilities from the retriever to the foundation model. To demonstrate this behavior, we use the mined factual report pairs as reference reports for the query report. We then use the retrieval metric Mean Reciprocal Rank (MRR) as an intermediate evaluation, shown in Fig-

ure 4. From the plot, we observe that as training progresses, the retrieval metric increases alongside two clinical metrics. This factually-oriented upward trend in our retriever’s performance in Figure 4a is also reflected in the foundation model’s performance in Figure 4b. This indicates that employing a factually-informed reference report selection strategy to train our multimodal retriever can also enhance the foundation model’s ability to generate factually accurate radiology reports.

5.5 Case Study

In this section, we present two examples from MIMIC-CXR to qualitatively analyze our retriever’s fact-aware capability, as illustrated in Table 3. In the first example, we observe that FactMM-RAG provides symptom observations consistent with the ground-truth report and generates more accurate factual details compared to Med-MARVEL, e.g., “post median sternotomy, atelectasis, not pneumothorax”; In the second example, we further observe that although both retrievers generate reports with diagnostic labels matching the ground-truth report, FactMM-RAG provides additional details compared to Med-MARVEL, such

Radiology Image	Med-MARVEL	FactMM-RAG	Reference
	Single portable view of the chest. <i>There are bilateral pleural effusions, moderate on the left and small on the right. There is also pulmonary vascular redistribution and hazy alveolar infiltrate. Cardiac silhouette is enlarged</i> but unchanged. Median sternotomy wires and mediastinal clips are again noted.	A left-sided pacemaker is in place with leads terminating in the right atrium and right ventricle. <i>The patient is status post median sternotomy and CABG.</i> The heart is moderately enlarged. <i>There is mild pulmonary edema. A small left pleural effusion is present. There is atelectasis at the left lung base. No pneumothorax is seen.</i>	<i>The patient is status post median sternotomy and CABG.</i> Left-sided dual-chamber pacemaker is noted with leads terminating in right atrium and right ventricle, unchanged. Cardiomegaly is similar. <i>There is continued mild to moderate pulmonary edema,</i> slightly improved compared to the prior exam. <i>Small layering bilateral pleural effusions</i> also may be slightly decreased in the interval. <i>Bibasilar airspace opacities likely reflect atelectasis. There is no pneumothorax.</i> No acute osseous abnormalities are visualized.
F1RadGraph	0.218	0.413	
CheXbert Observations	Cardiomegaly, Edema, Pleural Effusion	Cardiomegaly, Edema, Atelectasis, Pleural Effusion	Cardiomegaly, Edema, Atelectasis, Pleural Effusion
	<i>The heart is mildly enlarged. The aorta is mildly tortuous.</i> The mediastinal and hilar contours appear unchanged. <i>There is no pleural effusion or pneumothorax.</i> Streaky left basilar opacity suggests minor atelectasis. <i>There is no definite pleural effusion or pneumothorax.</i> The bones appear demineralized. There is mild-to-moderate rightward convex curvature centered along the mid thoracic spine.	<i>Heart size is mildly enlarged. The aorta is tortuous.</i> Mediastinal and hilar contours are otherwise unremarkable. <i>Pulmonary vasculature is normal.</i> Linear opacities in the left lower lobe are compatible with subsegmental atelectasis. <i>No focal consolidation, pleural effusion or pneumothorax is present. There are no acute osseous abnormalities.</i>	<i>Moderate enlargement of the cardiac silhouette</i> with a left ventricular predominance is unchanged. <i>The aorta remains tortuous,</i> and the hilar contours are stable. <i>Pulmonary vascularity is not engorged.</i> There is minimal atelectasis within the lung bases, <i>but no focal consolidation is present. No pleural effusion or pneumothorax is identified. There are no acute osseous abnormalities.</i>
F1RadGraph	0.333	0.526	
CheXbert Observations	Cardiomegaly, Atelectasis	Cardiomegaly, Atelectasis	Cardiomegaly, Atelectasis

Table 3: One case study from MIMIC-CXR. Cyan text indicates radiological consistency with the ground-truth report. Orange text highlights extra accurate details provided by FactMM-RAG compared to Med-MARVEL. Red text denotes observations missing in Med-MARVEL.

as “pulmonary vasculature is normal, no acute osseous abnormalities”. These characteristics confirm that adopting our fact-aware retriever can assist multimodal foundation models in generating more accurate radiology reports.

6 Conclusion

In this paper, we aim at improving radiology report generation by introducing a fact-informed medical multimodal retriever for retrieval-augmented generation. In particular, we utilize RadGraph to annotate chest radiograph reports and mine clinically-relevant pairs. We integrate factual information into a universal multimodal retriever, presenting FactMM-RAG, a fact-aware multimodal retrieval-augmented radiology report generation pipeline. FactMM-RAG outperforms all state-of-the-art retrievers evaluated by factual correctness and textual coherence for final report generation in MIMIC-CXR and CheXpert datasets. We further confirm the benefit of our multimodal retriever from the analysis of fact-aware capability control and propagation. Given the pervasive applications of machine learning in clinical diagnoses using chest X-rays, we hope our factual-informed approach

inspires further work in multimodal generative artificial intelligence in healthcare contexts.

7 Limitations

Despite the strong performance of our FactMM-RAG pipeline, we acknowledge potential limitations of our proposed method. In particular, our work only emphasizes chest radiology domains. It also worth exploring our retrieval-augmented factual report generation pipeline in broader medical domains, such as brain scan or histology datasets.

Another concern lies in the chosen evaluation metrics, F1RadGraph and F1CheXbert. F1CheXbert reflects high-level observational accuracy, while F1RadGraph assesses the correctness of radiology entities and clinical relationships. However, other radiologically-specific metrics, such as report conciseness and clarity, should also be considered (Sureka et al., 2014). Ideally, we should incorporate methods of evaluation directly aligned with human evaluations or involve domain expertise itself in our pair-mining and final evaluation procedure.

588
589
590
591
592

593
594
595
596
597
598
599
600
601
602
603

604
605
606
607
608
609
610

611
612
613
614
615

616
617
618
619
620
621
622
623

624
625
626
627
628
629
630
631
632

633
634
635
636
637
638
639
640

641
642
643
644
645

References

Muhammad Aurangzeb Ahmad, Ilker Yaramis, and Taposh Dutta Roy. 2023. [Creating trustworthy llms: Dealing with hallucinations in healthcare ai](#). *Preprint*, arXiv:2311.01463.

Sebastian Borgeaud, Arthur Mensch, Jordan Hoffmann, Trevor Cai, Eliza Rutherford, Katie Millican, George van den Driessche, Jean-Baptiste Lespiau, Bogdan Damoc, Aidan Clark, Diego de Las Casas, Aurelia Guy, Jacob Menick, Roman Ring, Tom Hennigan, Saffron Huang, Loren Maggiore, Chris Jones, Albin Cassirer, Andy Brock, Michela Paganini, Geoffrey Irving, Oriol Vinyals, Simon Osindero, Karen Simonyan, Jack W. Rae, Erich Elsen, and Laurent Sifre. 2022. [Improving language models by retrieving from trillions of tokens](#). *Preprint*, arXiv:2112.04426.

Pierre Chambon, Christian Bluethgen, Jean-Benoit Delbrouck, Rogier Van der Sluijs, Małgorzata Połacin, Juan Manuel Zambrano Chaves, Tanishq Mathew Abraham, Shivanshu Purohit, Curtis P. Langlotz, and Akshay Chaudhari. 2022. [Roentgen: Vision-language foundation model for chest x-ray generation](#). *Preprint*, arXiv:2211.12737.

Wenhu Chen, Hexiang Hu, Xi Chen, Pat Verga, and William W. Cohen. 2022. [Murag: Multimodal retrieval-augmented generator for open question answering over images and text](#). *Preprint*, arXiv:2210.02928.

Zhihong Chen, Yaling Shen, Yan Song, and Xiang Wan. 2021. [Cross-modal memory networks for radiology report generation](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 5904–5914, Online. Association for Computational Linguistics.

Zhihong Chen, Maya Varma, Jean-Benoit Delbrouck, Magdalini Paschali, Louis Blankemeier, Dave Van Veen, Jeya Maria Jose Valanarasu, Alaa Youssef, Joseph Paul Cohen, Eduardo Pontes Reis, Emily B. Tsai, Andrew Johnston, Cameron Olsen, Tanishq Mathew Abraham, Sergios Gatidis, Akshay S. Chaudhari, and Curtis Langlotz. 2024. [Chexagent: Towards a foundation model for chest x-ray interpretation](#). *Preprint*, arXiv:2401.12208.

Jean-Benoit Delbrouck, Pierre Chambon, Christian Bluethgen, Emily Tsai, Omar Almusa, and Curtis Langlotz. 2022. [Improving the Factual Correctness of Radiology Report Generation with Semantic Rewards](#). In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 4348–4360, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.

Jean-Benoit Delbrouck, Maya Varma, Pierre Chambon, and Curtis Langlotz. 2023. [Overview of the RadSum23 Shared Task on Multi-modal and Multi-anatomical Radiology Report Summarization](#). In *The 22nd Workshop on Biomedical Natural Language*

Processing and BioNLP Shared Tasks, pages 478–482, Toronto, Canada. Association for Computational Linguistics. 646
647
648

Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. 2021. [An image is worth 16x16 words: Transformers for image recognition at scale](#). *Preprint*, arXiv:2010.11929. 649
650
651
652
653
654
655

Mark Endo, Rayan Krishnan, Viswesh Krishna, Andrew Y. Ng, and Pranav Rajpurkar. 2021. [Retrieval-based chest x-ray report generation using a pre-trained contrastive language-image model](#). In *Proceedings of Machine Learning for Health*, volume 158 of *Proceedings of Machine Learning Research*, pages 209–219. PMLR. 656
657
658
659
660
661
662

Yunfan Gao, Yun Xiong, Xinyu Gao, Kangxiang Jia, Jinliu Pan, Yuxi Bi, Yi Dai, Jiawei Sun, Meng Wang, and Haofen Wang. 2024. [Retrieval-augmented generation for large language models: A survey](#). *Preprint*, arXiv:2312.10997. 663
664
665
666
667

Kelvin Guu, Kenton Lee, Zora Tung, Panupong Pasupat, and Ming-Wei Chang. 2020. [Realm: Retrieval-augmented language model pre-training](#). *Preprint*, arXiv:2002.08909. 668
669
670
671

Ziniu Hu, Ahmet Iscen, Chen Sun, Zirui Wang, Kai-Wei Chang, Yizhou Sun, Cordelia Schmid, David A. Ross, and Alireza Fathi. 2023. [Reveal: Retrieval-augmented visual-language pre-training with multi-source multimodal knowledge memory](#). *Preprint*, arXiv:2212.05221. 672
673
674
675
676
677

Shih-Cheng Huang, Liyue Shen, Matthew P Lungren, and Serena Yeung. 2021. [GLORIA: A Multimodal Global-Local Representation Learning Framework for Label-Efficient Medical Image Recognition](#). In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3942–3951. 678
679
680
681
682
683

Jeremy Irvin, Pranav Rajpurkar, Michael Ko, Yifan Yu, Silviana Ciurea-Ilcus, Chris Chute, Henrik Marklund, Behzad Haghgoo, Robyn Ball, Katie Shpankaya, 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. 2019. [CheXpert: A Large Chest Radiograph Dataset with Uncertainty Labels and Expert Comparison](#). *Preprint*, arXiv:1901.07031. 684
685
686
687
688
689
690
691
692
693

Lucky Iyeke, Rebecca Moss, Rachel Hall, Jun Wang, Lovleen Sandhu, Benjamin Appold, Ella Kalontar, Dimitra Menoudakos, Mahesh Ramnarine, Samuel P. LaVine, Seungwoo Ahn, and Michelle Richman. 2022. [Reducing unnecessary 'admission' chest x-rays: An initiative to minimize low-value care](#). *Cureus*, 14(10):e29817. 694
695
696
697
698
699
700

Gautier Izacard, Patrick Lewis, Maria Lomeli, Lucas Hosseini, Fabio Petroni, Timo Schick, Jane 701
702

703	Dwivedi-Yu, Armand Joulin, Sebastian Riedel, and	Michael Moor, Qian Huang, Shirley Wu, Michihiro	759
704	Edouard Grave. 2022. Atlas: Few-shot learning	Yasunaga, Yash Dalmia, Jure Leskovec, Cyril Za-	760
705	with retrieval augmented language models . <i>Preprint</i> ,	kka, Eduardo Pontes Reis, and Pranav Rajpurkar.	761
706	arXiv:2208.03299.	2023. Med-flamingo: a multimodal medical few-shot	762
707	Saahil Jain, Ashwin Agrawal, Adriel Saporta,	learner . In <i>Proceedings of the 3rd Machine Learning</i>	763
708	Steven QH Truong, Du Nguyen Duong, Tan Bui,	<i>for Health Symposium</i> , volume 225 of <i>Proceedings of</i>	764
709	Pierre Chambon, Yuhao Zhang, Matthew P. Lungren,	<i>Machine Learning Research</i> , pages 353–367. PMLR.	765
710	Andrew Y. Ng, Curtis P. Langlotz, and Pranav Ra-	Sahal Shaji Mullappilly Hisham Cholakkal Rao	766
711	ipurkar. 2021. RadGraph: Extracting Clinical Enti-	Muhammad Anwer Salman Khan Jorma Laakso-	767
712	ties and Relations from Radiology Reports . <i>Preprint</i> ,	nen Omkar Thawkar, Abdelrahman Shaker and Fa-	768
713	arXiv:2106.14463.	had Shahbaz Khan. 2023. Xraygpt: Chest radio-	769
714	Alistair E. W. Johnson, Tom J. Pollard, Seth J.	graphs summarization using large medical vision-	770
715	Berkowitz, Nathaniel R. Greenbaum, Matthew P.	language models. <i>arXiv: 2306.07971</i> .	771
716	Lungren, Chih-ying Deng, Roger G. Mark, and	Ankit Pal and Malaikannan Sankarasubbu. 2024. Gem-	772
717	Steven Horng. 2019. MIMIC-CXR, a de-identified	ini goes to med school: Exploring the capabili-	773
718	publicly available database of chest radiographs with	ties of multimodal large language models on med-	774
719	free-text reports . <i>Scientific Data</i> , 6(1).	ical challenge problems hallucinations . <i>Preprint</i> ,	775
720	Vladimir Karpukhin, Barlas Oğuz, Sewon Min, Patrick	arXiv:2402.07023.	776
721	Lewis, Ledell Wu, Sergey Edunov, Danqi Chen,	Ankit Pal, Logesh Kumar Umapathi, and Malaikannan	777
722	and Wen tau Yih. 2020. Dense passage retrieval	Sankarasubbu. 2023. Med-halt: Medical domain hal-	778
723	for open-domain question answering . <i>Preprint</i> ,	lucination test for large language models . <i>Preprint</i> ,	779
724	arXiv:2004.04906.	arXiv:2307.15343.	780
725	Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio	Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya	781
726	Petroni, Vladimir Karpukhin, Naman Goyal, Hein-	Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sas-	782
727	rich Küttler, Mike Lewis, Wen tau Yih, Tim Rock-	try, Amanda Askell, Pamela Mishkin, Jack Clark,	783
728	täschel, Sebastian Riedel, and Douwe Kiela. 2021.	Gretchen Krueger, and Ilya Sutskever. 2021. Learn-	784
729	Retrieval-Augmented Generation for Knowledge-	ing Transferable Visual Models From Natural Lan-	785
730	Intensive NLP Tasks . <i>Preprint</i> , arXiv:2005.11401.	guage Supervision . <i>Preprint</i> , arXiv:2103.00020.	786
731	Christy Y. Li, Xiaodan Liang, Zhiting Hu, and Eric P.	Colin Raffel, Noam Shazeer, Adam Roberts, Katherine	787
732	Xing. 2018. Hybrid retrieval-generation reinforced	Lee, Sharan Narang, Michael Matena, Yanqi Zhou,	788
733	agent for medical image report generation . <i>Preprint</i> ,	Wei Li, and Peter J. Liu. 2023. Exploring the limits	789
734	arXiv:1805.08298.	of transfer learning with a unified text-to-text trans-	790
735	Chunyuan Li, Cliff Wong, Sheng Zhang, Naoto	former . <i>Preprint</i> , arXiv:1910.10683.	791
736	Usuyama, Haotian Liu, Jianwei Yang, Tristan Nau-	Vignav Ramesh, Nathan Andrew Chi, and Pranav Ra-	792
737	mann, Hoifung Poon, and Jianfeng Gao. 2023. Llava-	ipurkar. 2022. Improving radiology report generation	793
738	med: Training a large language-and-vision assis-	systems by removing hallucinated references to non-	794
739	tant for biomedicine in one day . <i>arXiv preprint</i>	existent priors . <i>Preprint</i> , arXiv:2210.06340.	795
740	<i>arXiv:2306.00890</i> .	Weijia Shi, Sewon Min, Michihiro Yasunaga, Min-	796
741	Chin-Yew Lin. 2004. Rouge: A package for automatic	joon Seo, Rich James, Mike Lewis, Luke Zettle-	797
742	evaluation of summaries . In <i>Annual Meeting of the</i>	moyer, and Wen tau Yih. 2023. Replug: Retrieval-	798
743	<i>Association for Computational Linguistics</i> .	augmented black-box language models . <i>Preprint</i> ,	799
744	Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae	arXiv:2301.12652.	800
745	Lee. 2023. Visual Instruction Tuning . <i>Preprint</i> ,	Akshay Smit, Saahil Jain, Pranav Rajpurkar, Anuj Pa-	801
746	arXiv:2304.08485.	reek, Andrew Y. Ng, and Matthew P. Lungren. 2020.	802
747	Kang Liu, Zhuoqi Ma, Mengmeng Liu, Zhicheng Jiao,	Chexbert: Combining automatic labelers and expert	803
748	Xiaolu Kang, Qiguang Miao, and Kun Xie. 2024.	annotations for accurate radiology report labeling	804
749	Factual serialization enhancement: A key innova-	using bert . <i>Preprint</i> , arXiv:2004.09167.	805
750	tion for chest x-ray report generation . <i>Preprint</i> ,	Annemiek M. Speets, Yolanda van der Graaf, Arno W.	806
751	arXiv:2405.09586.	Hoes, Sjouke Kalmijn, Arnold P. Sachs, Matthijs J.	807
752	Ilya Loshchilov and Frank Hutter. 2019. De-	Rutten, Jan W. Gratama, Annemiek D. Mon-	808
753	coupled weight decay regularization . <i>Preprint</i> ,	tauban van Swijndregt, and Willem P. Mali. 2006.	809
754	arXiv:1711.05101.	Chest radiography in general practice: Indications,	810
755	Yasuhide Miura, Yuhao Zhang, Emily Bao Tsai, Curtis P.	diagnostic yield, and consequences for patient man-	811
756	Langlotz, and Dan Jurafsky. 2021. Improving factual	agement . <i>The British Journal of General Practice:</i>	812
757	completeness and consistency of image-to-text radiol-	<i>The Journal of the Royal College of General Practi-</i>	813
758	ogy report generation . <i>Preprint</i> , arXiv:2010.10042.	<i>tioners</i> , 56(529):574–578.	814

815	Liwen Sun, Abhineet Agarwal, Aaron Kornblith, Bin Yu, and Chenyan Xiong. 2024. Ed-copilot: Reduce emergency department wait time with language model diagnostic assistance . <i>Preprint</i> , arXiv:2402.13448.	872
816		873
817		874
818		875
819		876
820	Binit Sureka et al. 2014. Seven c’s of effective radiology reporting . <i>The Journal of National Accreditation Board for Hospitals & Healthcare Providers</i> , 1(1):17. Accessed 14 June 2024.	877
821		878
822		879
823		880
824	Tao Tu, Shekoofeh Azizi, Danny Driess, Mike Schaekermann, Mohamed Amin, Pi-Chuan Chang, Andrew Carroll, Chuck Lau, Ryutaro Tanno, Ira Ktena, 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. 2023. Towards generalist biomedical ai . <i>Preprint</i> , arXiv:2307.14334.	881
825		882
826		883
827		884
828		885
829		886
830		887
831		888
832		889
833		890
834		
835		
836	Zifeng Wang, Zhenbang Wu, Dinesh Agarwal, and Jimeng Sun. 2022. Medclip: Contrastive learning from unpaired medical images and text . <i>Preprint</i> , arXiv:2210.10163.	891
837		892
838		893
839		894
840		895
841	Chaoyi Wu, Xiaoman Zhang, Ya Zhang, Yanfeng Wang, and Weidi Xie. 2023. Towards generalist foundation model for radiology by leveraging web-scale 2d3d medical data . <i>Preprint</i> , arXiv:2308.02463.	896
842		897
843		898
844		899
845	Qianqian Xie, Jiayu Zhou, Yifan Peng, and Fei Wang. 2023. Factreranker: Fact-guided reranker for faithful radiology report summarization . <i>Preprint</i> , arXiv:2303.08335.	900
846		
847		
848	Michihiro Yasunaga, Armen Aghajanyan, Weijia Shi, Rich James, Jure Leskovec, Percy Liang, Mike Lewis, Luke Zettlemoyer, and Wen tau Yih. 2023. Retrieval-augmented multimodal language modeling . <i>Preprint</i> , arXiv:2211.12561.	901
849		902
850		903
851		904
852		905
853	Kihyun You, Jawook Gu, Jiyeon Ham, Beomhee Park, Jiho Kim, Eun K. Hong, Woonhyuk Baek, and Byungseok Roh. 2023. Cxr-clip: Toward large scale chest x-ray language-image pre-training . In <i>Medical Image Computing and Computer Assisted Intervention – MICCAI 2023</i> , pages 101–111. Springer Nature Switzerland.	906
854		907
855		908
856		909
857		910
858		911
859		912
860	Shi Yu, Zhenghao Liu, Chenyan Xiong, and Zhiyuan Liu. 2023a. Openmatch-v2: An all-in-one multimodality plm-based information retrieval toolkit . pages 3160–3164.	913
861		914
862		915
863		916
864	Zichun Yu, Chenyan Xiong, Shi Yu, and Zhiyuan Liu. 2023b. Augmentation-adapted retriever improves generalization of language models as generic plug-in . <i>Preprint</i> , arXiv:2305.17331.	917
865		918
866		919
867		920
868	Sheng Zhang, Yanbo Xu, Naoto Usuyama, Hanwen Xu, Jaspreet Bagga, Robert Tinn, Sam Preston, Rajesh Rao, Mu Wei, Naveen Valluri, Cliff Wong, Andrea Tupini, Yu Wang, Matt Mazzola, Swadheen	921
869		
870		
871		
	Shukla, Lars Liden, Jianfeng Gao, Matthew P. Lungren, Tristan Naumann, Sheng Wang, and Hoifung Poon. 2024. Biomedclip: a multimodal biomedical foundation model pretrained from fifteen million scientific image-text pairs . <i>Preprint</i> , arXiv:2303.00915.	
	Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q. Weinberger, and Yoav Artzi. 2020a. Bertscore: Evaluating text generation with bert . <i>Preprint</i> , arXiv:1904.09675.	
	Yuhao Zhang, Hang Jiang, Yasuhide Miura, Christopher D. Manning, and Curtis P. Langlotz. 2022. Contrastive Learning of Medical Visual Representations from Paired Images and Text . <i>Preprint</i> , arXiv:2010.00747.	
	Yuhao Zhang, Derek Merck, Emily Bao Tsai, Christopher D. Manning, and Curtis P. Langlotz. 2020b. Optimizing the factual correctness of a summary: A study of summarizing radiology reports . <i>Preprint</i> , arXiv:1911.02541.	
	Tianshuo Zhou, Sen Mei, Xinze Li, Zhenghao Liu, Chenyan Xiong, Zhiyuan Liu, Yu Gu, and Ge Yu. 2024. MARVEL: Unlocking the Multi-Modal Capability of Dense Retrieval via Visual Module Plugin . <i>Preprint</i> , arXiv:2310.14037.	
	Erdi Çallı, Ecem Sogancioglu, Bram van Ginneken, Kicky G. van Leeuwen, and Keelin Murphy. 2021. Deep learning for chest x-ray analysis: A survey . <i>Medical Image Analysis</i> , 72:102125.	
	A Appendix	
	A.1 Retriever Training Procedure	
	To training our fact-aware multimodal retriever, we not only use mined factual report pairs as positive reports to the query image, but also incorporate the query image’s corresponding report. Following (Yu et al., 2023a; Zhou et al., 2024), we also adopt modality-balanced hard negatives to train the retriever after in-batch negative training from the multimodal dense retrieval stage. We use AdamW (Loshchilov and Hutter, 2019) as our optimizer and training epochs = 15, early stopping epoch = 5, batch size = 32, learning rate = 5e-6, and the temperature hyperparameter $\tau = 0.01$. For our MARVEL backbone, we use T5-ANCE (Yu et al., 2023a) as the text encoder and vision transformer (Dosovitskiy et al., 2021) as the vision encoder.	
	A.2 RAG Finetuning Procedure	
	To create a RAG dataset for fine-tuning LLaVA, we search the nearest-neighbor document d_{txt}^* for a query image q_{img} using a retriever’s embeddings. We filter out any results that involve retrieving a	

Visual Question Answering:

Generate a radiology report from this image:
<image>

Retrieval Augmented Generation:

Here is a report of a related patient:
"<document>"
Generate a radiology report from this image:
<image>

Figure 5: Prompt templates for Visual Question Answering and Retrieval Augmented Generation

922 patient's own report, the same patient's other stud-
923 ies, or malformed reports in the training dataset
924 (specified by being less than 5 characters). We
925 apply the prompt templates in Figure 5, and fine-
926 tune LLaVA-1.5 for one epoch. Models are trained
927 using 8x NVIDIA RTX A6000, with epochs=1,
928 learning rate=2e-5, global batch size=128, from
929 vicuna-7b-v1.5 checkpoint. We save the check-
930 point after one full pass of the training dataset for
931 final evaluation.