

Retrieval-Augmented Multi-Modal Chain-of-Thoughts Reasoning for Large Language Models

Anonymous ACL submission

Abstract

The advancement of Large Language Models (LLMs) has brought substantial attention to the Chain of Thought (CoT) approach, primarily due to its ability to enhance the capability of LLMs on complex reasoning tasks. Moreover, the significance of CoT approaches extends to the application of LLMs for multi-modal tasks. However, the selection of optimal CoT demonstration examples in multi-modal reasoning remains less explored for LLMs due to the inherent complexity of multi-modal examples. In this paper, we introduce a novel approach that addresses this challenge by using retrieval mechanisms to dynamically and automatically select demonstration examples based on cross-modal and intra-modal similarities. Furthermore, we employ a Stratified Sampling method of categorising demonstration examples into groups based on their types and then retrieving examples from different groups respectively to promote the diversity of demonstration examples. Through a series of experiments on two popular benchmark datasets: ScienceQA and MathVista, we demonstrate that our approach significantly improves the performance of GPT-4 by 6% on ScienceQA and 12.9% on MathVista, and enhances the performance of GPT-4V on two datasets by 2.7%, substantially improving the performance of the most advanced LLMs and LMMs for complex multi-modal reasoning tasks.

1 Introduction

The field of Natural Language Processing (NLP) has experienced significant advancements due to the emergence of Large Language Models (LLMs), which have reshaped the landscape of many tasks for their extensive capabilities. A key technique that has contributed greatly to the success of LLMs is the Chain of Thought (CoT) technique, as documented in prior studies (Wei et al., 2022a; Kojima et al., 2022). This technique becomes particularly crucial, especially when applied to multi-modal

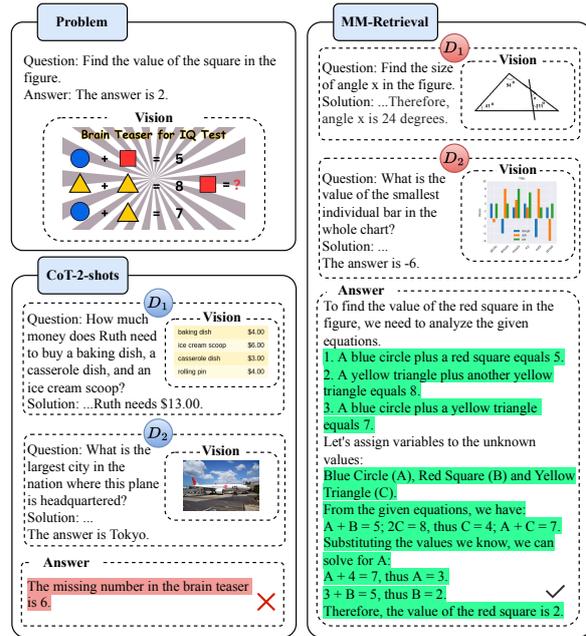


Figure 1: Our MM-Retrieval approach dynamically retrieves demonstrations based on the question. Compared with CoT, it has better adaptability and can stimulate the reasoning ability of LLMs. The red D_1 , D_2 represent demonstrations retrieved based on the question, while the blue D_1 , D_2 represent the fixed samples regardless of the question.

tasks. One of the most prominent applications is multi-modal question answering, which involves reasoning with both text and images (Zhang et al., 2023c; Lu et al., 2023b; Lyu et al., 2023; Li et al., 2023e). However, as researchers delve further into the integration of CoT with LLMs (Wang et al., 2022; Zhou et al., 2022; Zhang et al., 2022), the selection of appropriate demonstration examples to guide multi-modal reasoning emerges as a recurring challenge. Given that multi-modal examples often combine the complexities of both text and visual data, identifying the most relevant and informative examples is a non-trivial task (Bar et al., 2022; Li et al., 2023b,a).

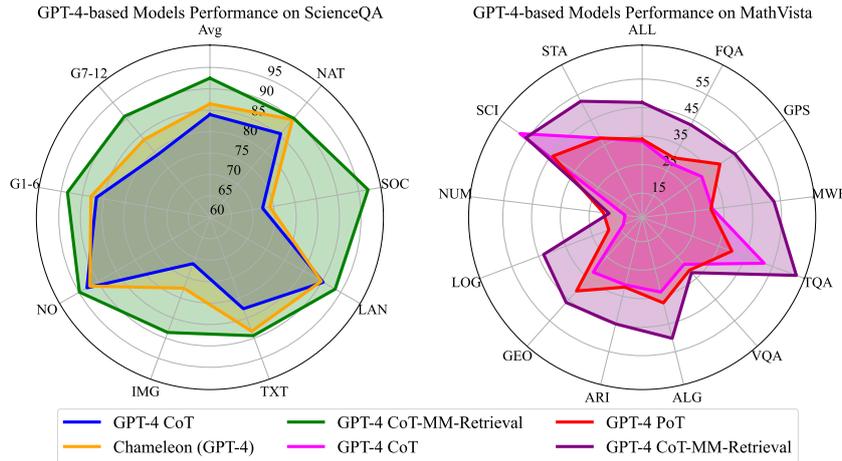


Figure 2: Results on different categories of ScienceQA (Lu et al., 2022) and MathVista (Lu et al., 2023a). Our proposed approach obtains substantial improvements over previous baseline models including *CoT* (Lu et al., 2023b), *PoT* (Lu et al., 2023a) and *Chameleon* (Lu et al., 2023b) on GPT-4 foundation models.

To address this issue, our paper presents a new method that utilizes retrieval mechanisms to select demonstration examples both dynamically and automatically. As shown in Figure 1, relevant and informative demonstration examples retrieved by our approach can elicit the reasoning ability of LLMs and lead to the correct answer. An overview of our proposed approach is shown in Figure 3. Our approach primarily focuses on utilizing cross-modal similarity and intra-modal similarity. We leverage these similarities to retrieve demonstration examples with the aim of enhancing the CoT reasoning process with more relevant demonstration examples on multi-modal tasks (Zhang et al., 2023a; Sun et al., 2023). To further ensure a more comprehensive and diverse selection, we innovatively utilize Stratified Sampling (Liberty et al., 2016). This simple but effective method involves sampling in sequence according to the group to which the retrieval sample belongs. By sampling from different groups, we aim to provide LLMs with a diverse range of examples, thus improving the overall quality of multi-modal reasoning.

To evaluate the effectiveness of our proposed approach, we performed extensive experiments on two benchmark multi-modal QA datasets, ScienceQA (Lu et al., 2022) and MathVista (Lu et al., 2023a). These experiments demonstrate that our approach substantially enhances the performance of LLMs, establishing a new state-of-the-art in multi-modal reasoning tasks. On the ScienceQA and MathVista dataset, our approach has shown substantial improvements, as illustrated in Figure 2. For the ScienceQA dataset, the ChatGPT-based

and GPT-4-based retrieval methods outperform the state-of-the-art Chameleon by 4.8% and 4.4% respectively. With more demonstrations, the optimal performance of ChatGPT-based and GPT-4-based methods can reach 86.4% and 92.5%. For GPT-4V, our method can achieve an average accuracy improvement of 2.7% compared to the zero-shot setting. Furthermore, our approach also demonstrates superior performance on MathVista dataset. Our approach based on ChatGPT and GPT-4 obtains substantial improvements of 8.4% and 13.6% respectively. Moreover, our approach can further boost the performance of the most advanced LLM - GPT-4V, improving its overall accuracy on MathVista by 2.7%, which demonstrates the effectiveness of our approach.

We also conducted comprehensive experiments on the contributions of each part of our method, including the visual information, retrieval mechanism, and Stratified Sampling. Additionally, we conducted detailed analysis to study the influence of varying the number of demonstration examples, providing valuable insights of how our proposed approach works with LLMs for multi-modal tasks.

2 Related Work

2.1 Retrieval-Augmented Generation for LLMs

Retrieval-augmented generation (RAG) for LLMs represents an important advancement in enhancing the generative capabilities of models by integrating external knowledge sources. Early works such as REALM (Guu et al., 2020) and RAG (Lewis

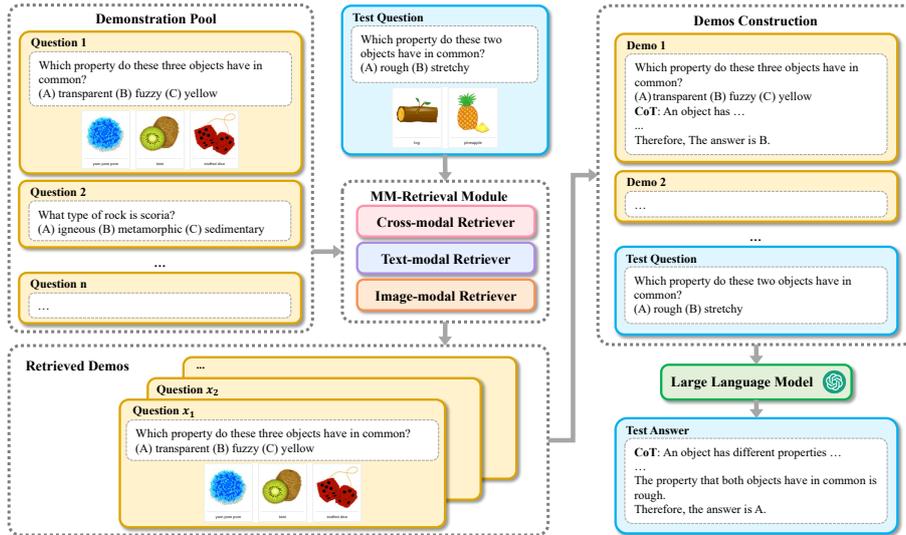


Figure 3: An overview of our proposed multi-modal retrieval method. We employ both cross-modality retrieval and intra-modality retrieval (text-modal and image-modal retrieval), to obtain relevant examples as retrieved demonstrations from demonstration pool. Then, these retrieved demonstrations are integrated with prompt and test question, serving as the input for LLMs.

et al., 2020) introduced the foundational methodology for incorporating external documents into the generation process. Subsequent research expanded the retrieval-augmented paradigm to multi-modal contexts. Such as MuRAG (Chen et al., 2022) and REVEAL (Hu et al., 2022), which augment language generation with both textual and visual information from external sources. Furthermore, recent studies such as FiD-Light (Hofstätter et al., 2022) and REPLUG (Shi et al., 2023) have focused on improving the efficiency and effectiveness of retrieval-augmented systems as well as exploring in-context retrieval-augmented mechanisms (Ram et al., 2023; de Jong et al., 2023).

2.2 In-Context Learning

In-Context Learning (ICL) uses LLMs for tasks based on a handful of examples embedded within the context (Devlin et al., 2019; Radford et al., 2019; Brown et al., 2020; Chowdhery et al., 2022), showing effectiveness in NLP and complex mathematical reasoning (Wei et al., 2022b). The setup for ICL involves generating responses from LLMs using context provided by task guidance and demonstration examples. It’s sensitive to prompt structure, example selection, and example order (Zhao et al., 2021; Wang et al., 2023a; Fu et al., 2022). The application of ICL extended to multi-modal tasks, including image segmentation, synthesis, and text-to-speech synthesis (Bar et al., 2022; Wang et al., 2023b,c,e; Tsimpoukelli et al., 2021; Alayrac

et al., 2022), with potential in structured spaces like graphs (Huang et al., 2023).

2.3 Chain-of-Thought Reasoning

Chain-of-Thought (CoT) reasoning guides LLMs to step-by-step reason, improving performance across arithmetic, symbolic, and logic tasks (Wei et al., 2022b; Kojima et al., 2022). Approaches included sampling multiple reasoning paths (Wang et al., 2022), dividing complex problems into sub-problems (Zhou et al., 2022), and dynamically selecting examples for few-shot prompting (Zhang et al., 2022; Shi et al., 2022). In addition to textual data, CoT has been adapted for tabular data (Ziqi and Lu, 2023) as well. Furthermore, its potential has been explored in multi-modal settings, demonstrating enhanced reasoning capabilities through the fusion of language and vision (Zhang et al., 2023c; Lu et al., 2023b,c). Studies like (Zhang et al., 2023c) proposed a two-stage CoT framework for multi-modal tasks, significantly improving reasoning accuracies on benchmarks like ScienceQA. Chameleon (Lu et al., 2023b) introduced plug-and-play modules for Large Multi-modal Models (LMMs), enabling complex reasoning by combining different tools.

3 Methodology

Our methodology is anchored in the CoT in-context learning paradigm, which effectively leverages LLMs. For every input query, we aim to harness

a set of relevant CoT demonstration examples to booster the LLM’s reasoning abilities. To achieve this, we introduce a novel approach that employs retrieval mechanisms for the dynamic and automatic selection of demonstration examples, as well as incorporating visual knowledge into the prompt. A detailed illustration of our approach is shown in Figure 4. Central to our methodology is the extraction of cross-modal similarity and intra-modal similarity, crosswise between textual q^t and visual context q^v of the test question q and examples in demonstration pool $Q = \{q_1, \dots, q_n\}$. Another distinctive feature of our methodology is the incorporation of Stratified Sampling. By categorizing demonstration examples into distinguishable groups according to their inherent attributes, we aim to expand the diversity of selected examples. Retrieving from diverse groups ensures that the LLMs receive a multifaceted set of demonstrations, ultimately enhancing the robustness of multi-modal reasoning.

3.1 Incorporation of Visual Information to LLMs

Our method works for both LLMs and LMMs, and our task is multi-modal QA task which contains image and corresponding text question. It poses a significant challenge for LLMs to answer correctly without transitioning the image modality into the textual modality through an auxiliary visual expert model. Therefore, it’s very important for LLM to acquire the visual information of the question through visual information model. Following the implementation of Chameleon and MathVista, our visual information model mainly includes two parts:

Image Captioning We employ an image captioning model to obtain textual description of the given image. The image captioning result is represented as: $\{V_c\}$, which is a segment of text expressing the main content of the image.

Optical Character Recognition Besides image captioning system, we also use the Optical Character Recognition (OCR) system to recognize textual characters in the given image, and the detected texts are represented as $\{V_o\}$.

So, the visual information we used is represented as $V = \{V_c, V_o\}$, which is the concatenation of the generated image caption and texts detected by OCR systems.

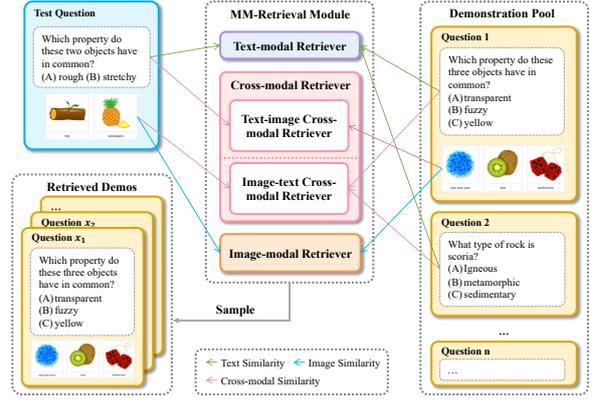


Figure 4: A detailed illustration of our multi-modal retrieval approach, where we use intra-modal similarity and cross-modal similarity to sample demonstration examples D from demonstration pool Q .

3.2 Retrieval Mechanism

Suppose we have a test example q to be answered, which consists of a visual context q^v (usually an image) and a textual context q^t (usually the description of the question). Each question in Q has the same components as q , so $q_i = \{q_i^v, q_i^t\}$ for $q_i \in Q$. Meanwhile, we also have a multi-modal question collection represented as $Q = \{q_1, \dots, q_n\}$, from where we can gather demonstration examples to help LLM answer the test example q . Using a retrieval function, demonstrations are extracted from Q to form the retrieved demonstration set D . The general retrieval process can be expressed as:

$$D = F_k(q, Q) = \arg \max_{q_i \in Q}^k \left(\frac{F_e(q) \cdot F_e(q_i)^T}{\|F_e(q)\| \|F_e(q_i)\|} \right) \quad (1)$$

where F_e denotes our encoder model for encoding q , $F_e(q) \in \mathbf{R}^{1 \times h}$ and $F_e(q_i) \in \mathbf{R}^{1 \times h}$. k indicates that we sample the top- k examples from Q that maximize the cosine similarity with q . Then the sampled top- k examples serve as the demonstration examples.

Specifically, considering the complexity of retrieval in multi-modal settings, we extend Equation 1 by employing q^v and q^t to retrieve demonstration examples from Q respectively:

$$D = F_k(q, Q) = F_{k_1}(q^v, Q^v) \cup F_{k_2}(q^t, Q^t) \cup F_{k_3}(q^v, Q^t) \cup F_{k_4}(q^t, Q^v) \quad (2)$$

where $F_{k_1}(q^v, Q^v)$ means we retrieve the top- k_1 demonstration examples from Q based on the

cosine similarity between $F_e(\mathbf{q}^v) \in \mathbf{R}^{1 \times h}$ and $F_e(\mathbf{q}_i^v) \in \mathbf{R}^{1 \times h}$, which represents comparison between the visual context of \mathbf{q} and example \mathbf{q}_i in the demonstration pool, the same applies for the others. It’s worth noting that $k = \sum_{i=1}^4 k_i$. The first two terms in the right side of Equation 2 represent retrieval based on intra-modal similarity whereas the latter two terms indicate cross-modal retrieval. F_e can be any appropriate encoders to obtain embeddings for \mathbf{q}^t (TEXT-ENCODER) and \mathbf{q}^v (VISUAL-ENCODER).

3.3 Sampling Method

Furthermore, to maintain diversity and relevance in our demonstration examples, we employ the approach of Stratified Sampling (Liberty et al., 2016). This approach allows us to sample demonstration examples from the four groups that we retrieved based on cross-modal and intra-modal similarity. For the sake of simplicity, we simplify $F_{k_1}(\mathbf{q}^v, \mathbf{Q}^v)$ to $D_{k_1}^{I2I}$, $F_{k_2}(\mathbf{q}^t, \mathbf{Q}^t)$ to $D_{k_2}^{T2T}$, $F_{k_3}(\mathbf{q}^v, \mathbf{Q}^t)$ to $D_{k_3}^{I2T}$ and $F_{k_4}(\mathbf{q}^t, \mathbf{Q}^v)$ to $D_{k_4}^{T2I}$. The Stratified Sampling process S is as follows:

$$D = S(D_{k_1}^{I2I}, D_{k_2}^{T2T}, D_{k_3}^{I2T}, D_{k_4}^{T2I}) = \{d_i^{k_1}, d_i^{k_2}, d_i^{k_3}, d_i^{k_4}\}_{i=1}^n \quad (3)$$

where $d_i^{k_1} \in D_{k_1}^{I2I}$, $d_i^{k_2} \in D_{k_2}^{T2T}$, $d_i^{k_3} \in D_{k_3}^{I2T}$ and $d_i^{k_4} \in D_{k_4}^{T2I}$. Moreover, to accommodate the different complex characteristics of multi-modal data (e.g. examples in different domain with various nature), we propose the adaptive use of Stratified Sampling when dealing with different types of questions \mathbf{q} . Specifically, we use an empirical function σ to decide whether to adopt Stratified sampling or not (i.e. σ can be the performance Δ on dev set or other heuristics). The detailed sampling strategy is shown in Table 3.

3.4 Final Prediction

With the retrieved demonstration examples, our next step is to predict the final answer for the given test question \mathbf{q} . To achieve this, we combine the test question \mathbf{q} with the retrieved demonstration set \mathbf{D} as well as the visual information \mathbf{V} . The goal is to provide the LLM with an enriched context that encompasses both the initial question and the insights from the relevant demonstration examples. The combination of the question and the demonstration examples can be denoted as $\mathbf{V} \oplus \mathbf{D} \oplus \mathbf{q}$,

which is the *prompt* for the LLM. When we input this *prompt* into the LLM, we obtain our final result:

$$\text{Answer} = \lambda(\mathbf{V} \oplus \mathbf{D} \oplus \mathbf{q}) \quad (4)$$

In this equation, λ represents the prediction process of the LLM. This final prediction step is essential, as it embodies the entire process we’ve established. It ensures that the LLM takes into account both the original question and the additional information from the demonstration set to generate a well-informed and accurate answer.

4 Experiments

4.1 Experimental Setup

Datasets In experiments, we utilized two benchmark datasets for multi-modal CoT reasoning: 1) **ScienceQA** (Lu et al., 2022) is a comprehensive benchmark designed to evaluate multi-modal reasoning ability, particularly for large-scale language models. 2) **MathVista** (Lu et al., 2023a) is designed to evaluate the mathematical reasoning capabilities of LLMs and LMMs in visual contexts.

Models In our experiments, we employ ChatGPT (OpenAI, 2022), GPT-4 (OpenAI, 2023) and GPT-4V (OpenAI, 2023) via OpenAI API¹. We utilize GPT-3.5-TURBO for ChatGPT, while GPT-4 is employed for GPT-4. As for the evaluation of GPT-4V, we use GPT-4-VISION-PREVIEW.

Implementation Details The TEXT-ENCODER and VISUAL-ENCODER used to encode textual and visual context of the CoT example are models pre-trained on large-scale corpora and images. Specifically, for intra-modality similarity, we encode texts and images using SentenceBERT (Reimers and Gurevych, 2019) and ViT (Dosovitskiy et al., 2021) (ViT-base-patch16-224), respectively. For cross-modality similarity, we encode texts and images using CLIP (Radford et al., 2021). Specifically, we followed Chameleon, which involved concatenating meta data and knowledge retrieval results with the current question as our baseline. Our text-based question method is subsequently developed based on this foundation. For the incorporation of visual context, we use BLIP (Li et al., 2023c) and GPT-4V (OpenAI, 2023) to obtain image captions for ScienceQA and MathVista, respectively. We chose the training set of ScienceQA as

¹<https://platform.openai.com/docs/model-index-for-researchers>

Table 1: Experimental results on ScienceQA (Lu et al., 2022). We evaluated the performance of our system by comparing it with various baseline models, which encompassed both supervised and unsupervised models. Results indicate that our proposed approach *CoT-MM-Retrieval* outperforms previous state-of-the-art models in terms of average accuracy and nearly all question categories on ScienceQA. It should be noted that * represents the best results of our method where we employ more demonstration examples.

Methods	Supervised	Avg	NAT	SOC	LAN	TXT	IMG	NO	G1-6	G7-12
Random Chance	✓	39.8	40.3	46.1	29.3	47.5	40.1	33.7	39.4	40.7
LLaVa (GPT-4) (Liu et al., 2023b)	✓	92.5	91.6	96.7	91.1	90.6	89.0	93.5	92.7	92.2
LLaMA-SciTune (CTOM) (Horawalavithana et al., 2023)	✓	90.0	89.3	95.6	87.0	93.1	86.7	91.8	84.4	91.3
ChatGPT CoT (Lu et al., 2023b)	✗	78.3	78.8	71.0	83.2	77.4	67.9	86.1	80.7	74.0
Chameleon (ChatGPT) (Lu et al., 2023b)	✗	79.9	81.6	70.6	84.0	79.8	70.8	86.6	81.9	76.5
ChatGPT CoT-MM-Retrieval	✗	84.7	84.4	86.4	83.8	83.1	79.2	87.3	86.9	80.6
ChatGPT CoT-MM-Retrieval*	✗	86.4	84.6	93.9	83.9	83.4	82.8	87.3	89.1	81.6
GPT-4 CoT (Lu et al., 2023b)	✗	84.0	85.5	72.4	90.3	82.7	71.5	92.9	86.7	79.0
Chameleon (GPT-4) (Lu et al., 2023b)	✗	86.5	89.8	74.1	89.8	88.3	77.6	92.1	88.0	83.7
GPT-4 CoT-MM-Retrieval	✗	90.9	90.2	89.2	93.5	89.3	85.1	94.9	92.1	88.7
GPT-4 CoT-MM-Retrieval*	✗	92.5	90.2	97.2	93.5	89.4	88.6	94.9	93.5	90.8
GPT-4V (zero-shot)	✗	90.4	92.7	83.2	91.3	91.6	85.3	92.9	91.3	88.7
GPT-4V CoT-MM-Retrieval	✗	93.1	94.5	89.5	93.1	93.7	89.7	94.5	94.1	91.3
Human Average (Lu et al., 2022)	✗	88.4	90.2	85.0	87.5	89.6	87.5	88.1	91.6	82.4

the demonstration pool. As for MathVista, we evaluated the performance on the test-mini. Given that the test-set does not have published answers and is more substantial in quantity, we chose to use it as the demonstration pool. Since there are no answers or rationale in MathVista’s test-set, we used the responses from GPT-4V zero-shot as the basis for our reasoning. To be specific, for GPT-4V MM-Retrieval, our text prompt is the same as LLMs, but we will add the image of the question after the text prompt.

4.2 Results

In our experiments conducted on the ScienceQA dataset (Lu et al., 2022), as shown in Table 1, various models were evaluated for their performance across different question types. The results of baseline models are taken from the ScienceQA leaderboard², where we distinguish *supervised* systems and *unsupervised* systems based on whether they are explicitly trained on the training set of ScienceQA. Our approach, termed *CoT-MM-Retrieval*, utilizes two demonstration examples, providing a fair comparison with the *Chameleon* (Lu et al., 2023b), while *CoT-MM-Retrieval** indicates our best performance with more demonstration examples. For the models based on ChatGPT: 1) The *Chameleon (ChatGPT)* (Lu et al., 2023b) slightly outperform the base *ChatGPT CoT* with an average of 79.9%. 2) Our approach *ChatGPT CoT-MM-Retrieval* based on ChatGPT with retrieval augmentation obtained an average accuracy of 84.7%,

²<https://scienceqa.github.io/>

outperforming previous state-of-the-art *Chameleon* by 4.8% 4) Among these, *ChatGPT CoT-MM-Retrieval** achieved the best performance with an average score of 86.4%. For the GPT-4 based models: 1) *Chameleon (GPT-4)* (Lu et al., 2023b), representing the previous state-of-the-art, achieved an average of 86.5% 2) Our method, *GPT-4 CoT-MM-Retrieval**, surpassed the *Chameleon (GPT-4)* by 6%, achieving an average score of 92.5%. It set a new state-of-the-art, especially in question types like SOC and NO, with scores of 97.2% and 94.9% respectively. For the model based on GPT-4V, our method *GPT-4V CoT-MM-Retrieval* which using the problem image, has surpassed GPT-4V in zero shot’s average score by 2.7%, indicating that our method is applicable not only to LLMs, but also to LMMs.

In our experiments conducted on the MathVista dataset (Lu et al., 2023a), as shown in Table 2. Our approach, termed *CoT-MM-Retrieval*, utilizes two demonstration examples, providing a fair comparison with *CoT* and *PoT*. For the models based on ChatGPT, our method *ChatGPT CoT-MM-Retrieval* surpassed the *ChatGPT CoT* by 8.4%, achieving an average score of 41.6%. For the models based on GPT-4, our method *GPT-4 CoT-MM-Retrieval* surpassed the *ChatGPT CoT* and *ChatGPT PoT* by 13.6% and 12.9% respectively. For the models based on GPT-4V, our method *GPT-4V CoT-MM-Retrieval* surpassed GPT-4V in zero shot average score by 2.7%, and among the 14 metrics, 8 of them have surpassed zero shot. It is worth noting that in MathVista, due to the difficulty of the

Table 2: Experimental results on MathVista (Lu et al., 2023a). We have evaluated the performance of our retrieval method in comparison with the traditional CoT method. Our proposed approach *CoT-MM-Retrieval* demonstrates exceptional performance across nearly all question categories within MathVista.

Methods	ALL	FQA	GPS	MWP	TQA	VQA	ALG	ARI	GEO	LOG	NUM	SCI	STA
Random Chance (Lu et al., 2023a)	17.9	15.5	24.1	4.5	23.4	24.3	25.8	13.8	22.7	13.4	8.8	15.8	14.3
ChatGPT CoT (Lu et al., 2023a)	33.2	26.0	31.7	35.5	48.1	30.2	32.4	32.3	33.0	16.2	17.4	54.9	36.2
ChatGPT CoT-MM-Retrieval	41.6	33.1	40.9	50.0	59.5	30.7	44.5	40.8	41.8	10.8	18.8	54.1	46.2
GPT-4 CoT (Lu et al., 2023a)	33.2	27.9	31.7	31.2	51.9	28.5	33.5	30.9	32.2	13.5	12.5	58.2	37.9
GPT-4 PoT (Lu et al., 2023a)	33.9	30.1	39.4	30.6	39.9	31.3	37.4	31.7	41.0	18.9	20.1	44.3	37.9
GPT-4 CoT-MM-Retrieval	46.8	43.1	45.7	52.7	63.9	32.4	50.2	45.0	46.4	43.2	18.1	55.7	52.5
GPT-4V (zero-shot)	49.3	42.8	54.8	53.8	60.1	38.6	54.1	46.7	54.0	29.7	25.7	60.7	54.2
GPT-4V CoT-MM-Retrieval	52.0	44.6	47.1	59.7	66.5	35.8	50.9	47.6	48.5	37.8	26.4	58.2	55.2
Human Performance (Lu et al., 2023a)	60.3	59.7	48.4	73.0	63.2	55.9	50.9	59.2	51.4	40.7	53.8	64.9	63.9

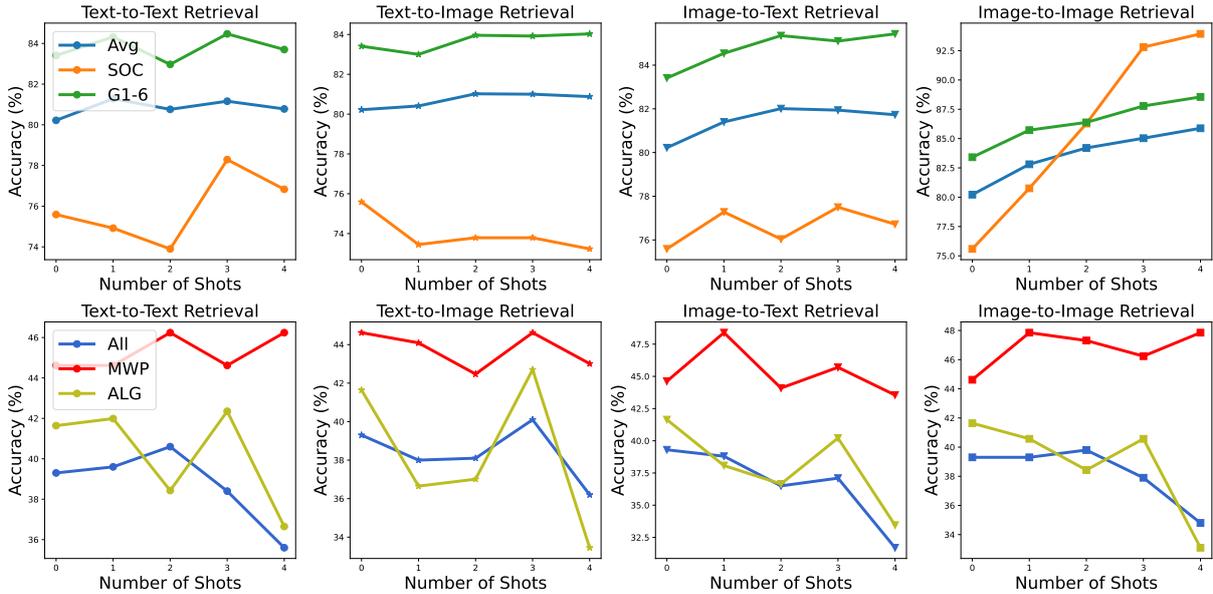


Figure 5: Ablation study of four retrieval methods: *Text-to-Text Retrieval*, *Text-to-Image Retrieval*, *Image-to-Text Retrieval*, *Image-to-Image Retrieval* on ScienceQA (upper) and MathVista (bottom). We inspect the performance of each retrieval approach under different amount of demonstration examples.

Table 3: Sampling strategy in ScienceQA and MathVista.

Dataset	$q^v \notin q$	$q^v \in q$
ScienceQA	$S(D_{k_4}^{T2I}, D_{k_2}^{T2T})$	$D_{k_1}^{I2I}$
MathVista	\emptyset	$S(D_{k_2}^{T2T}, D_{k_1}^{I2I})$

math problems, the current GPT-4V cannot surpass humans in average scores.

4.3 Ablation Studies

We conduct analysis towards the effect of different retrieval methods and the amount of demonstration examples (shots in few-shot learning) in Equation 2. The results are shown in Figure 5.

All four retrieval methods: 1) Text-to-Text (T2T), 2) Text-to-Image (T2I), 3) Image-to-

Text (I2T), 4) Image-to-Image (I2I) are explored with increasing shots (from 0 to 4 for varying k_1, k_2, k_3, k_4) to study their impact on the model’s performance. The performance metrics are provided for various question types, allowing us to discern patterns and variations across different categories. The results in Figure 5 firstly show that adding demonstration examples in the context can improve the overall accuracy especially for ScienceQA and MathVista. We can also observe from Figure 5: 1) Text-to-Text Retrieval: The accuracy for T2T retrieval remains fairly consistent as the number of shots increases on ScienceQA. Specifically, the average accuracy ranges between 80.8% and 81.3%. And the average accuracy ranges between 35.6% and 40.6% on MathVista. The highest accuracy for this method is achieved with 1 shot (81.3%) and 2 shots(40.6%) on ScienceQA and

422
423
424
425
426
427
428
429
430
431
432
433
434
435
436
437
438
439

Table 4: Ablation Studies of the three major components of our proposed approach: *Visual Knowledge, Retrieval Mechanism, Stratified Sampling*. We present results on ScienceQA and MathVista where \times means we do not employ the corresponding method in our approach and \checkmark means we use the corresponding method in the system.

Vision Knowledge	Retrieval	Stratified Sampling	ScienceQA			MathVista		
			Avg	SOC	G1-6	ALL	MWP	ALG
\times	\times	\times	77.8	71.2	80.5	28.0	12.4	35.9
\checkmark	\times	\times	80.7	74.6	83.7	39.4	45.7	39.9
\checkmark	\checkmark (T2T)	\times	80.8	73.9	83.0	40.6	46.2	38.4
\checkmark	\checkmark (T2I)	\times	81.0	73.8	84.0	38.1	42.5	37.0
\checkmark	\checkmark (I2T)	\times	82	76	85.4	36.5	44.1	36.7
\checkmark	\checkmark (I2I)	\times	84.2	86.3	86.4	39.8	47.3	38.4
\checkmark	\checkmark	\checkmark	84.7	86.4	86.9	41.6	50.0	44.5

MathVista, suggesting that adding more demonstration examples does not always guarantee performance improvement. 2) Text-to-Image Retrieval: The performance is similar to T2T, with average accuracy ranging from 80.4% to 81% on ScienceQA and from 36.2% to 40.1% on MathVista. For this method, the highest accuracy is achieved with 2 shots (81%) on ScienceQA, and 3 shots(40.1%) on MathVista. 3) Image-to-Text Retrieval: The accuracy for this method is slightly more varied than the previous two, ranging from 81.4% to 82% on ScienceQA and from 31.7% 38.8% on MathVista. Here, 2 shots provide the best average performance at 82% on ScienceQA, and 1 shot (38.8%) is the highest accuracy. 4) Image-to-Image Retrieval: This performance is similar to the Text-to-Text Retrieval on MathVista, with average accuracy ranges from 34.8% to 39.8%. The highest accuracy for this method is 2 shots(39.8%). On ScienceQA, this retrieval method showcases the most interesting trend. The accuracy improves significantly with increasing shots, starting from 82.8% with 1 shot and reaching 85.9% with 4 shots. On ScienceQA, the G1-6 type consistently performs well, with accuracy usually above 84%. The choice of retrieval method and the number of shots plays a crucial role in determining model performance. Our ablation study results demonstrate the robustness and adaptability of our proposed retrieval strategies across different modalities and varying number of shots. Notably, the consistent performance across diverse question types emphasizes the efficacy of our approach in enhancing the reasoning capabilities of LLMs.

5 Conclusion

In this paper, we proposed a novel method to address the challenge of selecting suitable demonstration examples for multi-modal reasoning for

LLMs (Lu et al., 2022). By integrating retrieval mechanisms with LLMs and emphasizing the modality connection between text and images, our approach aims to improve the efficacy of LLMs for multi-modal Chain-of-Thoughts (CoT) (Wei et al., 2022b; Zhang et al., 2023c) reasoning. Furthermore, the introduction of Stratified Sampling in our methodology ensures that LLMs are exposed to a varied and comprehensive set of demonstration examples. In our experiments on the ScienceQA dataset (Lu et al., 2022) and MathVista dataset (Lu et al., 2023a), our method consistently outperformed existing state-of-the-art models like Chameleon (Lu et al., 2023b) and ChatGPT PoT (Lu et al., 2023a). These empirical results validate our hypothesis that combining LLMs with tailored retrieval mechanisms, like the ones we propose, can significantly enhance multi-modal reasoning performance. As CoT techniques continue to gain traction in the NLP community, our work underscores the importance of effective demonstration example selection.

Future research should focus on refining the retrieval processes and extending the methodologies developed in this study to additional multi-modal tasks. These tasks could include those where the outputs from Large Language Models (LLMs) encompass multiple modalities, such as text-to-image and text-to-video generation (Liu et al., 2023a; Wang et al., 2023d). Additionally, application in specialized domains, like the medical field (Li et al., 2023d), presents a promising direction. Concurrently, in the development of increasingly sophisticated multi-modal LLMs, particularly those employing CoT reasoning, addressing the issue of hallucination is important (Ji et al., 2023; Zhang et al., 2023b). We believe that our work lays a strong foundation for these future explorations.

516 Limitations

517 In this paper, we present a new approach for aug-
518 menting LLMs with multi-modal retrieval for CoT
519 examples. However, there are a few limitations for
520 our work. Firstly, we only test our approach on
521 two datasets: ScienceQA and MathVista. These
522 two datasets are mainly about science and math
523 topics which do not have a broad coverage of
524 other complex reasoning tasks. Therefore, our
525 approach should also be evaluated on other com-
526 plex reasoning tasks. Secondly, due to the limita-
527 tion of resources we only conduct experiments on
528 close-source systems and do not carry out experi-
529 ments on open-source LLMs, making it harder and
530 more expensive to reproduce our results. Moreover,
531 due to the nature of these close-source LLMs we
532 can not fully eliminate the risk of data contami-
533 nation. Therefore, our approach should also be
534 evaluated comprehensively on more representative
535 languages.

536 References

537 Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc,
538 Antoine Miech, Iain Barr, Yana Hasson, Karel
539 Lenc, Arthur Mensch, Katherine Millican, Malcolm
540 Reynolds, et al. 2022. Flamingo: a visual language
541 model for few-shot learning. *Advances in Neural
542 Information Processing Systems*, 35:23716–23736.

543 Amir Bar, Yossi Gandelsman, Trevor Darrell, Amir
544 Globerson, and Alexei Efros. 2022. Visual prompt-
545 ing via image inpainting. *Advances in Neural Infor-
546 mation Processing Systems*, 35:25005–25017.

547 Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie
548 Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind
549 Neelakantan, Pranav Shyam, Girish Sastry, Amanda
550 Askell, Sandhini Agarwal, Ariel Herbert-Voss,
551 Gretchen Krueger, Tom Henighan, Rewon Child,
552 Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu,
553 Clemens Winter, Christopher Hesse, Mark Chen, Eric
554 Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess,
555 Jack Clark, Christopher Berner, Sam McCandlish,
556 Alec Radford, Ilya Sutskever, and Dario Amodei.
557 2020. [Language models are few-shot learners](#). In *Ad-
558 vances in Neural Information Processing Systems 33:
559 Annual Conference on Neural Information Process-
560 ing Systems 2020, NeurIPS 2020, December 6-12,
561 2020, virtual*.

562 Wenhu Chen, Hexiang Hu, Xi Chen, Pat Verga, and
563 William W. Cohen. 2022. [Murag: Multimodal
564 retrieval-augmented generator for open question an-
565 swering over images and text](#).

566 Aakanksha Chowdhery, Sharan Narang, Jacob Devlin,
567 Maarten Bosma, Gaurav Mishra, Adam Roberts,
568 Paul Barham, Hyung Won Chung, Charles Sutton,

Sebastian Gehrmann, et al. 2022. PaLM: Scaling
language modeling with pathways. *arXiv preprint
arXiv:2204.02311*. 569
570
571

Michiel de Jong, Yury Zemlyanskiy, Nicholas FitzGer-
ald, Joshua Ainslie, Sumit K. Sanghai, Fei Sha, and
William Cohen. 2023. Pre-computed memory or on-
the-fly encoding? a hybrid approach to retrieval aug-
mentation makes the most of your compute. *arXiv
preprint*, abs/2301.10448. 572
573
574
575
576
577

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and
Kristina Toutanova. 2019. [BERT: Pre-training of
deep bidirectional transformers for language under-
standing](#). In *Proc. of NAACL-HLT*, pages 4171–4186,
Minneapolis, Minnesota. Association for Computa-
tional Linguistics. 578
579
580
581
582
583

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](#). 584
585
586
587
588
589
590

Yao Fu, Hao Peng, Ashish Sabharwal, Peter Clark, and
Tushar Khot. 2022. Complexity-based prompting for
multi-step reasoning. *CoRR*, abs/2210.00720. 591
592
593

Kelvin Guu, Kenton Lee, Zora Tung, Panupong Pasu-
pat, and Ming-Wei Chang. 2020. [Realm: Retrieval-
augmented language model pre-training](#). *arXiv
preprint*, abs/2002.08909. 594
595
596
597

Sebastian Hofstätter, Jiecao Chen, Krishnan Raman,
and Hamed Zamani. 2022. Fid-light: Efficient and
effective retrieval-augmented text generation. *Pro-
ceedings of the 46th International ACM SIGIR Con-
ference on Research and Development in Information
Retrieval*. 598
599
600
601
602
603

Sameera Horawalavithana, Sai Munikoti, Ian Stewart,
and Henry Kvinge. 2023. [Scitune: Aligning large
language models with scientific multimodal instruc-
tions](#). 604
605
606
607

Ziniu Hu, Ahmet Iscen, Chen Sun, Zirui Wang, Kai-
Wei Chang, Yizhou Sun, Cordelia Schmid, David A.
Ross, and Alireza Fathi. 2022. [Reveal: Retrieval-
augmented visual-language pre-training with multi-
source multimodal knowledge memory](#). In *2023
IEEE/CVF Conference on Computer Vision and Pat-
tern Recognition (CVPR)*. 608
609
610
611
612
613
614

Qian Huang, Hongyu Ren, Peng Chen, Gregor Kržmanc,
Daniel Zeng, Percy Liang, and Jure Leskovec. 2023.
[Prodigy: Enabling in-context learning over graphs](#).
arXiv preprint arXiv:2305.12600. 615
616
617
618

Ziwei Ji, Nayeon Lee, Rita Frieske, Tiezheng Yu, Dan
Su, Yan Xu, Etsuko Ishii, Ye Jin Bang, Andrea
Madotto, and Pascale Fung. 2023. Survey of halluci-
nation in natural language generation. *ACM Comput-
ing Surveys*, 55(12):1–38. 619
620
621
622
623

624	Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. 2022. Large language models are zero-shot reasoners. <i>arXiv preprint arXiv:2205.11916</i> .	question answering. In <i>The 36th Conference on Neural Information Processing Systems (NeurIPS)</i> .	679 680
628	Patrick Lewis, Ethan Perez, Aleksandara Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Kuttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, Sebastian Riedel, and Douwe Kiela. 2020. Retrieval-augmented generation for knowledge-intensive nlp tasks. <i>arXiv preprint</i> , abs/2005.11401.	Pan Lu, Baolin Peng, Hao Cheng, Michel Galley, Kai-Wei Chang, Ying Nian Wu, Song-Chun Zhu, and Jianfeng Gao. 2023b. Chameleon: Plug-and-play compositional reasoning with large language models .	681 682 683 684
634	Bo Li, Yuanhan Zhang, Liangyu Chen, Jinghao Wang, Fanyi Pu, Jingkang Yang, Chunyuan Li, and Ziwei Liu. 2023a. Mimic-it: Multi-modal in-context instruction tuning .	Yujie Lu, Pan Lu, Zhiyu Chen, Wanrong Zhu, Xin Eric Wang, and William Yang Wang. 2023c. Multimodal procedural planning via dual text-image prompting. <i>arXiv preprint arXiv:2305.01795</i> .	685 686 687 688
638	Bo Li, Yuanhan Zhang, Liangyu Chen, Jinghao Wang, Jingkang Yang, and Ziwei Liu. 2023b. Otter: A multi-modal model with in-context instruction tuning .	Chenyang Lyu, Minghao Wu, Longyue Wang, Xinting Huang, Bingshuai Liu, Zefeng Du, Shuming Shi, and Zhaopeng Tu. 2023. Macaw-11m: Multi-modal language modeling with image, audio, video, and text integration. <i>arXiv preprint arXiv:2306.09093</i> .	689 690 691 692 693
642	Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. 2023c. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. <i>arXiv preprint arXiv:2301.12597</i> .	OpenAI. 2022. Chatgpt .	694
646	Yingshu Li, Yunyi Liu, Zhanyu Wang, Xinyu Liang, Lingqiao Liu, Lei Wang, Leyang Cui, Zhaopeng Tu, Longyue Wang, and Luping Zhou. 2023d. A comprehensive study of gpt-4v’s multimodal capabilities in medical imaging. <i>medRxiv</i> , pages 2023–11.	OpenAI. 2023. GPT-4 technical report. <i>ArXiv</i> , abs/2303.08774.	695 696
651	Yunxin Li, Longyue Wang, Baotian Hu, Xinyu Chen, Wanqi Zhong, Chenyang Lyu, and Min Zhang. 2023e. A comprehensive evaluation of gpt-4v on knowledge-intensive visual question answering. <i>arXiv preprint arXiv:2311.07536</i> .	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 .	697 698 699 700 701 702
656	Edo Liberty, Kevin Lang, and Konstantin Shmakov. 2016. Stratified sampling meets machine learning . In <i>Proceedings of The 33rd International Conference on Machine Learning</i> , volume 48 of <i>Proceedings of Machine Learning Research</i> , pages 2320–2329, New York, New York, USA. PMLR.	Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners.	703 704 705
662	Bingshuai Liu, Longyue Wang, Chenyang Lyu, Yong Zhang, Jinsong Su, Shuming Shi, and Zhaopeng Tu. 2023a. On the cultural gap in text-to-image generation. <i>arXiv preprint arXiv:2307.02971</i> .	Ori Ram, Yoav Levine, Itay Dalmedigos, Dor Muhlgay, Amnon Shashua, Kevin Leyton-Brown, and Yoav Shoham. 2023. In-context retrieval-augmented language models. <i>Transactions of the Association for Computational Linguistics</i> , 11:1316–1331.	706 707 708 709 710
666	Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. 2023b. Visual instruction tuning. <i>arXiv preprint arXiv:2304.08485</i> .	Nils Reimers and Iryna Gurevych. 2019. Sentence-bert: Sentence embeddings using siamese bert-networks .	711 712
669	Pan Lu, Hritik Bansal, Tony Xia, Jiacheng Liu, Chunyuan Li, Hannaneh Hajishirzi, Hao Cheng, Kai-Wei Chang, Michel Galley, and Jianfeng Gao. 2023a. Mathvista: Evaluating mathematical reasoning of foundation models in visual contexts. <i>arXiv preprint arXiv:2310.02255</i> .	Freda Shi, Mirac Suzgun, Markus Freitag, Xuezhi Wang, Suraj Srivats, Soroush Vosoughi, Hyung Won Chung, Yi Tay, Sebastian Ruder, Denny Zhou, et al. 2022. Language models are multilingual chain-of-thought reasoners . <i>ArXiv preprint</i> , abs/2210.03057.	713 714 715 716 717
675	Pan Lu, Swaroop Mishra, Tony Xia, Liang Qiu, Kai-Wei Chang, Song-Chun Zhu, Oyvind Tafjord, Peter Clark, and Ashwin Kalyan. 2022. Learn to explain: Multimodal reasoning via thought chains for science	Weijia Shi, Sewon Min, Michihiro Yasunaga, Minjoon Seo, Rich James, Mike Lewis, Luke Zettlemoyer, and Wen-tau Yih. 2023. Replug: Retrieval-augmented black-box language models. <i>arXiv preprint</i> , abs/2301.12652.	718 719 720 721 722
676		Yanpeng Sun, Qiang Chen, Jian Wang, Jingdong Wang, and Zechao Li. 2023. Exploring effective factors for improving visual in-context learning .	723 724 725
677		Maria Tsimpoukelli, Jacob L Menick, Serkan Cabi, SM Eslami, Oriol Vinyals, and Felix Hill. 2021. Multimodal few-shot learning with frozen language models. <i>Advances in Neural Information Processing Systems</i> , 34:200–212.	726 727 728 729 730

731	Boshi Wang, Sewon Min, Xiang Deng, Jiaming Shen,	Zhuosheng Zhang, Aston Zhang, Mu Li, Hai Zhao,	785
732	You Wu, Luke Zettlemoyer, and Huan Sun. 2023a.	George Karypis, and Alex Smola. 2023c. Multi-	786
733	Towards understanding chain-of-thought prompting:	modal chain-of-thought reasoning in language mod-	787
734	An empirical study of what matters. In <i>Proceedings</i>	els. <i>arXiv preprint arXiv:2302.00923</i> .	788
735	<i>of the 61st Annual Meeting of the Association for</i>		
736	<i>Computational Linguistics (Volume 1: Long Papers)</i> ,	Zihao Zhao, Eric Wallace, Shi Feng, Dan Klein, and	789
737	pages 2717–2739, Toronto, Canada. Association for	Sameer Singh. 2021. Calibrate before use: Im-	790
738	Computational Linguistics.	proving few-shot performance of language models.	791
		In <i>Proc. of ICML</i> , volume 139 of <i>Proceedings of</i>	792
739	Xinlong Wang, Wen Wang, Yue Cao, Chunhua Shen,	<i>Machine Learning Research</i> , pages 12697–12706.	793
740	and Tiejun Huang. 2023b. Images speak in images:	PMLR.	794
741	A generalist painter for in-context visual learning. In		
742	<i>Proceedings of the IEEE/CVF Conference on Com-</i>	Denny Zhou, Nathanael Schärli, Le Hou, Jason Wei,	795
743	<i>puter Vision and Pattern Recognition</i> , pages 6830–	Nathan Scales, Xuezhi Wang, Dale Schuurmans,	796
744	6839.	Olivier Bousquet, Quoc Le, and Ed Chi. 2022.	797
		Least-to-most prompting enables complex reason-	798
745	Xinlong Wang, Xiaosong Zhang, Yue Cao, Wen Wang,	ing in large language models. <i>ArXiv preprint</i> ,	799
746	Chunhua Shen, and Tiejun Huang. 2023c. Seggpt:	abs/2205.10625.	800
747	Segmenting everything in context. <i>arXiv preprint</i>		
748	<i>arXiv:2304.03284</i> .	Jin Ziqi and Wei Lu. 2023. Tab-CoT: Zero-shot tabular	801
		chain of thought. In <i>Findings of the Association for</i>	802
749	Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc Le,	<i>Computational Linguistics: ACL 2023</i> , pages 10259–	803
750	Ed Chi, Sharan Narang, Aakanksha Chowdhery, and	10277, Toronto, Canada. Association for Computa-	804
751	Denny Zhou. 2022. Self-consistency improves chain	tional Linguistics.	805
752	of thought reasoning in language models. <i>arXiv</i>		
753	<i>preprint arXiv:2203.11171</i> .		
754	Zhanyu Wang, Longyue Wang, Zhen Zhao, Minghao		
755	Wu, Chenyang Lyu, Huayang Li, Deng Cai, Lup-		
756	ing Zhou, Shuming Shi, and Zhaopeng Tu. 2023d.		
757	Gpt4video: A unified multimodal large language		
758	model for Instruction-followed understanding and		
759	safety-aware generation.		
760	Zhendong Wang, Yifan Jiang, Yadong Lu, Yelong		
761	Shen, Pengcheng He, Weizhu Chen, Zhangyang		
762	Wang, and Mingyuan Zhou. 2023e. In-context learn-		
763	ing unlocked for diffusion models. <i>arXiv preprint</i>		
764	<i>arXiv:2305.01115</i> .		
765	Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten		
766	Bosma, Ed Chi, Quoc Le, and Denny Zhou. 2022a.		
767	Chain of thought prompting elicits reasoning in large		
768	language models. <i>arXiv preprint arXiv:2201.11903</i> .		
769	Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten		
770	Bosma, Ed Chi, Quoc Le, and Denny Zhou. 2022b.		
771	Chain of thought prompting elicits reasoning in large		
772	language models. <i>ArXiv preprint</i> , abs/2201.11903.		
773	Yuanhan Zhang, Kaiyang Zhou, and Ziwei Liu. 2023a.		
774	What makes good examples for visual in-context		
775	learning?		
776	Yue Zhang, Yafu Li, Leyang Cui, Deng Cai, Lemao Liu,		
777	Tingchen Fu, Xinting Huang, Enbo Zhao, Yu Zhang,		
778	Yulong Chen, Longyue Wang, Anh Tuan Luu, Wei		
779	Bi, Freda Shi, and Shuming Shi. 2023b. Siren’s song		
780	in the ai ocean: A survey on hallucination in large		
781	language models.		
782	Zhuosheng Zhang, Aston Zhang, Mu Li, and Alex		
783	Smola. 2022. Automatic chain of thought prompting		
784	in large language models. <i>CoRR</i> , abs/2210.03493.		

A Appendices

A.1 Case Study

We present a case study to demonstrate the effectiveness of our approach in Figure 6, Figure 7, Figure 8 and Figure 9.

In Figure 6 and Figure 7, compared to CoT-2-shots that uses fixed samples regardless of the question, our method is able to retrieve Demonstrations based on the current question, and control the diversity of samples through hierarchical sampling, thus effectively stimulating the reasoning ability of LLMs.

In Figure 8, we illustrate our approach’s capability when dealing with a question with an image, such as determining the interaction between two magnets. Our approach enhances the LLM’s reasoning by retrieving related questions that provide deeper understanding of magnetic properties, such as attraction and repulsion based on pole alignment, enriching the LLM’s knowledge base for more accurate inference.

Figure 9 highlights our approach’s adaptability to questions without visual context, such as deducing alphabetical order in a dictionary. It showcases the approach’s use of varied demonstrations—incorporating both visual and non-visual information—to strengthen the LLM’s ability to discern patterns and sort information. This demonstrates the approach’s versatility in improving reasoning by broadening the range of demonstrations, which in turn equips the LLM to handle a diverse set of textual queries effectively.

Problem

Question: Find the value of the square in the figure.

Answer: The answer is 2.

Vision

Brain Teaser for IQ Test

$\text{Blue Circle} + \text{Red Square} = 5$
 $2 \times \text{Yellow Triangle} = 8$
 $\text{Blue Circle} + \text{Yellow Triangle} = 7$
 $\text{Red Square} = ?$

CoT-2-shots

Question: How much money does Ruth need to buy a baking dish, a casserole dish, and an ice cream scoop?

(Unit: \$)

Solution: Find the total cost of a baking dish, a casserole dish, and an ice cream scoop.
 $\$4.00 + \$3.00 + \$6.00 = \13.00
 Ruth needs \$13.00.

Vision

baking dish	\$4.00
ice cream scoop	\$6.00
casserole dish	\$3.00
rolling pin	\$4.00

Question: What is the largest city in the nation where this plane is headquartered?

Options:
 (A) hong kong
 (B) osaka
 (C) shanghai
 (D) tokyo

Solution: The caption mentions that the text on the image says "Japan. Endless Discovery". This indicates that the plane is headquartered in Japan. Among the Japanese cities, Tokyo is the largest city.
 The answer is (D).

Vision

The missing number in the brain teaser is 6. ❌

MM-Retrieval

Question: Find the size of angle x in the figure.

Solution: First, let's consider the large triangle...So, the third angle in the large triangle is 45 degrees...we can find angle x by subtracting the 111-degree angle and the 45-degree angle from 180 degrees:
 $180 \text{ degrees} - 111 \text{ degrees} - 45 \text{ degrees} = 24 \text{ degrees}.$
 Therefore, angle x is 24 degrees.

Vision

Question: What is the value of the smallest individual bar in the whole chart?

Solution: The smallest individual bar in the chart is the orange bar at the "proud" category, which appears to have a value of -6.
 The answer is -6.

Vision

To find the value of the red square in the figure, we need to analyze the given equations.

- A blue circle plus a red square equals 5.
- A yellow triangle plus another yellow triangle equals 8.
- A blue circle plus a yellow triangle equals 7.

Let's assign variables to the unknown values:

- Blue circle = A
- Red square = B
- Yellow triangle = C

From equation 1, we have $A + B = 5$.
 From equation 2, we have $2C = 8$, which simplifies to $C = 4$.
 From equation 3, we have $A + C = 7$.
 Substituting the values we know, we can solve for A:
 $A + 4 = 7$
 $A = 3$
 Now that we know $A = 3$, we can substitute it into equation 1 to solve for B:
 $3 + B = 5$
 $B = 2$
 Therefore, the value of the red square is 2. ✓

Figure 6: An example of our method on MathVista. The left side depicts the common 2-shot CoT approach for LLM, while the right side showcases our MM-Retrieval method retrieving demonstration examples as relevant as possible to enhance the reasoning ability of LLM.

Problem

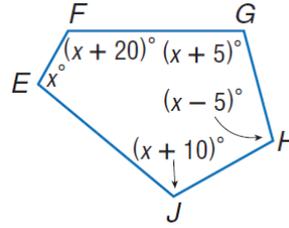
Question: Find the measure of $\angle 3$.

Options:

- (A) 97
- (B) 102
- (C) 107
- (D) 122

Answer: The answer is (A).

Vision



CoT-2-shots

Question: How much money does Ruth need to buy a baking dish, a casserole dish, and an ice cream scoop?

(Unit: \$)

Solution: Find the total cost of a baking dish, a casserole dish, and an ice cream scoop.

$$\$4.00 + \$3.00 + \$6.00 = \$13.00$$

Ruth needs \$13.00.

Vision

baking dish	\$4.00
ice cream scoop	\$6.00
casserole dish	\$3.00
rolling pin	\$4.00

Question: What is the largest city in the nation where this plane is headquartered?

Options:

- (A) hong kong
- (B) osaka
- (C) shanghai
- (D) tokyo

Solution: The caption mentions that the text on the image says "Japan. Endless Discovery". This indicates that the plane is headquartered in Japan. Among the Japanese cities, Tokyo is the largest city.

The answer is (D).

Vision



To find the measure of angle H, we need to use the equation given in the image:

$$a = (x + 20)^\circ + (x + 5)^\circ + (x - 5)^\circ + (x + 10)^\circ + (x)^\circ$$

Since the sum of all the angles in a hexagon is 720° , we can set the equation equal to 720 and solve for x:

$$(x + 20)^\circ + (x + 5)^\circ + (x - 5)^\circ + (x + 10)^\circ + (x)^\circ = 720$$

Combining like terms, we get: $5x + 30 = 720$

Subtracting 30 from both sides, we get: $5x = 690$

Dividing both sides by 5, we get: $x = 138$

Now that we have the value of x, we can substitute it back into the equation to find the measure of angle H:

$$(x - 5)^\circ = (138 - 5)^\circ = 133^\circ$$

Therefore, the measure of angle H is 133° .

The correct option is not provided in the question. ❌

MM-Retrieval

Question: Find the measure of $\angle A$.

Options:

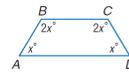
- (A) 30
- (B) 60
- (C) 90
- (D) 120

Solution: ...Since we are looking for the measure of angle A, which is marked as x° , we have found that:

$$\angle A = x = 60^\circ$$

The answer is (B).

Vision



Question: Find the measure of $\angle W$.

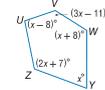
Options:

- (A) 38
- (B) 38
- (C) 76
- (D) 152

Solution: Now that we have the value of x, we can find $\angle W$ by substituting x into the expression for angle W: $\angle W = (x + 8)$ Therefore, the measure of angle W is 76 degrees.

The answer is (C).

Vision



To find the measure of angle H, denoted as $\angle H$, we need to use the fact that the sum of the interior angles of a pentagon is 540 degrees.

We can set up an equation using the expressions given for each angle in the pentagon:

$$x + (x + 20) + (x + 5) + (x - 5) + (x + 10) = 540$$

Combining like terms, we get: $5x + 30 = 540$

Subtracting 30 from both sides: $5x = 510$ Dividing both sides by 5: $x = 102$

Now that we have the value of x, we can find $\angle H$ by substituting x into the expression for angle H: $H = (x - 5)$

$$H = (102 - 5) H = 97$$

Therefore, the measure of angle H is 97 degrees.

The answer is (A). ✓

Figure 7: An example of our method on MathVista. The left side illustrates the conventional 2-shot CoT technique used for LLM, while the right side highlights our MM-Retrieval method, which aims to retrieve demonstration examples that are as relevant as possible to bolster the reasoning capabilities of LLM.

Instruction: Given the question (and the context), select the answer from the options ["A", "B", "C", "D", "E"]. You should give concise and step-by-step solutions. Finally, conclude the answer in the format of "the answer is [ANSWER]", where [ANSWER] is one from the options ["A", "B", "C", "D", "E"]. For example, "the answer is A", "the answer is B", "the answer is C", "the answer is D", or "the answer is E". If the answer is not in the options, select the most possible option.

Demonstration 1:
Question: Will these magnets attract or repel each other?
Context: Two magnets are placed as shown.
Options: (A) repel (B) attract
Solution: Magnets can pull or push on each other without touching. When magnets attract, they pull together. When magnets repel, they push apart. ... Therefore, the answer is A.



Demonstration 2:
 ...

Current Question:
Question: Will these magnets attract or repel each other?
Context: Two magnets are placed as shown.
Options: (A) repel (B) attract
Solution:



Answer:
 Magnets can pull or push on each other without touching. When magnets attract, they pull together. When magnets repel, they push apart. Whether a magnet attracts or repels other magnets depends on the positions of its poles, or ends. Every magnet has two poles: north and south. Here are some examples of magnets. The north pole of each magnet is labeled N, and the south pole is labeled S. If opposite poles are closest to each other, the magnets attract. The magnets in the pair below attract. If the same, or like, poles are closest to each other, the magnets repel. The magnets in both pairs below repel. To predict if these magnets will attract or repel, look at which poles are closest to each other. The south pole of one magnet is closest to the north pole of the other magnet. Opposite poles attract. So, these magnets will attract each other. Therefore, the answer is **B**.

Figure 8: An example of our method on ScienceQA. For the question with visual context, our approach retrieves demonstration examples as relevant as possible to enhance the reasoning ability of LLM.

Instruction: Given the question (and the context), select the answer from the options ["A", "B", "C", "D", "E"]. You should give concise and step-by-step solutions. Finally, conclude the answer in the format of "the answer is [ANSWER]", where [ANSWER] is one from the options ["A", "B", "C", "D", "E"]. For example, "the answer is A", "the answer is B", "the answer is C", "the answer is D", or "the answer is E". If the answer is not in the options, select the most possible option.

Demonstration 1:
Question: Which property do these two objects have in common?
Options: (A) breakable (B) soft
Solution: An object has different properties. ... Therefore, the answer is B.



Demonstration 2:
Question: Which word would you find on a dictionary page with the following guide words? advance - and
Options: (A) amazed (B) arm
Solution: Guide words appear on each page of a dictionary. ... Therefore, the answer is A.

Current Question:
Question: Which word would you find on a dictionary page with the following guide words? shot - suit
Options: (A) service (B) stockade
Solution:

Answer:
 To find the word that would be on a dictionary page with the guide words "shot" and "suit", we need to look for a word that falls alphabetically between these two words. Comparing the first letters of the guide words, "s" comes before "s". Since the first letters are the same, we need to look at the second letters. Comparing the second letters, "h" comes before "u". Therefore, any word that starts with "sh" would come before any word that starts with "su". Looking at the options, "service" starts with "se" and "stockade" starts with "st". Since "stockade" starts with "st" and comes after "shot" alphabetically, it would be the word found on the dictionary page with the guide words "shot" and "suit". Therefore, the answer is B.

Figure 9: An example of our method on ScienceQA. For the question without visual context, our method retrieves demonstration examples with both visual context and textual context, enhancing the diversity of demonstration examples to improve the reasoning ability of LLM.