

Multimodal Prompt Retrieval for Generative Visual Question Answering

Timothy Ossowski¹, Junjie Hu^{1,2}

¹Department of Computer Science, ²Department of Biostatistics and Medical Informatics
University of Wisconsin, Madison, WI, USA

ossowski@wisc.edu, junjie.hu@wisc.edu

Abstract

Recent years have witnessed impressive results of pre-trained vision-language models on knowledge-intensive tasks such as visual question answering (VQA). Despite the recent advances in VQA, existing methods mainly adopt a discriminative formulation that predicts answers within a pre-defined label set, leading to easy overfitting on low-resource domains with limited labeled data (e.g., medicine) and poor generalization under domain shift to another dataset. To tackle this limitation, we propose a novel generative model enhanced by multimodal prompt retrieval (MPR) that integrates retrieved prompts and multimodal features to generate answers in free text. Our generative model enables rapid zero-shot dataset adaptation to unseen data distributions and open-set answer labels across datasets. Our experiments on medical VQA tasks show that MPR outperforms its non-retrieval counterpart by up to 30% accuracy points in a few-shot domain adaptation setting.¹

1 Introduction

Visual question answering (VQA) is a popular multimodal machine learning problem that challenges a model to answer a question posed about an image. As encouraged by recent advances in VQA, pioneering studies have investigated the application of VQA systems to low-resourced, knowledge-intensive domains such as medicine (Lin et al., 2021), where collecting domain-specific annotations is extremely costly and time-consuming. In particular, medical VQA has attracted increasing research interests (Hasan et al., 2018), with the target of supporting clinical decision-making such as acting as an auxiliary virtual “diagnostic radiologist” (Kovaleva et al., 2020).

Despite recent progress in general VQA leveraging pre-training (Chen et al., 2022), retrieval (Wu

et al., 2022), or knowledge bases (Narasimhan and Schwing, 2018; Shevchenko et al., 2021), several challenges still exist for medical VQA. First, medical VQA systems still suffer from a stark lack of high-quality labeled data. As a result, it is essential to leverage domain adaptation techniques (Zhou et al., 2019) that rapidly adapt models trained from a similar dataset to a target dataset. Second, as the medical domain covers a wide variety of complex diseases, there exists a large distribution shift across medical datasets, significantly increasing the complexity of learning medical images and texts by deep neural models. However, many existing medical VQA methods mainly focus on in-domain evaluation, testing systems on a held-out test set under the same data distribution of the training data. Moreover, these methods often augment their model architecture with dataset-specific components such as an answer-type classifier (Zhan et al., 2020), separate models for each question-type (Khare et al., 2021), or specific pre-trained medical encoders (Moon et al., 2022). These dataset-specific designs hinder the application of these medical VQA models across datasets in new domains. Furthermore, existing medical VQA approaches (Tanwani et al., 2022; Eslami et al., 2021) often adopt a discriminative model architecture that predicts a fixed set of answers, limiting model generalization to different answer sets.

To tackle these challenges, we propose a domain-agnostic generative VQA model with multimodal prompt retrieval (MPR) that retrieves relevant VQA examples to construct multimodal prompts and generates arbitrary free text as the answers, removing the restriction of predicting a fixed label set. To augment the retrieval data, we also investigate a data augmentation strategy to create a synthetic medical VQA dataset from medical image-captioning data. Our experiments on two medical VQA datasets demonstrate the effective adaptation of our proposed method to a new target medical dataset,

¹Our code is publicly available at <https://github.com/tossowski/MultimodalPromptRetrieval>

while also showing similar in-domain performance of our models to existing discriminative baselines.

Our contributions are summarized below:

- We introduce a multimodal prompt retrieval module that improves VQA generalization across different data distributions even with noisy synthetic data and smaller retrieval datasets.
- We investigate a zero-shot dataset adaptation setting for medical VQA systems across datasets, encouraging future research on in-context prediction of VQA systems for dataset adaptation.
- We propose a novel prompt-based generative VQA model, which enables more flexible answer outputs and controllable generation guided by multimodal prompts.

2 Preliminaries

This section provides descriptions of the VQA task and the challenges faced in the medical domain.

Problem Setup Formally, given a VQA dataset of n tuples $\mathcal{D} = \{(v_i, x_i, y_i)\}_{i=1}^n$, we aim to learn a model to predict an answer y_i given a question x_i and an image v_i . Conventionally, a model consists of an image and text encoder that maps the inputs v_i and x_i to the latent space of \mathcal{V} and \mathcal{X} respectively:

$$\mathbf{v}_i = \text{ImgEncoder}(v_i) \in \mathcal{V} \quad (1)$$

$$\mathbf{x}_i = \text{TextEncoder}(x_i) \in \mathcal{X} \quad (2)$$

Most prior works learn a discriminative model f_θ that directly estimates a probability distribution over all possible answers in a pre-defined label set, i.e., $f_\theta : \mathcal{V}, \mathcal{X} \rightarrow \mathcal{Y}$. In contrast, we adopt a generative model g_ϕ that predicts words in a vocabulary Σ to generate a varying length text string $z \in \Sigma^+$, and apply a deterministic function to map the answer string z to the closest answer label $y \in \mathcal{Y}$.

Dataset Adaptation: We also focus on a dataset adaptation setting where a model is trained on a source labeled dataset \mathcal{D}_{src} and further adapted to a target dataset \mathcal{D}_{tgt} with a different label set, i.e., $\mathcal{Y}_{\text{src}} \neq \mathcal{Y}_{\text{tgt}}$. Thus, it is nontrivial for a discriminative model f_θ to perform adaptation over different label sets. For adaptation with generative models, we consider two strategies of using target labeled data for (a) in-context prediction without updating the source-trained models g_ϕ and (b) continued fine-tuning g_ϕ . While our method focuses on in-context prediction (§3), we also compare these two strategies in our experiments (§5).

Types of Medical VQA: According to the annotations of popular medical VQA tasks (e.g., SLAKE (Liu et al., 2021) and VQA-RAD (Lau et al., 2018)), there are two answer types $\mathcal{A}_{\text{type}}$: *closed* answers where the set of possible answers are disclosed in a question (e.g., yes-no questions); and *open* answers that can be free-form texts. Besides, there are multiple different question types $\mathcal{Q}_{\text{type}}$ such as organ, abnormality, or modality, indicating the medicinal category for which the question is intended. Prior medical VQA models (Zhan et al., 2020; Eslami et al., 2021) use a binary classifier to distinguish the two answer types based on questions and apply two discriminative models to predict answers, while we propose to predict both types of answers by a single generative model in this work.

3 Methods

In this section, we start by introducing the text and image encoding for retrieval (§3.1), then describe the prompt construction from retrieval (§3.3), and prompt integration in our generative model (§3.4).

Overview: For each $(v, x, y) \in \mathcal{D}_{\text{src}}$ during training, we propose to retrieve similar tuples from the training dataset \mathcal{D}_{src} , integrate the retrieved tuples for prediction, and update the model. We also assume to have access to a target labeled dataset \mathcal{D}_{tgt} for dataset adaptation. Note that we mainly describe the in-context prediction using \mathcal{D}_{tgt} here and leave the discussion of fine-tuning on \mathcal{D}_{tgt} to the experiments. When predicting a target test example at test time, we directly apply our source-trained model to retrieve labeled tuples from \mathcal{D}_{tgt} and perform prediction. The key insight is that even if the source-trained model is not directly trained on target data, the retrieved tuples may contain the correct answer to the given target question, potentially improving model predictions in the target dataset.

3.1 Multimodal Prompt Encoding

For a VQA dataset, we can easily construct a mapping by using the image-question pair as the key and the answer as the value. Therefore we can use a multimodal encoder to encode the image-question pairs into multimodal features and perform K-Nearest Neighbors (KNN) search to find the most similar VQA tuples in the feature space.

Question-Image Encoding: Before model training, we use a pre-trained CLIP model (Radford

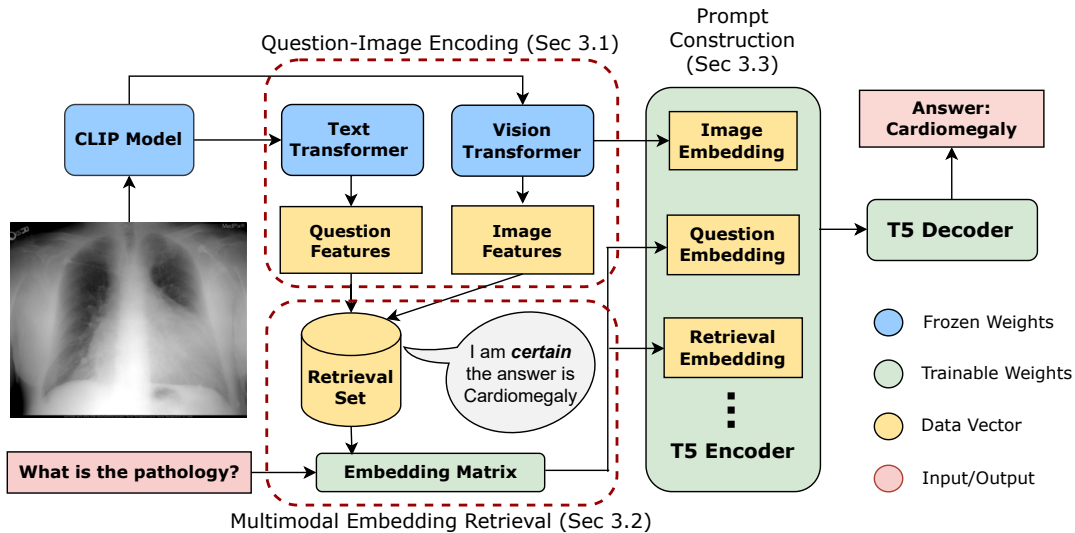


Figure 1: An overview of multimodal prompt retrieval (MPR). There are three primary components of the prompt we use for the encoder indicated by the three yellow boxes contained in the T5 encoder block. These components can be optionally omitted or further extended with additional data.

et al., 2021) to encode image-question pairs in a retrieval dataset \mathcal{R} , where $\mathcal{R} = \mathcal{D}_{\text{src}}$ during training and $\mathcal{R} = \mathcal{D}_{\text{tgt}}$ at testing. Specifically, we first preprocess each image by downsampling it to the 224×224 resolution and adopt CLIP’s vision transformer (Dosovitskiy et al., 2021) to obtain image features \mathbf{v}_{CLS} of the image patch [CLS] token that summarizes the image content. Similarly, we process each question using CLIP’s corresponding text transformer to obtain question features \mathbf{x}_{EOT} from the [EOT] token. These question and image features are concatenated to form a holistic vector representation $\mathbf{p} = [\mathbf{v}_{\text{CLS}}; \mathbf{x}_{\text{EOT}}]$ of a question-image pair. These question-image vectors are paired along with the corresponding answers to construct the retrieval mapping set $\mathcal{M} = \{(\mathbf{p}_i, y_i)\}_{i=1}^m$ of m items.

Retrieval Set Augmentation on Image-Caption

Data: As many VQA datasets in a low-resourced target domain (e.g., medicine) often contain a limited amount of labeled examples, we propose a data augmentation method to create a synthetic VQA set \mathcal{D}_{syn} from image-caption pairs and augment the retrieval set \mathcal{R} . First, we determine a desired set of question types $\mathcal{Q}_{\text{type}}$ and answer types $\mathcal{A}_{\text{type}}$ described in §2. For each combination of question and answer types $t \in \mathcal{Q}_{\text{type}} \times \mathcal{A}_{\text{type}}$, we manually prepare a collection of question templates \mathcal{T}_t along with a corresponding collection of keywords \mathcal{W}_t . We then iterate through all the image-caption pairs and identify if the caption contains any keywords $w \in \mathcal{W}_t$. If any keywords match, we create a question by sampling a template from \mathcal{T}_t uniformly at random and filling it with the matched keyword as the answer. Example templates from several ques-

tion and answer types can be found in Appendix A.

3.2 Multimodal Embedding Retrieval

To answer a question x about an image v , we propose to retrieve its top- k most similar examples from the retrieval mapping set \mathcal{M} (as constructed in §3.1). Specifically, we first encode the query question-image pair into an embedding \mathbf{p} by the CLIP model and compute the cosine similarity between the query embedding \mathbf{p} and each question-image embedding in \mathcal{M} . Therefore, we can obtain the k nearest neighbors of image-question pairs in \mathcal{M} , denoted as $\mathcal{K} = \{(\mathbf{p}_i, y_i)\}_{i=1}^k$. Note that if the size of \mathcal{M} is large, KNN search can be implemented with efficient algorithms such as Maximum Inner Product Search (Shrivastava and Li, 2014). The retrieved pairs are used to construct the retrieval prompt (detailed in §3.3).

3.3 Prompt Construction

Inspired by the prompt tuning method (Lester et al., 2021) that appends several prompt embeddings to the original input before feeding to the transformer layers of the encoder, we propose to construct multimodal prompt embeddings to augment a question input, as shown in Figure 1. Specifically, given an image-question pair, the inputs to our model consist of three main components: image, question, and retrieval embeddings. Our model concatenates these embeddings as inputs to the subsequent stack of encoder layers in a T5 model (Raffel et al., 2020). We begin the prompt with the image embedding, followed by the question and retrieval embeddings,

leaving experimentation with alternative concatenation orders to Appendix B.

Image Embedding: The image embedding is obtained using the same vision transformer of CLIP applied to construct the retrieval dataset. However, instead of using the [CLS] token which summarizes the image content, we use the intermediate output of the penultimate layer to obtain a collection of image token embeddings $\mathbf{v}_p \in \mathbb{R}^{l_v \times d}$, where l_v denotes the number of image tokens.

Question Embedding: To encode a question corresponding to an image, we use the embedding matrix of a pre-trained T5 encoder. Following the practice of T5, we include a short text snippet (e.g., “Answer the abnormality question:”) at the beginning of the question to instruct the model to perform a QA task. The combined text is first tokenized according to T5’s subword tokenization, followed by an embedding lookup to fetch the corresponding embedding vectors $\mathbf{x}_q \in \mathbb{R}^{l_q \times d}$ from T5’s input embedding matrix.

Retrieval Embedding: Based on the top- k similar examples retrieved \mathcal{K} , we define an ordered list of quantifier words $\mathcal{Q} = [q_1, \dots, q_M]$ (e.g., [very unlikely, ..., very likely, certainly]) and a text template T_{prompt} . We define a confidence score that counts the frequency of the retrieved answers in \mathcal{K} , and then select the most frequent answer y_r^* from \mathcal{K} in Eq. (3). We then apply a threshold function to select an appropriate quantifier q_r^* from \mathcal{Q} based on the confidence score of y_r^* by Eq. (4).

$$\mathbf{p}_r^*, y_r^* = \arg \max_{(\mathbf{p}, y) \in \mathcal{K}} \text{Freq}(y, \mathcal{K}) \quad (3)$$

$$q_r^* = q_i, \text{ if } \frac{i-1}{M} \leq \frac{\text{Freq}(y_r^*, \mathcal{K})}{k} < \frac{i}{M} \quad (4)$$

We then construct the retrieval prompt by filling in the template T_{prompt} with the quantifier q_r^* and the retrieved answer y_r^* . We detail example templates and prompt variants we explored in Appendix A. The same pre-trained T5 model used for the question prompt is used to tokenize the retrieval prompt and obtain the retrieval embeddings $\mathbf{x}_r \in \mathbb{R}^{l_r \times d}$.

3.4 Generative Visual Question Answering

Encoder: Following prompt construction, we obtain a combination of embeddings $[\mathbf{v}_p; \mathbf{x}_q; \mathbf{x}_r]$ which is further fed as inputs to the transformer encoder layers of a pre-trained T5 model, and obtain contextualized representations of the combined

sequence from the top encoder layer, where we denote as $\mathbf{X} = \text{Encoder}([\mathbf{v}_p; \mathbf{x}_q; \mathbf{x}_r])$. In this work, we use a moderately sized model with around 60 million parameters, T5-small, and leave models with more parameters for future exploration.

Decoder: While most prior works use a discriminative architecture for medical VQA, we experiment with a decoder to predict free-form text. A transformer decoder from T5 is used to predict words in the vocabulary autoregressively. As each answer label y has a corresponding text string z of varying length, we formulate the likelihood of an answer string z given an image \mathbf{v} and a question \mathbf{x} by the following conditional probability:

$$P_{\text{gen}}(z|\mathbf{X}) = \prod_{j=0}^{|z|} P_{\phi}(z_j|\mathbf{X}, z_{<j}). \quad (5)$$

We finally optimize the generative model using a cross-entropy loss between the conditional probability $P_{\text{gen}}(z|\mathbf{X})$ and the ground-truth answer string z on the training dataset. This formulation allows for more flexible answers, which can easily change depending on the task, but may produce answers that are essentially the same with minor differences (e.g., extra whitespace, synonyms, etc.). To resolve these minor discrepancies, we utilize a simple string-matching heuristic that matches the longest continuous subsequence² between the generated answer and the closest possible label in the answer label set. Thus, our final generative model predicts answers as follows:

$$z^* = \arg \max_z P_{\text{gen}}(z|\mathbf{X}) \quad (6)$$

$$y^* = \text{LongestCommonString}(z^*, \mathcal{Y}). \quad (7)$$

Compared to the exact match between the generated answer string z^* and the ground-truth string z , we observe a 3-4% improvement in accuracy when using this heuristic on the VQA-RAD dataset and a 1% gain on the SLAKE dataset.

4 Experimental Setup

We perform our analysis on the VQA-RAD and SLAKE datasets, which are anonymous and preprocessed following prior works (Eslami et al., 2021; Zhan et al., 2020). We use an AdamW optimizer with an initial learning rate $1e^{-4}$ for T5 finetuning. We use a ViT-B/32 architecture for our CLIP models, and T5-small for answer generation. The

²<https://docs.python.org/3/library/difflib.html>

plateau learning rate scheduler is used to decay the learning rate by a factor of 10 if the validation loss does not decrease for 10 consecutive epochs.

All model training used a batch size of 16 and took 2-3 hours on average on a RTX 3090 GPU. All results were seeded with the best run of [Eslami et al. \(2021\)](#); [Zhan et al. \(2020\)](#) for reproducibility.³

4.1 Datasets

SLAKE The SLAKE dataset comprises 642 images and over 14,000 VQA pairs in English and Chinese. We only use the English portion to match the language of the T5 pretraining corpus. We use the provided train, validation and test splits, corresponding to 4918, 1053, and 1061 QA pairs. SLAKE consists of 10 different question types.⁴

VQA-RAD VQA-RAD is a high-quality dataset consisting of 315 patient scans and 3515 questions. We use the train and eval splits provided with the original data, following prior works ([Tanwani et al., 2022](#); [Eslami et al., 2021](#); [Nguyen et al., 2019](#)). VQA-RAD consists of 11 different question types.

Radiology Objects in Context (ROCO) The ROCO dataset ([Pelka et al., 2018](#)) has over 81,000 radiology image-caption pairs, making it a popular medical dataset for pretraining vision-language models ([Pelka et al., 2018](#)). Each image-caption pair also contains keywords and semantic types used by existing works for masked language modeling on salient spans ([Khare et al., 2021](#)).

Synthetic VQA Data Using the image-caption data from the ROCO dataset, we construct a large-scale synthetic VQA dataset consisting of over 50,000 question-answer pairs. Using our procedure (§3.1), we create question and keyword templates focusing on organ, modality, and plane questions.

4.2 Training Settings

Dataset Adaptation (DA): To evaluate the generalization of VQA models across datasets, we examine a setting where we train a model on a source-labeled dataset and use it to answer questions from a different target dataset with access to a target dataset. We further compare models using the target labeled examples for (a) in-context prediction without updating source-trained models or (b) continued fine-tuning.

³The performance was similar across 5 different seeds for our best model, with a standard deviation of about 1.5%.

⁴Details in Table 7 and 8 in the appendix.

In-domain Evaluation (IDE): In this setting, we adopt a standard split of each dataset into train/validation/test sets. We then train models on the train set, select the best checkpoints by the validation set, and evaluate models on the test set.

4.3 Baselines

Mixture of Enhanced Visual Features (MEVF) [Nguyen et al. \(2019\)](#) utilize model agnostic meta-learning (MAML) in conjunction with a convolutional denoising autoencoder (CDAE) to learn medical image latent feature representations.

Question Answering with Conditional Reasoning (QCR) [Zhan et al. \(2020\)](#) introduce novel task-conditioned, open, and closed reasoning modules to distinguish between answer types and improve open question accuracy.

PubMedCLIP [Eslami et al. \(2021\)](#) utilize the ROCO dataset to finetune a general CLIP model on medical image-caption pairs. They modify existing architectures with the finetuned vision encoder to achieve improved results.

MMBERT [Khare et al. \(2021\)](#) introduces a BERT-based method that utilizes pretraining on the ROCO dataset with a masked language modeling objective. The model predicts answers by performing an average pooling on the last layer features followed by a linear classification layer.

MPR_{disc} (Ours): MPR_{disc} refers to our discriminative variant by replacing a generative decoder with a prediction head to predict a finite set of answers. MPR_{disc_BAN} uses a prediction head similar to MPR_{disc}, but fuses the image and text features with a bilinear attention network ([Kim et al., 2018](#)).

MPR_{gen} (Ours): MPR_{gen} refers to our generative architecture which outputs flexible answers. MPR_{gen_PM} has the same architecture as MPR_{gen}, but is initialized with a pre-trained checkpoint from PubMedCLIP ([Eslami et al., 2021](#)).⁵

5 Results and Analysis

This section describes the results of our main experiments (§5.1) and fine-grained analysis thereafter.

5.1 In-context Prediction for Adaptation

First, we evaluate our proposed method’s generalization capability of in-context predictions. We

⁵<https://github.com/sarahESL/PubMedCLIP>

Context	Method	SLAKE \rightarrow VQA-RAD			VQA-RAD \rightarrow SLAKE		
		Open	Closed	Overall	Open	Closed	Overall
Image and Question	MPR _{gen_PM}	6.0	53.4	34.6	18.3	52.2	31.6
	MPR _{gen}	4.9	52.0	33.3	16.9	46.4	28.5
Image, Question, and Retrieval	MPR _{gen_PM}	42.9	76.2	63.0	45.1	67.3	53.8
	MPR _{gen}	41.8	74.4	61.4	38.4	57.7	46.0

Table 1: Performances of our generative prompting method in a domain adaptation setting with different levels of context. When provided with retrieval context, the models query for $k = 1$ relevant image-question pairs.




Figure 2: Domain adaptation accuracy for the top 5 most common question types in the SLAKE and VQA-RAD datasets. Percentages on the x-axis indicate each question type’s proportion of the dataset. MPR_{gen} and MPR_{gen_PM} use retrieval and query for $k = 1$ relevant image-question pairs. Retrieval-based models outperform our zero-shot baseline, and initializing MPR with a PubMedCLIP checkpoint helps.

define a k -shot setting where our model retrieves Top- k similar image-question pairs from the retrieval set. We compare the performance of the $k = 1$ setting with zero-shot MPR (i.e., MPR w/o retrieval) on two medical domain adaptation tasks.

Overall Accuracy: Table 1 compares the performances of our generative models under domain shift. Most notably, allowing the models to access a retrieval set universally improves performance, especially on questions with open answers. We also demonstrate that initializing our model with a PubMedCLIP pre-trained checkpoint results in higher accuracy than a general CLIP checkpoint. As the other discriminative baselines can only predict a fixed set of answers, they cannot perform adaptation over different answer sets. We only compare them for our in-domain analysis (§5.5).

Fine-grained Accuracy over QA Types: Figure 2 summarizes our model performances across individual QA types in a domain adaptation setting. We find that zero-shot MPR struggles with question types that require logical reasoning, such as Knowledge Graph (KG) or Position questions, while in-context retrieval increases model performance significantly in these question types. Using

a PubMedCLIP vision encoder further increases accuracy for these challenging question types.

5.2 Retrieval Sets for In-context Prediction

Source \rightarrow Target	Retrieval Set	Open	Closed	Overall
SLAKE \rightarrow VQA-RAD	None (Zero-shot)	6.0	53.4	34.6
	Synthetic	11.5	49.8	34.6
	VQA-RAD	42.9	76.2	63.0
	VQA-RAD + Synthetic	44.5	76.5	63.8
VQA-RAD \rightarrow SLAKE	None (Zero-shot)	16.9	46.4	28.5
	Synthetic	18.3	50.2	30.8
	SLAKE	45.1	67.3	53.8
	SLAKE + Synthetic	45.1	67.3	53.8

Table 2: Results of zero-/few-shot in-context prediction for domain adaptation with varying degrees of retrieval dataset access. We use MPR_{gen_PM} with $k = 1$ for all settings except $k = 50$ for the noisy synthetic dataset.

We also examine the effect of using different datasets for retrieval. Table 2 illustrates the zero-shot/few-shot accuracies when applying a source model to a target dataset with different retrieval datasets. Increasing the retrieval dataset’s quality improves the model’s adaptation capability to new questions. Without any retrieval, open question accuracy is as low as 6%. Providing access to a noisy synthetic retrieval dataset improves open question performance. Using a higher quality in-domain retrieval set further enhances performance in all




Figure 3: A generative model $\text{MPR}_{\text{gen_PM}}$ trained on the SLAKE dataset is evaluated on the VQA-RAD dataset over different numbers of few-shot retrievals.

categories, achieving over 30% improvement in open question accuracy compared to the zero-shot baselines. Combining in-domain retrieval data with noisy synthetic data further boosts accuracy in all three accuracy categories on VQA-RAD. However, we observed no further improvement when combining the synthetic and SLAKE datasets. With a manual investigation, we find that questions in SLAKE have much simpler synthetic variants than those in VQA-RAD. Therefore, SLAKE already provides the most similar examples during retrieval, and additional synthetic data provides minimal gains.

5.3 How Many Shots are Needed?

For adaptation at test time, we investigate the effect of varying the number of retrieved image-question pairs from the target dataset for constructing the retrieval prompts in Figure 3. Regardless of the number of pairs retrieved, the overall target accuracy of MPR_{gen} is always above the none-retrieval baseline (i.e., zero-shot). We hypothesize that accuracy peaks when $k = 1$ and stabilizes as k increases due to the small dataset size. MPR_{gen} outperforms a purely nearest neighbor-based approach when testing on the VQA-RAD dataset. However, on a syntactically simpler dataset (i.e., SLAKE), we also find that a nearest neighbor-based classifier can achieve higher accuracy than our model.

5.4 In-context Prediction vs Finetuning

While further finetuning neural models on the target dataset often successfully learns to adapt to the new distribution, this technique often results in catastrophic forgetting (Thompson et al., 2019). Figure 4 shows our experiments with further finetuning a source-trained model on a target dataset. First, we initialize three models with a $\text{MPR}_{\text{gen_PM}}$ checkpoint trained on SLAKE and adapt them to VQA-RAD. The first model is frozen, only using in-context prediction with retrieved target data (green). Another model is further finetuned on the target data without in-context prediction

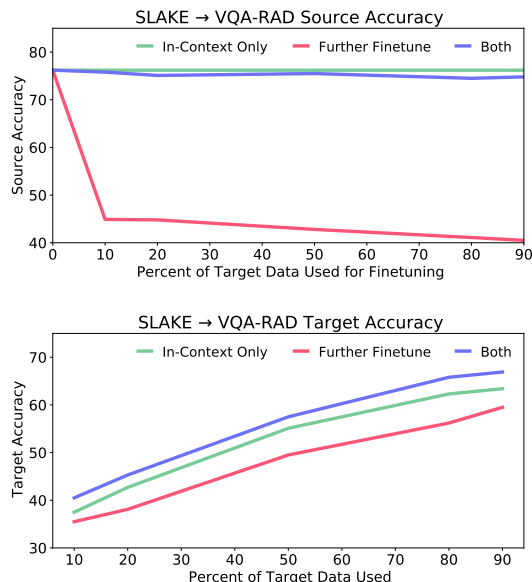


Figure 4: We vary the finetuning dataset size and observe the accuracy of finetuning compared to in-context prediction. All models start from a $\text{MPR}_{\text{gen_PM}}$ checkpoint and are either further finetuned (red + blue) or kept the same (green).

(red). The last model uses fine-tuning first and then does in-context predictions with retrieval (blue).

Several findings can be observed. First, we find that in-context prediction with $\text{MPR}_{\text{gen_PM}}$ can mitigate the forgetting issue and improve cross-dataset adaptation. Second, when target data is scarce, in-context prediction outperforms further finetuning. Although the finetuned model achieved higher test accuracy when using all the target data, it suffered significant performance loss in its original domain. Lastly, combining in-context prediction with further finetuning eliminates most of this forgetting with minimal target domain performance loss.

5.5 In-domain Evaluation

Overall Accuracy We also compare our proposed model with existing models for the in-domain setting on SLAKE and VQA-RAD. We highlight the overall, open, and closed test accuracy for each dataset. We also evaluate our method with three contexts to analyze the effect of each component of our prompting method in Table 3.

As expected, the model variants perform worse when we only provide questions as inputs. Under the same setting where both the question and image features are provided, our generative model is competitive with the state-of-the-art discriminative models. Besides, we also find that using an in-domain dataset for retrieval does not provide performance gains, indicating that models can easily fit a small in-domain dataset, and retrieving prompts from the same training set does not pro-

Context	Method	SLAKE			VQA-RAD		
		Open	Closed	Overall	Open	Closed	Overall
Question Only	MPR _{gen}	45.6	68.3	54.9	22.5	63.5	50.3
	MPR _{disc}	48.5	66.6	55.6	38.5	72.6	59.0
Image and Question	MPR _{gen}	71.5	76.7	73.5	57.7	77.6	69.7
	MPR _{gen_PM}	74.1	82.2	77.3	62.6	78.3	72.1
	MPR _{disc}	78.3	84.9	80.9	57.7	76.2	68.8
	MPR _{disc_BAN}	76.0	79.8	77.5	60.4	81.6	73.2
	PubMedCLIP (Eslami et al., 2021)	78.4	82.5	80.1	60.1	80.0	72.1
	MMBert (Khare et al., 2021)	-	-	-	63.1	77.9	72.0
	QCR (Zhan et al., 2020)	-	-	-	60.0	79.3	71.6
	MEVF (Nguyen et al., 2019)	-	-	-	43.9	75.1	62.7
Image, Question, and Retrieval	MPR _{gen}	73.0	79.8	75.7	57.7	77.3	69.5
	MPR _{gen_PM}	73.5	80.5	76.2	60.4	80.9	72.8
	MPR _{disc}	75.0	81.0	77.4	51.6	78.0	67.5
	MPR _{disc_BAN}	77.5	82.2	79.4	62.6	80.1	73.2

Table 3: Performances of our prompting method with different levels of context provided to the model. In most cases, our model is competitive with state-of-the-art methods even in the generative based cases. Retrieval context is provided by querying for the 15 most relevant image-question pairs. Bold values indicate the maximum in each column.

vide extra useful information.

Finegrained Accuracy Figure 5 in Appendix also shows the in-domain performance of our model variants across different question types for both datasets. The results indicate that all models generally struggle with questions requiring more complex reasoning, such as Position, Abnormality, and Knowledge Graph (KG) questions.

6 Related Work

Retrieval-Based VQA Retrieval-based methods typically combine parametric models with non-parametric external memory for prediction. This idea first surfaces in KNN-LMs (Khandelwal et al., 2020), which utilizes a static retrieval data store to help language models adapt rapidly to new domains without further training. Guu et al. (2020) extends this idea by introducing a parametric retriever that learns to attend to relevant documents during training. Recently, Gao et al. (2022) summarizes visual information into natural language to use as a query for dense passage retrieval. The retrieved passages allow for the VQA model to outperform existing works, especially on questions which require outside knowledge. Lin and Byrne (2022) consider training the retriever in an end-to-end manner similar to Lewis et al. (2020) and find that this results in higher answer quality and lower computational training cost.

Different from these methods, we propose to construct multimodal prompts from retrieval to perform zeroshot dataset adaptation. While dataset adaptation of VQA models has been investigated in Agrawal et al. (2023), we focus on the effect of retrieval on generalization capability.

Generative QA Generative QA models focus on predicting answers autoregressively based on the input question. In this setting, the model may either generate the response based solely on model parameters (closed book) (Khashabi et al., 2021; Roberts et al., 2020) or rely on additional retrieved contexts (open book) (Karpukhin et al., 2020; Lewis et al., 2020). Our prompt construction method is inspired by the retrieval augmented generator (RAG) model (Lewis et al., 2020), which retrieves relevant documents to answer questions. Instead of retrieving documents exclusively, we identify suitable image-question pairs to perform VQA.

VQA First introduced by Antol et al. (2015), most VQA systems learn a joint embedding space for images and text to answer questions (Malinowski et al., 2015; Gao et al., 2015). These approaches combine image and text features through either bilinear pooling or attention-based mechanisms (Yang et al., 2016; Lu et al., 2016; Anderson et al., 2018; Guo et al., 2021). To help models understand the relationships between objects in an image, graph convolutional neural networks were introduced for VQA (Norcliffe-Brown et al., 2018; Li et al., 2019). Current methods often combine supplemental knowledge with fusion-based approaches to achieve state-of-the-art performance (Shevchenko et al., 2021; Marino et al., 2021; Wu et al., 2022; Chen et al., 2022). We take a similar approach by using supplementary knowledge to construct context-aware prompts.

7 Conclusion

In this work, we propose a flexible prompt-based method for VQA in the medical domain. While our approach is designed for low-resource domains,

the generative architecture of our model, in combination with a retrieval component, enables generalization to other fields. Our results are on par with state-of-the-art accuracies on the SLAKE and VQA-RAD datasets and show promising zero-shot and few-shot transfer results across different medical datasets. We hope these results can offer a baseline to compare with future work on knowledge-intensive and reasoning tasks.

8 Limitations

When evaluating our model in a cross-dataset adaptation setting, our experiments indicate the importance of using a retrieval dataset. It is challenging to procure high-quality and volume retrieval datasets, especially in low-resource domains such as the medical field. Fortunately, the VQA-RAD and SLAKE datasets we evaluate on contain professionally annotated medical images. We also overcome the lack of data by creating a synthetic dataset from the medical ROCO image-captioning dataset.

Additionally, our model struggles with questions requiring multi-step reasoning, such as knowledge graph, abnormality, and position questions. Although performances in these question types are not far below the overall accuracy, future work may consider supplementary knowledge-based retrieval to assist in these challenging question types.

9 Ethics Statement

Although medical VQA provides exciting opportunities for future AI-assisted clinical diagnostic tools, there are several ethical challenges associated with these approaches.

Patient Safety and Model Transparency Since the model decision process for deep learning models is difficult to understand, these models should only be used as an auxiliary tool. This obscure decision process is crucial to clarify in the medical domain, in which a poor diagnosis or choice of treatment can significantly affect patient lives. For example, medical experts found that cancer treatment recommendation software often gave unsafe or incorrect treatment advice in a recent study (Ross and Swetlitz, 2018).

Dataset Biases The fairness of medical AI systems depends on the distribution of people in its training dataset. To ensure that AI algorithms display fairness to all races, genders, and ethnic

groups, practitioners should verify that the training dataset contains an equal representation of all groups. Before deploying our architecture or other deep learning-based models to a clinical setting, practitioners should ensure that their patient's background is adequately represented in the training data.

References

- Aishwarya Agrawal, Ivana Kajić, Emanuele Bugliarello, Elnaz Davoodi, Anita Gergely, Phil Blunsom, and Aida Nematzadeh. 2023. Rethinking evaluation practices in visual question answering: A case study on out-of-distribution generalization. In *Proceedings of the 2023 European Chapter of the Association for Computational Linguistics (Findings)*.
- Peter Anderson, Xiaodong He, Chris Buehler, Damien Teney, Mark Johnson, Stephen Gould, and Lei Zhang. 2018. Bottom-up and top-down attention for image captioning and visual question answering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6077–6086.
- Stanislaw Antol, Aishwarya Agrawal, Jiasen Lu, Margaret Mitchell, Dhruv Batra, C Lawrence Zitnick, and Devi Parikh. 2015. Vqa: Visual question answering. In *Proceedings of the IEEE international conference on computer vision*, pages 2425–2433.
- 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. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*.
- Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. 2021. An image is worth 16x16 words: Transformers for image recognition at scale. In *International Conference on Learning Representations*.
- Sedigheh Eslami, Gerard de Melo, and Christoph Meinel. 2021. Does clip benefit visual question answering in the medical domain as much as it does in the general domain? *arXiv preprint arXiv:2112.13906*.
- Feng Gao, Qing Ping, Govind Thattai, Aishwarya Reganti, Ying Nian Wu, and Prem Natarajan. 2022. Transform-retrieve-generate: Natural language-centric outside-knowledge visual question answering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5067–5077.
- Haoyuan Gao, Junhua Mao, Jie Zhou, Zhiheng Huang, Lei Wang, and Wei Xu. 2015. Are you talking to a machine? dataset and methods for multilingual image

- question. *Advances in neural information processing systems*, 28.
- Dalu Guo, Chang Xu, and Dacheng Tao. 2021. Bilinear graph networks for visual question answering. *IEEE Transactions on Neural Networks and Learning Systems*.
- Kelvin Guu, Kenton Lee, Zora Tung, Panupong Pasupat, and Mingwei Chang. 2020. Retrieval augmented language model pre-training. In *International Conference on Machine Learning*, pages 3929–3938. PMLR.
- Sadid A Hasan, Yuan Ling, Oladimeji Farri, Joey Liu, Henning Müller, and Matthew P Lungren. 2018. Overview of imageclef 2018 medical domain visual question answering task. In *Conference and Labs of the Evaluation Forum (Working Notes)*.
- Vladimir Karpukhin, Barlas Oğuz, Sewon Min, Patrick Lewis, Ledell Wu, Sergey Edunov, Danqi Chen, and Wen-tau Yih. 2020. Dense passage retrieval for open-domain question answering. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing*.
- Urvashi Khandelwal, Omer Levy, Dan Jurafsky, Luke Zettlemoyer, and Mike Lewis. 2020. Generalization through memorization: Nearest neighbor language models. In *International Conference on Learning Representations*.
- Yash Khare, Viraj Bagal, Minesh Mathew, Adithi Devi, U Deva Priyakumar, and CV Jawahar. 2021. Mmbert: multimodal bert pretraining for improved medical vqa. In *2021 IEEE 18th International Symposium on Biomedical Imaging (ISBI)*, pages 1033–1036. IEEE.
- Daniel Khashabi, Sewon Min, Tushar Khot, Ashish Sabharwal, Oyvind Tafjord, Peter Clark, and Hannaneh Hajishirzi. 2021. Unifiedqa: Crossing format boundaries with a single qa system. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing (Findings)*.
- Jin-Hwa Kim, Jaehyun Jun, and Byoung-Tak Zhang. 2018. Bilinear attention networks. *Advances in neural information processing systems*, 31.
- Olga Kovaleva, Chaitanya Shivade, Satyananda Kashyap, Karina Kanjaria, Joy Wu, Deddeh Ballah, Adam Coy, Alexandros Karargyris, Yufan Guo, David Beymer Beymer, Anna Rumshisky, and Vandana Mukherjee Mukherjee. 2020. [Towards visual dialog for radiology](#). In *Proceedings of the 19th SIGBioMed Workshop on Biomedical Language Processing*, pages 60–69, Online. Association for Computational Linguistics.
- Jason J Lau, Soumya Gayen, Asma Ben Abacha, and Dina Demner-Fushman. 2018. A dataset of clinically generated visual questions and answers about radiology images. *Scientific data*, 5(1):1–10.
- Brian Lester, Rami Al-Rfou, and Noah Constant. 2021. [The power of scale for parameter-efficient prompt tuning](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 3045–3059, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, et al. 2020. Retrieval-augmented generation for knowledge-intensive nlp tasks. *Advances in Neural Information Processing Systems*, 33:9459–9474.
- Linjie Li, Zhe Gan, Yu Cheng, and Jingjing Liu. 2019. Relation-aware graph attention network for visual question answering. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 10313–10322.
- Weizhe Lin and Bill Byrne. 2022. Retrieval augmented visual question answering with outside knowledge. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*.
- Zhihong Lin, Donghao Zhang, Qingyi Tao, Danli Shi, Gholamreza Haffari, Qi Wu, Mingguang He, and Zongyuan Ge. 2021. Medical visual question answering: A survey. *ArXiv*, abs/2111.10056.
- Bo Liu, Li-Ming Zhan, Li Xu, Lin Ma, Yan Yang, and Xiao-Ming Wu. 2021. Slake: a semantically-labeled knowledge-enhanced dataset for medical visual question answering. In *2021 IEEE 18th International Symposium on Biomedical Imaging (ISBI)*, pages 1650–1654. IEEE.
- Jiasen Lu, Jianwei Yang, Dhruv Batra, and Devi Parikh. 2016. Hierarchical question-image co-attention for visual question answering. *Advances in neural information processing systems*, 29.
- Mateusz Malinowski, Marcus Rohrbach, and Mario Fritz. 2015. Ask your neurons: A neural-based approach to answering questions about images. In *Proceedings of the IEEE international conference on computer vision*, pages 1–9.
- Kenneth Marino, Xinlei Chen, Devi Parikh, Abhinav Gupta, and Marcus Rohrbach. 2021. Krisp: Integrating implicit and symbolic knowledge for open-domain knowledge-based vqa. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14111–14121.
- Jong Hak Moon, Hyungyung Lee, Woncheol Shin, Young-Hak Kim, and Edward Choi. 2022. Multi-modal understanding and generation for medical images and text via vision-language pre-training. *IEEE Journal of Biomedical and Health Informatics*.
- Medhini Narasimhan and Alexander G Schwing. 2018. Straight to the facts: Learning knowledge base retrieval for factual visual question answering. In *Proceedings of the European conference on computer vision (ECCV)*, pages 451–468.

- Binh D Nguyen, Thanh-Toan Do, Binh X Nguyen, Tuong Do, Erman Tjiputra, and Quang D Tran. 2019. Overcoming data limitation in medical visual question answering. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 522–530. Springer.
- Will Norcliffe-Brown, Stathis Vafeias, and Sarah Parisot. 2018. Learning conditioned graph structures for interpretable visual question answering. *Advances in neural information processing systems*, 31.
- Obioma Pelka, Sven Koitka, Johannes Rückert, Felix Nensa, and Christoph M Friedrich. 2018. Radiology objects in context (roco): a multimodal image dataset. In *Intravascular Imaging and Computer Assisted Stenting and Large-Scale Annotation of Biomedical Data and Expert Label Synthesis*, pages 180–189. Springer.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. 2021. [Learning transferable visual models from natural language supervision](#). In *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 8748–8763. PMLR.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, Peter J Liu, et al. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of Machine Learning Research*, 21(140):1–67.
- Adam Roberts, Colin Raffel, and Noam Shazeer. 2020. How much knowledge can you pack into the parameters of a language model? In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing*.
- Casey Ross and Ike Swetlitz. 2018. Ibm’s watson super-computer recommended ‘unsafe and incorrect’ cancer treatments, internal documents show. *Stat*, 25.
- Violetta Shevchenko, Damien Teney, Anthony Dick, and Anton van den Hengel. 2021. Reasoning over vision and language: Exploring the benefits of supplemental knowledge. *arXiv preprint arXiv:2101.06013*.
- Anshumali Shrivastava and Ping Li. 2014. [Asymmetric lsh \(alsh\) for sublinear time maximum inner product search \(mips\)](#). In *Advances in Neural Information Processing Systems*, volume 27. Curran Associates, Inc.
- Ajay K Tanwani, Joelle Barral, and Daniel Freedman. 2022. Reptsnet: Combining vision with language for automated medical reports. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 714–724. Springer.
- Brian Thompson, Jeremy Gwinnup, Huda Khayrallah, Kevin Duh, and Philipp Koehn. 2019. Overcoming catastrophic forgetting during domain adaptation of neural machine translation. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 2062–2068.
- Jialin Wu, Jiasen Lu, Ashish Sabharwal, and Roozbeh Mottaghi. 2022. Multi-modal answer validation for knowledge-based vqa. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 2712–2721.
- Zichao Yang, Xiaodong He, Jianfeng Gao, Li Deng, and Alex Smola. 2016. Stacked attention networks for image question answering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 21–29.
- Li-Ming Zhan, Bo Liu, Lu Fan, Jiabin Chen, and Xiao-Ming Wu. 2020. Medical visual question answering via conditional reasoning. In *Proceedings of the 28th ACM International Conference on Multimedia*, pages 2345–2354.
- Yangyang Zhou, Xin Kang, and Fuji Ren. 2019. Tual at imageclef 2019 vqa-med: a classification and generation model based on transfer learning. In *Conference and Labs of the Evaluation Forum (Working Notes)*.

Appendix

Q Type	A Type	Question Templates & Answer Templates (\mathcal{T}_t)	Answer Templates
Organ	Open	Q: What part of the body is being imaged? What is the organ shown in this image? ...	A: {Brain, Chest, ...}
Organ	Closed	Q: Does the picture contain {}? Is this a study of the {}? ...	A: {Yes, No}
Organ System	Open	Q: What organ system is pictured? What system is this pathology in? ...	A: {Respiratory System, Cardiovascular System, ...}
Organ System	Closed	Q: Is this an image of the {}? Is the {} shown? ...	A: {Yes, No}
Modality	Open	Q: What kind of scan is this? How was this image taken? ...	A: {MRI, X-ray, ...}
Modality	Closed	Q: Is this a {}? Is the image a {}? ...	A: {Yes, No}
Plane	Open	Q: What image plane is this? How is the image oriented? ...	A: {Axial, Coronal, ...}
Plane	Closed	Q: Is this a {} plane? Is the image a {} section? ...	A: {Yes, No}

Table 4: Example templates for different question categories and question types.

Q&A Types ($t \in \mathcal{Q}_{\text{type}} \times \mathcal{A}_{\text{type}}$)	Keywords (\mathcal{W}_t)
Organ & Open	Heart, Lungs, Lung, Liver, Breasts, Chest, Cardiovascular System, Respiratory System ...
Plane & Open	Axial, Coronal, Supratentorial, Posteroanterior ...
Modality & Open	MRI, T1, T2, CT, X-ray, Ultrasound, Flair ...

Table 5: Example keywords for different question types.

A Templates

We use the question and keyword templates in Tables 4 and 5 to construct a synthetic retrieval set. For open questions, the question template is static. However, the answer to open questions may be any of the keywords $w \in \mathcal{W}_t$. Closed question templates have a slot that is filled in by one of the keywords $w \in \mathcal{W}_t$, and the answer to these questions is either yes or no:

Prompt Construction Order	Prompt Template (T_{prompt})	Possible Quantifiers (\mathcal{Q})	Open	Closed	Overall
Question, Retrieval, Image	I believe the answer is {quantifier} {answer}	very unlikely, unlikely, maybe, likely, very likely, certainly	39.6	65.0	54.9
Image, Retrieval, Question	I believe the answer is {quantifier} {answer}	very unlikely, unlikely, maybe, likely, very likely, certainly	39.0	65.3	54.9
Image, Question, Retrieval	{answer} is {quantifier} the answer	very unlikely, unlikely, maybe, likely, very likely, certainly	37.9	68.6	56.4
Image, Question, Retrieval	I believe the answer is {quantifier} {answer}	very unlikely, unlikely, maybe, likely, very likely, certainly	39.6	65.3	55.1

Table 6: Using different variants of our prompt results in similar performances. This table shows the open, closed, and overall accuracy of $\text{MPR}_{\text{gen_PM}}$ in a domain adaptation setting from SLAKE to VQA-RAD, retrieving k=1 nearest question-image pairs.

B Prompt Variations

During the prompt construction process, we experiment with different prompt ordering and wording of retrieval prompts. We illustrate different template wording in the Prompt Template column of Table 6. Each prompt contains a quantifier that is filled in with an expression ranging from “very unlikely” to “certainly” based on the confidence score of y^* , described in Section 3.2. We found that performance does not significantly change when changing these aspects of the prompt, and we ultimately decided to use settings in the last row of the table.

C Dataset Information

Tables 7 and 8 report descriptive statistics about the datasets used in our experiments. Although the synthetic data contains more question-answer pairs than SLAKE and VQA-RAD, it has noisier labels and more limited question types. SLAKE and VQA-RAD have larger question-answer diversity and share several question types, such as Organ, Position, and Abnormality questions.

Figure 5 displays our model’s in-domain accuracies on SLAKE and VQA-RAD. The models perform best on Color, Attribute, and Size questions in VQA-RAD. The discriminative variants have better accuracy overall, but can not be directly applied to other datasets. We observed lower accuracy on Modality and Organ questions in VQA-RAD, which we attribute to the diversity of question and answer phrasing in these VQA-RAD question types.

Dataset	Train Split	Validation Split	Test Split	Number of Question Types
SLAKE	4918	1053	1061	10
VQA-RAD	3064	-	451	11
ROCO (image-caption only)	65460	8183	8182	-
Synthetic	56526	-	-	3

Table 7: Statistics about the datasets in our experiments. The synthetic data contains three different question types because other types require domain-specific knowledge (e.g. abnormality, knowledge graph, etc).

SLAKE		VQA-RAD		Synthetic	
Question Type	Percentage of Test Data	Question Type	Percentage of Test Data	Question Type	Percentage of Data
Shape	0.66	Color	0.87	Modality	30.61
Color	3.20	Quantity/Counting	1.31	Plane	30.65
Quantity	4.90	Organ	2.18	Organ	38.74
Plane	5.47	Attribute	4.36	-	-
Size	6.13	Other	5.66	-	-
Modality	10.18	Plane	5.66	-	-
Knowledge Graph	13.95	Modality	7.19	-	-
Abnormality	14.14	Size	10.02	-	-
Position	17.53	Abnormality	12.20	-	-
Organ	23.85	Position	13.29	-	-
-	-	Presence	37.25	-	-

Table 8: Composition of the datasets in our experiments. We use the English portion of the SLAKE dataset.




Figure 5: Accuracy for each question type for the SLAKE and VQA-RAD test datasets. Percentages on the x-axis indicate each question type’s proportion of the dataset.

D Attention Visualization

Transformer-based models utilize attention to calculate dependencies between inputs which may be important for prediction. Since our MPR model uses a pretrained T5 encoder to combine features from several sources, a visualization of its attention scores may indicate which parts of the image contribute to

its answers. Figure 6 illustrates the encoder self-attention in different attention heads and layers of our model when asked a challenging position question. The results suggest that some attention modules may attend to the entire input image (Layer 1, Head 4), whereas others may look for local dependencies by attending to adjacent tokens (Layer 2, Head 7).

In addition to encoder self-attention, cross attention in encoder-decoder transformer architectures may also illustrate which tokens from the input prompt contribute the most when generating answer tokens. Since the prompt to our model consists of image tokens, we visualize which image regions have the highest attention scores in 8.

Existing work has shown the effectiveness of using retrieval from a data store to rapidly adapt language models to new domains (Khandelwal et al., 2020). KNN LMs uses a blending parameter $\lambda \in [0, 1]$ to control the influence of retrieved information towards prediction:

$$\lambda \text{pkNN}(y|x) + (1 - \lambda) \text{pLM}(y|x) \quad (8)$$

This method assumes the availability of a language model $\text{pLM}(y|x)$ and a retrieval model $\text{pkNN}(y|x)$ which can predict the next vocabulary token y given context x . However, given our retrieval set which consists of variable length answers and image data, it is difficult to estimate $\text{pkNN}(y|x)$ directly from our data store. Consequently, we augment our model input with retrieval prompts to allow for the implicit learning of retrieval reliance in an end-to-end manner. Figure 7 shows the average cross attention scores to the retrieval portion of the prompt when evaluating on test data. The results demonstrate that when the model prediction matches the retrieved answer, the attention scores to the corresponding prompt section are significantly higher. Based on this observation, we believe the model has learned how to weigh the retrieved information through end-to-end training.




Figure 6: Selected encoder self-attention visualizations across different encoder layers and attention heads. ITK represents an image-token from the input image.




Figure 7: Average attention score to the retrieval prompt for each attention module in our T5 encoder. When the model predicts the retrieved answer, attention scores to the retrieved information is significantly higher.

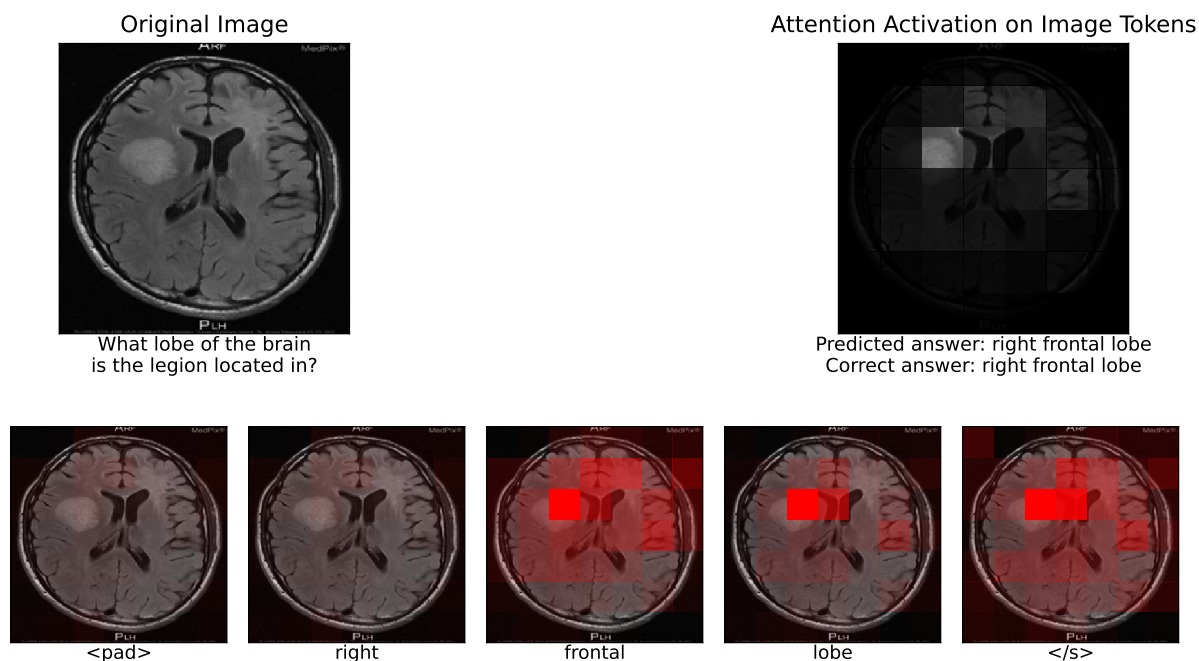


Figure 8: *Top*: The original image and question input to the model, followed by the average attention scores for each image patch, with darker patches corresponding to lower scores. *Bottom*: When predicting each word in the answer span, which input image regions are attended to.

ACL 2023 Responsible NLP Checklist

A For every submission:

- A1. Did you describe the limitations of your work?
Section 8
- A2. Did you discuss any potential risks of your work?
Section 9
- A3. Do the abstract and introduction summarize the paper’s main claims?
Section 1
- A4. Have you used AI writing assistants when working on this paper?
Left blank.

B Did you use or create scientific artifacts?

Section 3

- B1. Did you cite the creators of artifacts you used?
Section 3, Section 4.3
- B2. Did you discuss the license or terms for use and / or distribution of any artifacts?
All datasets and artifact models used in this study are publicly available and we provide links/citations to them.
- B3. Did you discuss if your use of existing artifact(s) was consistent with their intended use, provided that it was specified? For the artifacts you create, do you specify intended use and whether that is compatible with the original access conditions (in particular, derivatives of data accessed for research purposes should not be used outside of research contexts)?
Section 3
- B4. Did you discuss the steps taken to check whether the data that was collected / used contains any information that names or uniquely identifies individual people or offensive content, and the steps taken to protect / anonymize it?
Section 4
- B5. Did you provide documentation of the artifacts, e.g., coverage of domains, languages, and linguistic phenomena, demographic groups represented, etc.?
Section 3, Section 4
- B6. Did you report relevant statistics like the number of examples, details of train / test / dev splits, etc. for the data that you used / created? Even for commonly-used benchmark datasets, include the number of examples in train / validation / test splits, as these provide necessary context for a reader to understand experimental results. For example, small differences in accuracy on large test sets may be significant, while on small test sets they may not be.
Section 4, Appendix C

C Did you run computational experiments?

Section 5

- C1. Did you report the number of parameters in the models used, the total computational budget (e.g., GPU hours), and computing infrastructure used?
Section 3.4, Section 4

The Responsible NLP Checklist used at ACL 2023 is adopted from NAACL 2022, with the addition of a question on AI writing assistance.

- C2. Did you discuss the experimental setup, including hyperparameter search and best-found hyperparameter values?

Section 4

- C3. Did you report descriptive statistics about your results (e.g., error bars around results, summary statistics from sets of experiments), and is it transparent whether you are reporting the max, mean, etc. or just a single run?

Section 4, Section 5

- C4. If you used existing packages (e.g., for preprocessing, for normalization, or for evaluation), did you report the implementation, model, and parameter settings used (e.g., NLTK, Spacy, ROUGE, etc.)?

Section 4. Code and instructions to reproduce our results are provided in our GitHub repository.

D Did you use human annotators (e.g., crowdworkers) or research with human participants?

Left blank.

- D1. Did you report the full text of instructions given to participants, including e.g., screenshots, disclaimers of any risks to participants or annotators, etc.?

No response.

- D2. Did you report information about how you recruited (e.g., crowdsourcing platform, students) and paid participants, and discuss if such payment is adequate given the participants' demographic (e.g., country of residence)?

No response.

- D3. Did you discuss whether and how consent was obtained from people whose data you're using/curating? For example, if you collected data via crowdsourcing, did your instructions to crowdworkers explain how the data would be used?

No response.

- D4. Was the data collection protocol approved (or determined exempt) by an ethics review board?

No response.

- D5. Did you report the basic demographic and geographic characteristics of the annotator population that is the source of the data?

No response.