

Research Article

Document Haystacks: Vision-Language Reasoning Over Piles of 1000+ Documents

Jun Chen¹, Dannong Xu², Junjie Fei¹, Chun-Mei Feng³, Mohamed Elhoseiny¹

1. King Abdullah University of Science and Technology, Saudi Arabia; 2. University of Sydney, Australia; 3. Institute of High Performance Computing, Singapore, Singapore

Large multimodal models (LMMs) have achieved impressive progress in vision-language understanding, yet they face limitations in real-world applications requiring complex reasoning over a large number of images. Existing benchmarks for multi-image question-answering are limited in scope, each question is paired with only up to 30 images, which does not fully capture the demands of large-scale retrieval tasks encountered in the real-world usages. To reduce these gaps, we introduce two document haystack benchmarks, dubbed DocHaystack and InfoHaystack, designed to evaluate LMM performance on large-scale visual document retrieval and understanding. Additionally, we propose V-RAG, a novel, vision-centric retrieval-augmented generation (RAG) framework that leverages a suite of multimodal vision encoders, each optimized for specific strengths, and a dedicated question-document relevance module. V-RAG sets a new standard, with a 9% and 11% improvement in Recall@1 on the challenging DocHaystack-1000 and InfoHaystack-1000 benchmarks, respectively, compared to the previous best baseline models. Additionally, integrating V-RAG with LMMs enables them to efficiently operate across thousands of images, yielding significant improvements on our DocHaystack and InfoHaystack benchmarks. Our code and datasets are available at <https://github.com/Vision-CAIR/dochaystacks>

Dannong Xu and Junjie Fei equally contributed to this work.

Corresponding authors: Dannong Xu, daxu8019@uni.sydney.edu.au; Junjie Fei, junjie.fei@kaust.edu.sa

1. Introduction

Large Multimodal Models (LMMs)^{[1][2][3][4]} have made remarkable progress in the vision-language understanding. However, these models still face challenges when tasked with reasoning over extensive collections of images or documents^[5], limiting their effectiveness in real-world applications, such as visual search or querying over large sets of images or documents, like those stored on personal devices or in photo albums. However, there lacks such proper benchmarks to evaluate these capabilities. To address this gap, we introduce the DocHaystack and InfoHaystack benchmarks, designed to evaluate LMMs on large-scale image retrieval and understanding capabilities, pushing the boundaries of LMM performance in complex, real-world scenarios.

The existing multi-image retrieving and reasoning benchmarks are primarily constructed on a small scale, as highlighted in works such as^{[6][7]}. Each question in these benchmarks is paired with only up to 30 images as illustrated in Figure 1 (a). However, this limited scope does not align well with real-world scenarios, which often require retrieval and reasoning across hundreds or thousands of images or documents. In contrast, our established benchmarks, depicted in Figure 1 (b), allow for querying questions from a large-scale collection of up to 1,000 documents, necessitating that models retrieve and reason over an extensive set of documents for each question. This scale better simulates practical applications and their demands.

The main challenge in constructing such benchmarks is collecting specific questions while ensuring there are no ambiguous answers across a large set of images. Existing datasets, such as those in DocVQA and InfographicVQA^{[8][9]}, contain numerous “general” questions, like “What is the table number?”, where answers could be derived from multiple images, leading to non-unique responses. To address this, we implemented a rigorous data filtering pipeline. First, we employed both a large language model (LLM) and human annotators to systematically filter out “general” questions based on carefully defined criteria. Additionally, we used the LLM to exclude questions relying on generic knowledge, such as “What is the capital of Missouri?”, which can be answered without image context. This approach ensures that the questions in the benchmark can only be answered through specific visual cues from the provided images, maintaining the benchmark’s integrity for evaluating image-based understanding.

To enable the current LMMs effectively reason over a large number of images, we propose a vision-centric retrieval-augmented generation (RAG) framework, named V-RAG. V-RAG combines multiple multimodal vision encoders, leveraging each encoder’s unique strengths to enhance retrieval accuracy.

Additionally, it incorporates an LMM-filter module to assess the relevance of each document to the query, refining the retrieval process by ensuring that only relevant documents are prioritized. This integrated approach allows V-RAG to navigate extensive document collections efficiently. Experimental results demonstrate that V-RAG achieves 9% and 11% improvement in Recall@1 on the DocHaystack-1000 and InfoHaystack-1000 compared to previous best text-to-image retrieval methods. Additionally, we found that integrating V-RAG brings GPT-4o over a 55% acc improvement on DocHaystack-200 and a 34% acc improvement on InfoHaystack-200, indicating the effectiveness of our V-RAG.

Our contributions are as follows:

- We introduce the Document Haystack benchmarks, including DocHaystack-100/200/1000 and InfoHaystack-100/200/1000, with the most challenging setup consisting of 1,000 documents for each inquiry. These benchmarks advance document retrieval and reasoning tasks by requiring models to navigate and reason across extensive document collections, surpassing prior benchmarks limited to smaller retrieval tasks.
- We propose a vision-centric retrieval-augmented generation framework, V-RAG, which enhances the retrieval capabilities of LMMs. V-RAG achieves substantial improvements over previous best text-to-image retrieval methods by 9% and 11% on DocHaystack-1000 and InfoHaystack-1000, respectively.

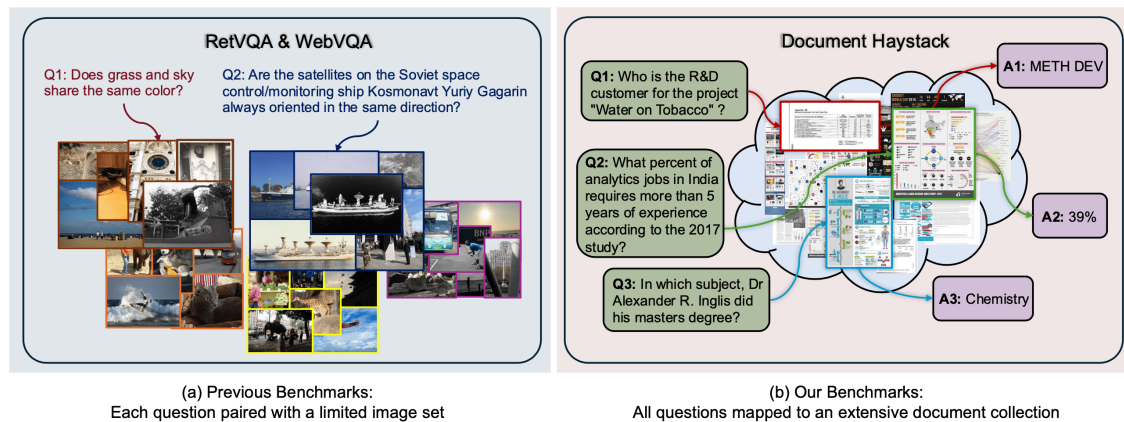


Figure 1. Comparison between previous and proposed benchmarks. Given a question as input, all benchmarks aim to retrieve relevant images from an image pool to correctly answer the question. Unlike prior benchmarks like RetVQA^[6] and WebVQA^[10], which structure their datasets by pairing each question with a limited set of images (typically ≤ 30), our benchmarks, DocHaystack and InfoHaystack, map each question to a substantially larger document collection, scaling up to 1,000 visual documents. This expanded scope more accurately represents large-scale document retrieval scenarios and offers a greater challenge in retrieval accuracy and visual question answering.

2. Related Works

VQA benchmarks.

VQA play a critical role in assessing a model's ability to understand and reason across visual contexts^{[11][12][13]}. Traditional VQA datasets typically measure a model's comprehension of object attributes^{[14][15]}, spatial relationships^[14], as well as its understanding of documents^{[8][9]}, charts^[16], mathematics^{[17][18][19]}, and open knowledges^{[20][21]}. Additionally, these benchmarks explore models' knowledge across varied fields, including science and the arts^{[22][23]}. This broad array of benchmarks has greatly advanced vision-language models by cultivating diverse visual comprehension skills, particularly for modern foundation models in vision-language understanding^{[24][1][2][25][26][22][4][27][28]}. Notably, these benchmarks have primarily focused on question answering within single image or document. In contrast, our benchmark shifts the focus towards retrieval and comprehensive understanding across a large collection of visual documents, presenting new challenges and expanding the scope of visual question answering.

Several previous efforts have tackled the challenge of visual question answering and reasoning across multiple images^{[6][10][29][7][30][31]}. For instance, datasets such as MultimodalQA^[7] and ISVQA^[29] require models to have multi-image reasoning abilities. Meanwhile, WebQA^[10] and RetVQA^[6] involve an additional step where models must first retrieve relevant images from a limited image pool before answering visual questions based on these results. However, these benchmarks are generally constrained to relatively small image pools, where each question is paired with an image set containing up to 30 images. In contrast, our proposed benchmarks, DocHaystack and InfoHaystack, significantly expand this scope by requiring models to retrieve and reason from a much larger set of up to 1,000 documents, presenting a notably greater challenge in retrieval and multi-image reasoning.

Large multimodal models (LMMs).

LMMs have achieved substantial advancements in understanding and reasoning across single or multiple images^{[1][4][2][3][25][26]}. These models have significantly enhanced vision-language understanding across numerous dimensions and applications^{[23][11][19][17]}. LMMs benefit primarily from large-scale image-text alignment and extensive language modeling, which emerge them with advanced understanding and reasoning abilities. However, despite these breakthroughs, LMMs still encounter challenges when handling large-scale image or document sets^[5]. This difficulty is due to the inherent complexity of processing such complex data. To address this, retrieval-based approaches have been developed to extend the capacity of vision-language models, augmenting their ability to process and reason over a larger number of images.

Retrieval-augmented generation (RAG).

RAG integrates retrieval systems^{[32][33][34][35][13]}, with generative models, enhancing them with additional knowledge. While RAG has been extensively explored in language domains^{[36][37][38][39]}, its application in vision-language contexts is also advancing. In vision-language RAG, models like MuRAG^[40] leverage image-text memory to retrieve top-k neighbors by comparing inner-product similarities. RetVQA^[6] uses an image-question relevance encoder, combining BERT^[41] and Faster R-CNN^[42] to filter relevant images, while MIRAGE^[5] employs a CLIP-based encoder^[32] to train a retriever. These frameworks extend model capabilities, enabling retrieval and reasoning across hundreds or thousands of images. In contrast, we propose V-RAG, a vision-centric RAG framework that integrates multiple vision encoders to more effectively capture image features, and introduces a LMM-based

question-document relevance comparison module. Our results demonstrate that V-RAG surpasses existing methods on our DocHaystack and InfoHaystack benchmarks, setting a new standard for large-scale visual retrieval and reasoning.

3. DocHaystack and InfoHaystack Benchmarks

To support effective retrieval and reasoning across extensive document collections, we present two new benchmarks—DocHaystack and InfoHaystack—designed to ensure each question yields a unique, document-specific answer. Derived from DocVQA^[8] and InfographicVQA^[9], these benchmarks address the challenge of answer ambiguity by selectively curating questions that can only be answered by a single document within a large dataset.

Benchmark construction pipeline.

There exists many general questions in the existing benchmarks and lead to multiple answers for different document context. For example, general questions like “What is the table number?” may apply to various documents and yield multiple valid answers, while a targeted question like “Who is the reviewer for the article titled ‘An antithyroid factor in milk?’” is likely to produce a unique answer, as only a single document or a limited set of documents would contain that information. Therefore, our benchmark construction follows a structured three-step filtering pipeline, illustrated in Figure 2, to ensure high-quality, unique-answer questions. First, we employ a large language model (LLM) to filter out general questions that could generate multiple answers across documents. Next, a manual review step further checks the questions to ensure the data quality. Finally, a generic-knowledge filtering stage refines the dataset further, retaining only questions closely tied to specific document content.

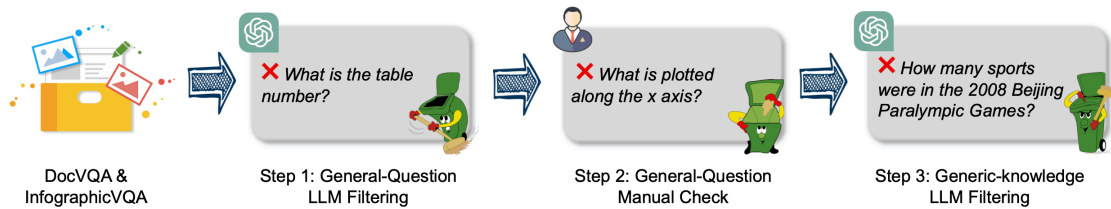


Figure 2. Data Curation Pipeline. Our benchmarks are curated based on the DocVQA and InfographicVQA datasets, following a three-step filtering process to obtain document-specific question-answer pairs. In Step 1, we filter out general questions (e.g., “What is the table number?”), as these could be answered by multiple documents and lack specificity. Step 2 involves a manual review by human annotators to further remove general questions. In Step 3, we eliminate generic-knowledge questions (e.g., “How many sports were in the 2008 Beijing Paralympic Games?”) that can be answered directly by large language models without requiring image input.”

This carefully designed pipeline, combining LLM-based filtering and human review, effectively curates questions that drive accurate, document-specific retrieval. By focusing on reducing answer ambiguity, DocHaystack and InfoHaystack enhance the precision of retrieval and reasoning in large-scale document processing tasks, providing a valuable tool for the evaluation of retrieval systems. We discuss this data curation pipeline in details as follows:

General-question LLM filtering.

We begin by using the LLM, GPT-4o^[1], to filter out general questions through a set of well-crafted instructions. Leveraging the LLM’s strong contextual understanding, this initial filtering step allows us to efficiently process large volumes of data, identifying broad or ambiguous questions that may yield multiple answers across documents. This automated approach significantly enhances the benchmark construction’s efficiency and quality.

To guide the LLM, we first define the task, providing clear distinctions between general and specific questions along with illustrative examples. With this framework, the LLM can then assess each question and determine if it is general or specific. The instructional format is as follows: LLM i

You are an evaluator tasked with identifying if a question is specific or general. A general question seeks commonly known or widely applicable information without unique identifiers, e.g., “Who is the person standing in the ground?” A specific question, however, requests unique information about a particular individual, event, or

object, e.g., “What is the Social Security Number of Charles Yarbrough?” Based on these definitions, determine if the following question is general or specific: {question}.

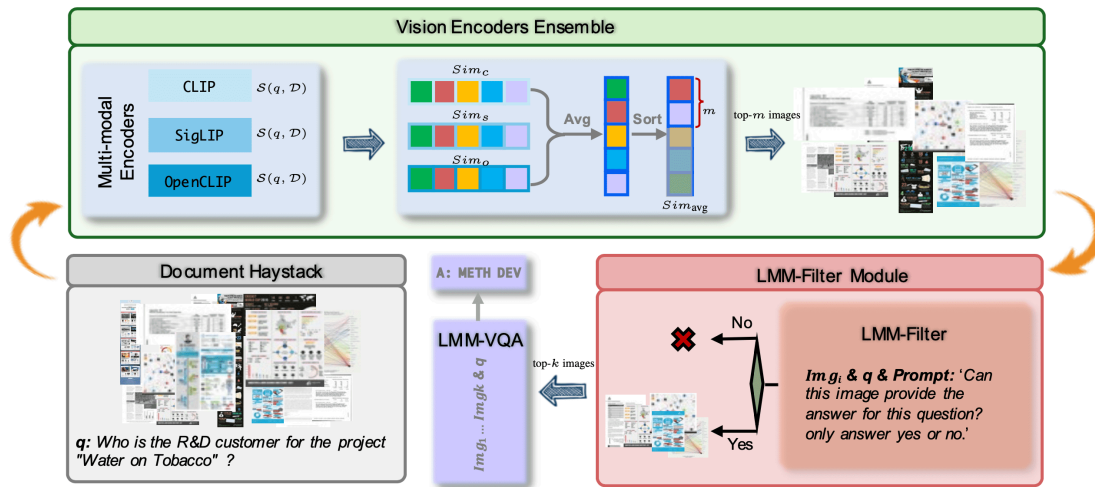


Figure 3. The V-RAG pipeline workflow. In the top section, a vision encoder ensemble is used, combining multiple vision models—CLIP, SigLIP, and OpenCLIP—to process a large document haystack. Each encoder computes similarity scores, which are averaged into Sim_{avg} . The top m documents, based on these scores, are selected for further analysis. In the bottom right, the LMM-Filter Module utilizes a pretrained LMM to assess whether each selected document can potentially answer the posed question. This filtering step removes documents that do not match, retaining only relevant ones. Finally, the top k most relevant images are input into the LMM along with the original question q to generate a specific answer.

General-question manual review.

After the initial LLM filtering, we conduct a manual review of the questions that were classified as specific. This manual process involves two key steps to ensure answer uniqueness and benchmark quality.

In the first step, we examine each question to confirm it contains unique identifiers—such as names, dates, titles, or other specific attributes—suggesting a document-specific answer. This careful check helps identify questions with clear, unique markers that direct the retrieval process to a single document.

In the second step, we verify the uniqueness of each answer to eliminate any remaining ambiguity. Although specific identifiers are present, questions may still be prone to ambiguity, such as with common names or recurring book titles. To address this, we employ a refined verification process. First,

we use Optical Character Recognition (OCR)^[43] to extract all text from images in the dataset. We then search for occurrences of the unique identifiers retained from the first step across other documents. If matches are found, a manual review is conducted to ensure no alternative valid answers exist. This comprehensive approach minimizes the possibility of a single question mapping to multiple answers, enhancing the precision and reliability of our benchmarks.

Generic-knowledge filtering.

In DocVQA and InfographicVQA tasks, certain questions—such as “*How many sports were in the 2008 Beijing Paralympic Games?*”—can be answered based on general knowledge accessible to a large language model, without relying on the image content. This introduces a language bias when using LMMs for visual question answering, as it shifts the focus away from image-based reasoning. To address this, we filter out these general-knowledge questions, ensuring that evaluation emphasizes vision-based understanding and that models rely primarily on visual content to generate accurate answers.

To implement this, we developed an LLM-based evaluation pipeline that detects and excludes such questions. For each question, we prompt an LLM with “*{question}, answer briefly.*”. After receiving a response, we compare it to the ground-truth answer using another LLM. If the response matches the ground truth, we classify the question as general knowledge-related and remove it, thereby isolating questions that truly require visual document understanding. As shown in Table 1, GPT-4o accurately answers 26.4% of DocVQA questions and 54.9% of InfographicVQA questions directly, a rate significantly higher than that of open-source LLMs. Therefore, we select GPT-4o to filter out the questions that can be directly answered by the GPT-4o model. Overall, this process is to ensure that the evaluation reflects the necessity of vision-based comprehension.

	GPT-4o	LLaVA-OneVision	Qwen2-VL
DocVQA	26.4%	4.7%	3.4%
InfographicVQA	54.9%	13.4%	11.3%

Table 1. Percentage of questions answerable by LMMs without vision input. We evaluate GPT-4o, LLaVA-Onevision, and Qwen2-VL on their ability to answer questions directly from our dataset without requiring vision input. The reported percentage reflects the proportion of examples that can be answered solely through language understanding.

Final dataset profile.

After a rigorous three-stage data filtering process, we retained 109 questions from DocVQA and 155 questions from InfographicVQA, associated with 59 and 66 documents that provide the evidence, respectively. To assess retrieval performance at scale, we introduce two benchmarks: DocHaystack-1000 and InfoHaystack-1000, where each question requires retrieving relevant content from a set of 1,000 documents. Given the challenge this scale presents to current LMMs, particularly in terms of context length limitations, we also construct two smaller benchmarks: DocHaystack-100/200 and InfoHaystack-100/200. These benchmarks allow direct input of all associated images into the context, enabling evaluation of models' long-context comprehension ability. For training set, we also construct a dataset comprising 2,835 questions similarly, with 899 from DocVQA and 1,936 from InfographicVQA, to support robust learning and generalization for the multi-image reasoning tasks.

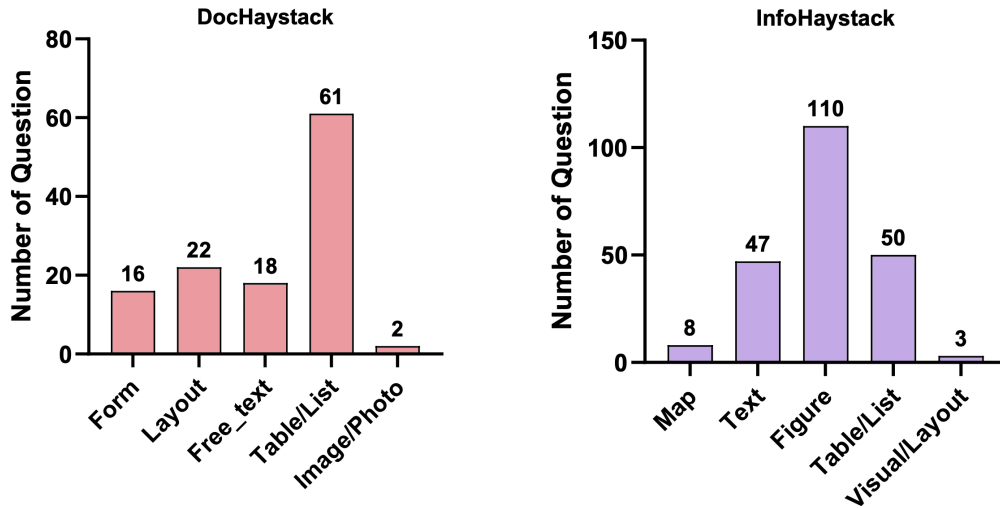


Figure 4. Question type analysis. We analyze the distribution of question types of DocHaystack and InfoHaystack. Each benchmark categorizes the data into 5 different types.

Question type analysis.

The types of questions represent the types of the evidence required for accurate answers. In Figure 4, we illustrate the distribution of question types across our dataset to provide insights into its structure. Following the classification system used in DocVQA and InfographicVQA, we categorize questions accordingly (note that a single question may fall into multiple categories). As shown in the figure, the DocHaystack benchmark places a greater emphasis on Table/List and Layout understanding, whereas InfoHaystack primarily targets Figure, Text, and Table comprehension.

4. Method

Current large multimodal models (LMMs) face substantial challenges when reasoning across hundreds or thousands of images, due not only to context length limitations but also to the inherent complexity of the task. This issue is particularly pronounced in our benchmarks, which contain 1k document files requiring high-resolution images to capture and interpret small-font text effectively. To enable LMMs to perform reasoning over a substantial number of documents, we introduce a vision-centric retrieval-augmented generation (V-RAG) framework. V-RAG efficiently retrieves a reduced set of relevant documents, allowing the LMM to focus on a manageable subset for deeper understanding, as illustrated in Figure 3. In the following section, we provide a detailed description of the V-RAG pipeline.

Task definition.

Given a question q and a collection of N documents $\mathcal{D} = \{D_1, \dots, D_N\}$, the V-RAG framework aims to retrieve the top- k most relevant documents to support LMMs understanding and answering the question q . V-RAG accomplishes this through a two-step retrieval process designed to effectively identify and rank relevant documents for each question.

Vision encoder ensemble.

Document files often contain a mix of text, symbols, and visual elements across various scales, requiring vision encoders to capture a comprehensive understanding of these complex structures. To efficiently handle this diversity, we represent each document as an image and utilize an ensemble of vision encoders, including CLIP^[32], SigLIP^[33], and OpenCLIP^[34], each bringing distinct strengths to the image understanding, as depicted in Figure 3. For example, the ConvNext encoder^[44] from OpenCLIP^[34] is particularly effective for high-resolution image encoding. We compute the similarity score between each question q and all documents in the document set \mathcal{D} according to Equation 1, with similarity scores from each encoder represented as Sim_c , Sim_o , and Sim_s respectively.

$$\mathcal{S}(q, \mathcal{D}) = \cos(\phi_t(q), \phi_v(D_j)) \mid D_j \in \mathcal{D}, \quad (1)$$

where \mathcal{S} denotes the computing the similarity between the query q and a collection of documents \mathcal{D} . \cos denotes the cosine similarity. ϕ_t denotes the text encoder, and ϕ_v denotes the vision encoder.

To derive a final relevance score, we calculate the average similarity Sim_{avg} for each question-image pair by combining Sim_c , Sim_o , and Sim_s . We then rank the images based on Sim_{avg} in descending order, selecting the top- m most relevant images according to their similarity scores.

	DocHaystack-100			DocHaystack-200			DocHaystack-1000		
	R@1	R@3	R@5	R@1	R@3	R@5	R@1	R@3	R@5
BM25 (OCR)	63.30	75.23	79.82	65.14	71.56	75.23	56.88	66.06	69.72
Jina-CLIP ^[45]	16.51	31.19	41.28	9.17	24.77	30.28	3.67	7.34	12.84
Nomic-Embed-Vision ^[46]	16.51	24.77	28.44	13.76	21.10	25.69	1.83	2.75	6.42
CLIP ^[32]	46.79	65.14	69.72	44.04	58.72	65.14	23.85	41.28	45.87
SigLIP ^[33]	51.38	67.89	76.15	47.71	63.30	70.64	33.03	49.54	57.80
OpenCLIP ^[34]	58.72	75.23	79.82	56.88	70.64	75.23	34.86	49.54	57.80
V-RAG (ours)	81.65	88.99	88.99	77.98	84.40	84.40	66.06	77.98	78.90
	InfoHaystack-100			InfoHaystack-200			InfoHaystack-1000		
	R@1	R@3	R@5	R@1	R@3	R@5	R@1	R@3	R@5
BM25 (OCR)	56.77	65.81	70.97	51.61	65.16	69.03	38.71	51.61	58.06
Jina-CLIP	43.23	51.61	58.06	36.77	46.45	51.61	23.87	33.55	37.42
Nomic-Embed-Vision	34.84	50.32	56.77	30.97	43.23	48.39	20.65	30.97	35.48
CLIP	69.68	78.71	85.81	65.16	77.42	81.94	45.81	64.52	70.32
SigLIP	58.06	71.61	80.00	55.48	67.74	76.77	39.35	55.48	61.94
OpenCLIP	72.26	85.16	92.90	66.45	81.94	89.03	53.55	65.81	72.90
V-RAG (ours)	79.35	90.97	92.90	74.84	88.39	88.39	64.52	74.19	78.06

Table 2. Retrieval Results. We compare our V-RAG model with other text-to-image and text-to-text (using OCR) retrieval methods across both benchmarks. V-RAG consistently outperforms baseline models on Recall@1, Recall@3, and Recall@5 metrics. Notably, V-RAG leverages an ensemble of text-to-image models along with a large multimodal model in a two-stage filtering approach. Top-performing values in each column are highlighted in bold.

LMM-filter module. To refine the selection of top- m relevant images further, we introduce a LMM-based

question-image relevance assessment module. This module evaluates the relevance between each question and the top- m images identified in the first filtering step. Specifically, we pair each image with the question text and input them into an open-source vision-language model, prompting, “Can this image provide answers to this question? Respond only with yes or no”. We only retain the question-image pairs that are identified as “yes” from LMM, and remove other irrelevant images.

LMM-VQA module. Achieving high top-1 ranking accuracy in image retrieval is challenging, so we retain the top- k images from the LMM-filtered ranking list and present them to the LMM-VQA to improve the likelihood of including relevant images. We input these top- k images alongside the question into the LMM-VQA (see Figure 3), which then generates the answer directly. To enhance robustness against visual distractors, the LMM-VQA can be further optimized, as analyzed in the experiment section.

5. Experiments

In the experiments section, we will primarily describe our training setup, covering evaluation metrics, baseline models, and the fine-tuning procedure for the LMM-VQA model. We also present the main experimental results along with an ablation study to provide further insights.

5.1. Training setup

Metric.

In our evaluation of the DocHaystack and InfoHaystack benchmarks, we employ a model-based assessment by leveraging GPT-4o-mini^[1] to accurately determine whether the model predictions match target answers. This method uses a carefully structured prompt to facilitate GPT-4o-mini’s evaluation of answer correctness. We empirically found that the model-based evaluation achieves higher consistency and alignment with human judgment. Additional details on the prompt design are provided in the Appendix.

For the document retrieval evaluation, we report the baseline results using Recall@1, Recall@3, and Recall@5 metrics. These metrics enable a thorough assessment of retrieval accuracy across varying levels of precision.

Model	DocHaystack			InfoHaystack		
	100	200	1000	100	200	1000
LLaVA-OV ^[4]	-	-	-	-	-	-
GPT-4o ^[1]	27.52	23.85	-	23.87	20.00	-
Gemini ^[2]	50.46	48.62	-	29.03	21.94	-
Qwen2-VL ^[3]	41.28	12.84	-	20.00	14.19	-
MIRAGE ^[5]	3.67	3.67	2.75	7.74	7.10	6.45
LLaVA-OV+V-RAG	69.72	65.14	55.05	43.22	41.94	36.77
GPT-4o+V-RAG	81.65	72.48	66.97	65.16	63.23	56.77
Gemini+V-RAG	73.39	65.14	58.72	57.42	57.42	47.10
Qwen2-VL+V-RAG	82.57	74.31	66.06	65.81	65.81	60.00
Qwen2-VL-f.t.+V-RAG	86.24	79.82	73.39	67.10	67.74	60.00

Table 3. The VQA results for the DocHaystack and InfoHaystack. We evaluate with many closed-source and open-source multimodal model, and also integrating them with our V-RAG retrieval framework. - denotes that those models can not be inferred due to their token context constraints. To enable GPT-4o and Qwen2-VL to process hundreds of images, we employ low-resolution mode and adjust image size for compatibility.

Baselines.

In our experiment, we have evaluated several open and closed-sourced vision-language models on the retrieval and VQA performance. For the large multimodal model, we used the gpt-4o-2024-08-06 version of GPT-4o^[1], the LLaVA-OneVision-Qwen2-7b-OV-HF version of LLaVA-OneVision^[4], and the Qwen2-VL-7B-Instruct version of Qwen2-VL^[27]. For computing the text-to-image similarities, we employed the Jina-CLIP-v1^[45] variant, Nomic-Embed-Vision-v1.5^[46] variant, CLIP^[32] ViT-L/14@336 variant, for SigLIP^[33], the ViT-SO400M/14@384 variant, and for OpenCLIP^[34], the ConvNeXt-XXL@1024 variant as well as text-based method, BM25.

In our V-RAG setting, we apply LLaVA-OneVision-Qwen2-7b-OV-HF for the LMM-filter module and Qwen2-VL-7B-Instruct for the LMM-VQA module. We select m as 60 and k as 5 in our experiment.

Optimizing the LMM-VQA module.

To improve the robustness of the LMM-VQA model in handling visual question answering with multiple distractor images, we further fine-tune the model using our curated training data.

During this fine-tuning process, we introduce 1–10 randomly sampled distractor images for each question, creating a challenging setting that encourages the model to focus on relevant content amid a mix of positive and negative images. The fine-tuning is conducted with a batch size of 32 and a peak learning rate of $1e-4$ over a single epoch. Additionally, we leverage LoRA^[47] with a rank of 8 to efficiently adapt the model's parameters during training.

5.2. Main Experimental Results

We evaluated a range of open-source and closed-source vision-language models for VQA tasks. We also evaluate several text-to-image and text-to-text (with OCR) retrieval models to evaluate their retrieval capabilities on our benchmarks. More detailed performance analysis are described in the following sections.

Retrieving results.

The retrieval results in Table 2 demonstrate the superiority of our proposed V-RAG framework over several baseline methods across both DocHaystack and InfoHaystack benchmarks. V-RAG consistently achieves the highest Recall@1, Recall@3, and Recall@5 scores on most categories, indicating its robust retrieval capabilities. Notably, V-RAG outperforms text-based retrieving models such as BM25 and also the text-to-image retrieval models like jina-clip, CLIP, SigLIP, and OpenCLIP by substantial margins, especially on the DocHaystack-100 subset, where it reaches Recall@1 of 81.65% and Recall@5 of 88.99%. This pattern continues for larger datasets (DocHaystack-1000), where V-RAG remains competitive, achieving Recall@1 of 66.06%. It achieves the top performance across all recall metrics on DocHaystack. For InfoHaystack benchmarks, V-RAG also outperforms other models, particularly on InfoHaystack-100 and InfoHaystack-200, where it receives Recall@1 of 74.84% and 64.52%, higher than previous best by 8% and 11%, respectively. This consistent performance advantage highlights the effectiveness of V-RAG's

ensemble of multiple vision encoders, allowing it to capture more granular details and improve retrieval accuracy over large multimodal models.




Figure 5. Top-k selection ablation analysis for LMM-VQA. We demonstrate the results for LLaVA, Qwen2-VL, GPT-4o and also the finetuned Qwen2-VL model on the DocHaystack-100/1000 and InfoHaystack-100/1000 benchmarks. All the models are integrated with our V-RAG framework. We show the VQA accuracy performance for each ablation.

CLIP	SigLIP	OpenCLIP	VLM-filter	DocHaystack-1000			InfoHaystack-1000		
				R@1	R@3	R@5	R@1	R@3	R@5
✓	✗	✗	✗	23.85	41.28	45.87	45.81	64.52	70.32
✗	✓	✗	✗	33.03	49.54	57.80	39.35	55.48	61.94
✗	✗	✓	✗	34.86	49.54	57.80	53.55	65.81	72.90
✓	✓	✗	✗	40.37	59.63	62.39	59.35	67.74	74.19
✓	✓	✓	✗	42.20	66.06	77.48	56.13	70.97	78.06
✓	✓	✓	✓	66.06	77.98	78.90	64.52	74.19	78.06

Table 4. Ablation study on the V-RAG framework components. We quantify the impact of each module for the Recall@1, Recall@3 and Recall@5 retrieval performance on the DocHaystack-1000 and InfoHaystack-1000 for our V-RAG framework.

Visual question answering (VQA) results.

The table presents VQA results for the DocHaystack and InfoHaystack benchmarks across varying dataset sizes (100, 200, 1000) using different multimodal models, both independently and in combination with the V-RAG framework. The results show that Qwen2-VL fine-tuned with V-RAG (Qwen2-VL-ft.+V-RAG) achieves the highest scores across most benchmarks, with particularly notable performance on DocHaystack-100 (86.24) and InfoHaystack-100 (67.10), indicating superior retrieval and VQA capabilities in these scenarios. When V-RAG is added to other models, substantial improvements are observed, demonstrating the framework's efficacy in enhancing retrieval accuracy. For instance, GPT-4o's performance increases significantly with V-RAG, particularly for DocHaystack-100 and -200. The analysis highlights that V-RAG integration generally boosts performance across models, with Qwen2-VL-ft.+V-RAG standing out as the top performer on both benchmarks, especially for the larger 1000-document tasks where retrieval accuracy is more challenging. This suggests that V-RAG's vision-centric, retrieval-augmented approach is highly effective for large-scale multimodal document understanding.

The table also shows that the DocHaystack-1000 and InfoHaystack-1000 present significant challenges for current LMMs. The drop in performance for larger document sets, with top accuracy only reaching 73.39% for DocHaystack-1000 and 60.00% for InfoHaystack-1000, underscores the difficulty our benchmarks.

5.3. Ablation Studies

Ablation study on Top-k Selection.

This figure presents the top-k selection ablation analysis for LMM-VQA across four models: LLaVA-OV, Qwen2-VL, GPT-4o, and the fine-tuned Qwen2-VL (Qwen2-VL-ft.), evaluated on the DocHaystack-100/1000 and InfoHaystack-100/1000 benchmarks. The analysis reports VQA accuracy as a function of top-k selection (Top 1, Top 3, and Top 5). Overall, accuracy tends to improve with larger k-values, suggesting that offering more retrieval options positively impacts model performance. However, for LLaVA-OV, there is a marked decrease in performance at top-5, indicating that this model struggles to process multiple images at this scale.

Ablation study on the V-RAG framework components.

The ablation study in Table 4 highlights the contributions of each component in the V-RAG framework on the DocHaystack-1000 and InfoHaystack-1000 benchmarks. Using CLIP alone yields low performance (e.g., Recall@1 of 23.85% on DocHaystack-1000 and 45.81% on InfoHaystack-1000), indicating its limited retrieval capability on its own. Adding SigLIP and OpenCLIP incrementally improves results.

The highest performance is achieved when all three encoders are combined with the VLM-filter module, leading to Recall@1 scores of 66.06% on DocHaystack-1000 and 64.52% on InfoHaystack-1000. This setup also achieves the top Recall@1, Recall@3 and Recall@5 values, demonstrating that the VLM-filter is essential for refining the ensemble outputs and significantly improving retrieval accuracy. These results confirm that each module contributes to V-RAG's overall effectiveness.

6. Conclusion

In this work, we introduced the DocHaystack and InfoHaystack benchmarks to evaluate LMMs for retrieving and reasoning across large-scale documents. Our benchmarks providing a more rigorous and realistic assessment of large multimodal models in real-world, large-scale retrieval scenarios. To tackle these challenges, we proposed V-RAG, a vision-centric retrieval-augmented generation framework that significantly enhances retrieval precision and overall VQA performance. V-RAG achieves this through an ensemble of vision encoders and a specialized relevance filtering module, enabling improved accuracy across diverse visual inputs. Experimental results indicate that integrating V-RAG enables both open-source and closed-source LMMs to achieve superior performance in large-scale image retrieval and complex reasoning tasks.

References

1. ^{a, b, c, d, e, f, g}OpenAI (2024). "GPT-4o: Enhanced Multimodal Language Model". OpenAI Research. Available from: <https://openai.com/index/hello-gpt-4o/>.
2. ^{a, b, c, d}Google AI. "Gemini: Google's Multimodal AI Model". Google AI Research. 2024. Available from: <https://fireflies.ai/blog/gemini-vs-gpt-4>.
3. ^{a, b, c}Wang P, Bai S, Tan S, Wang S, Fan Z, Bai J, Chen K, Liu X, Wang J, Ge W, et al. Qwen2-VL: Enhancing Vision-Language Model's Perception of the World at Any Resolution. *arXiv preprint arXiv:2409.12191*. 2024.

4. ^aLi B, ^bZhang Y, ^cGuo D, ^dZhang R, ^eLi F, ^fZhang H, ^gZhang K, ^hZhang P, ⁱLi Y, ^jLiu Z, ^kLi C (2024). "LLaVA-One Vision: Easy Visual Task Transfer". *arXiv preprint arXiv:2408.03326*.
5. ^aWu TH, ^bBiamby G, ^cQuenum J, ^dGupta R, ^eGonzalez JE, ^fDarrell T, ^gChan DM (2024). "Visual Haystacks: A Vision-Centric Needle-In-A-Haystack Benchmark". *arXiv preprint arXiv:2407.13766*.
6. ^aPenamakuri AS, ^bGupta M, ^cDas Gupta M, ^dMishra A. Answer mining from a pool of images: towards retrieval-based visual question answering. *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*. 2023:1052-1058.
7. ^aTalmor A, ^bMin S, ^cJia R, ^dElazar Y, ^eHasson US, ^fChen D (2021). "MultimodalQA: Complex Question Answering over Text, Tables, and Images". In: *Proceedings of the International Conference on Learning Representations (ICLR)*.
8. ^aMathew M, ^bKaratzas D, ^cJawahar CV (2021). "DocVQA: A Dataset for VQA on Document Images". In: *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pp. 2200–2209.
9. ^aMathew M, ^bBagal V, ^cPérez Tito R, ^dKaratzas D, ^eValveny E, ^fJawahar CV. InfographicVQA. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2021. p. 4390–4399.
10. ^aChang Y, ^bNarang M, ^cSuzuki H, ^dCao G, ^eGao J, ^fBisk Y (2022). "WebQA: Multihop and Multimodal QA". *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 2022: 14178–14188.
11. ^aGoyal Y, ^bKhot T, ^cSummers-Stay D, ^dBatra D, ^eParikh D. Making the V in VQA matter: Elevating the role of image understanding in Visual Question Answering. In: *Conference on Computer Vision and Pattern Recognition (CVPR)*; 2017.
12. ^aBordes F, ^bPang RY, ^cAjay A, ^dLi AC, ^eBardes A, ^fPetryk S, ^gMañas O, ^hLin Z, ⁱMahmoud A, ^jJayaraman B, et al. An introduction to vision-language modeling. *arXiv preprint arXiv:2405.17247*. 2024.
13. ^aFeng CM, ^bBai Y, ^cLuo T, ^dLi Z, ^eKhan S, ^fZuo W, ^gXu X, ^hGoh RSM, ⁱLiu Y (2023). "VQA4CIR: Boosting Composed Image Retrieval with Visual Question Answering". *arXiv preprint arXiv:2312.12273*. Available from: <https://arxiv.org/abs/2312.12273>.
14. ^aHudson DA, ^bManning CD (2019). "GQA: A New Dataset for Real-World Visual Reasoning and Compositional Question Answering". *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pages 6700–6709.
15. ^aKrishna R, ^bZhu Y, ^cGroth O, ^dJohnson J, ^eHata K, ^fKravitz J, ^gChen S, ^hKalantidis Y, ⁱLi L-J, ^jShamma DA, ^kBernstein MS, ^lFei-Fei L (2017). "Visual Genome: Connecting Language and Vision Using Crowdsourced Dense Image Annotations". *International Journal of Computer Vision*. 123 (1): 32–73.

16. [△]Masry A, Long DX, Tan JQ, Joty S, Hoque E (2022). "ChartQA: A Benchmark for Question Answering about Charts with Visual and Logical Reasoning". *Findings of the Association for Computational Linguistics: ACL 2022*. pp. 2263–2279.
17. [△][♢]Lu P, Bansal H, Xia T, Liu J, Li C, Hajishirzi H, Cheng H, Chang KW, Galley M, Gao J (2023). "MathVista: Evaluating Mathematical Reasoning of Foundation Models in Visual Contexts". *arXiv preprint arXiv:2310.02255*. Available from: <https://arxiv.org/abs/2310.02255>.
18. [△]Wang K, Pan J, Shi W, Lu Z, Zhan M, Li H (2024). "Measuring Multimodal Mathematical Reasoning with MATH-Vision Dataset". *arXiv preprint arXiv:2402.14804*.
19. [△][♢]Zhang R, Jiang D, Zhang Y, Lin H, Guo Z, Qiu P, Zhou A, Lu P, Chang KW, Gao P, Li H (2024). "MathVerse: Does Your Multi-modal LLM Truly See the Diagrams in Visual Math Problems?" *arXiv preprint arXiv:2403.14624*.
20. [△]Marino K, Rastegari M, Farhadi A, Mottaghi R (2019). "OK-VQA: A Visual Question Answering Benchmark Requiring External Knowledge". *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. pages 3195--3204.
21. [△]Schwenk D, Khandelwal A, Clark C, Marino K, Mottaghi R (2022). "A-OKVQA: A Benchmark for Visual Question Answering using World Knowledge". In: *European Conference on Computer Vision (ECCV)*. pp. 1–17.
22. [△][♢]Lu P, Xia T, Shi W, El Kholy A, Lin XV, Gao J, Chen X, Chang KW (2022). "Learn to Explain: Multimodal Reasoning via Thought Chains for Science Question Answering". In: *Advances in Neural Information Processing Systems*.
23. [△][♢]Yue X, Ni Y, Zhang K, Zheng T, Liu R, Zhang G, Stevens S, Jiang D, Ren W, Sun Y, et al. MMMU: A Massive Multi-discipline Multimodal Understanding and Reasoning Benchmark for Expert AGI. *arXiv preprint arXiv:2311.16502*. 2024.
24. [△]OpenAI. "GPT-4V: Multimodal Language Model with Vision Capabilities". OpenAI Research. 2024. Available from: <https://openai.com/index/gpt-4/>.
25. [△][♢]Zhu D, Chen J, Shen X, Li X, Elhoseiny M. "MiniGPT-4: Enhancing Vision-Language Understanding with Advanced Large Language Models." In: *International Conference on Learning Representations (ICLR)*; 2024. Available from: <https://openreview.net/forum?id=1tZbq88f27>.
26. [△][♢]Chen J, Zhu D, Shen X, Li X, Liu Z, Zhang P, Krishnamoorthi R, Chandra V, Xiong Y, Elhoseiny M (2023). "MiniGPT-v2: Large Language Model as a Unified Interface for Vision-Language Multi-task Learning". *arXiv preprint arXiv:2310.15339*.

27. ^a ^bBai Y, Zhang Z, Zhang Y, Li Y, Zhang Y, Zhang Y, Zhang Y, Zhang Y, Zhang Y, Zhang Y, et al. Qwen-VL: A Frontier Vision-Language Model with Larger-scale Vision Pre-training and Aligned Cross-modal Instruction Tuning. *arXiv preprint arXiv:2310.06726*. 2023.
28. ^ΔWang K, Zhang Y, Li H (2023). "Mini-Gemini: An Efficient and Versatile Vision-Language Model". *arXiv preprint arXiv:2310.12345*. Available from: [arXiv:2310.12345](https://arxiv.org/abs/2310.12345).
29. ^a ^bBansal A, Sikka K, Sharma G, Chellappa R (2020). "Visual Question Answering on Image Sets". In: *Proceedings of the European Conference on Computer Vision (ECCV)*. pp. 35–51.
30. ^ΔWang W, Zhang S, Ren Y, Duan Y, Li T, Liu S, Hu M, Chen Z, Zhang K, Lu L, et al. Needle in a multimodal haystack. *arXiv preprint arXiv:2406.07230*. 2024.
31. ^ΔTanaka R, Nishida K, Nishida K, Hasegawa T, Saito I, Saito K (2023). "SlideVQA: A Dataset for Document Visual Question Answering on Multiple Images". *Proceedings of the AAAI Conference on Artificial Intelligence*. 37 (11): 13636–13645.
32. ^a ^b ^c ^d ^eRadford A, Kim JW, Hallacy K, Ramesh A, Goh G, Agarwal S, Sastry G, Askell A, Mishkin P, Clark J, et al. Learning transferable visual models from natural language supervision. In: *Proceedings of the 38th International Conference on Machine Learning*. PMLR; 2021. p. 8748–8763. Available from: <https://arxiv.org/abs/2103.00020>.
33. ^a ^b ^c ^dZhai X, Mustafa B, Kolesnikov A, Beyer L (2023). "Sigmoid Loss for Language Image Pre-Training". In: *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*. pp. 11975–11985.
34. ^a ^b ^c ^d ^eIlharco G, Wortsman M, Wightman R, Gordon C, Carlini N, Taori R, Dave A, Shankar V, Namkoong H, Miller J, Hajishirzi H, Farhadi A, Schmidt L. OpenCLIP. 2021. doi:[10.5281/zenodo.5143773](https://doi.org/10.5281/zenodo.5143773).
35. ^ΔBai Y, Xu X, Liu Y, Khan S, Khan F, Zuo W, Goh RSM, Feng CM. Sentence-level prompts benefit composed image retrieval. In: *Proceedings of the International Conference on Learning Representations (ICLR)*; 2024. Spotlight Presentation.
36. ^ΔAsai A, Wu Z, Wang Y, Sil A, Hajishirzi H (2023). "Self-RAG: Learning to Retrieve, Generate, and Critique through Self-Reflection". *arXiv preprint arXiv:2310.11511*.
37. ^ΔGuu K, Lee K, Tung Z, Pasupat P, Chang MW. "REALM: Retrieval-Augmented Language Model Pre-Training." In: *Proceedings of the 37th International Conference on Machine Learning (ICML)*. PMLR; 2020. p. 3929–3938.
38. ^ΔLuo H, Chuang YS, Gong Y, Zhang T, Kim Y, Wu X, Fox D, Meng H, Glass J (2023). "SAIL: Search-Augmented Instruction Learning". *arXiv preprint arXiv:2305.15225*. Available from: <https://arxiv.org/abs/2305.15225>.

39. ^ΔKarpukhin V, Oğuz B, Min S, Lewis P, Wu L, Edunov S, Chen D, Yih W (2020). "Dense Passage Retrieval for Open-Domain Question Answering". *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. Association for Computational Linguistics; 2020. p. 6769–6781.
40. ^ΔChen W, Hu H, Chen X, Verga P, Cohen WW. "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 (EMNLP)*. Association for Computational Linguistics; 2022. p. 5558–5570.
41. ^ΔDevlin J, Chang MW, Lee K, Toutanova K (2019). "BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding." In: *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. Association for Computational Linguistics. pp. 4171–4186. doi:[10.18653/v1/N19-1423](https://doi.org/10.18653/v1/N19-1423). Available from: <https://aclanthology.org/N19-1423>.
42. ^ΔRen S, He K, Girshick R, Sun J (2015). "Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks". *Advances in Neural Information Processing Systems*. 28. Available from: <https://arxiv.org/abs/1506.01497>.
43. ^ΔRay Smith et al. Tesseract OCR Engine [Internet]. 2024. Available from: <https://github.com/tesseract-ocr/tesseract>. Accessed: 2024-11-06.
44. ^ΔLiu Z, Mao H, Wu CY, Feichtenhofer C, Darrell T, Xie S (2022). "A ConvNet for the 2020s". In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 2022: 11976–11986. doi:[10.1109/CVPR52688.2022.01168](https://doi.org/10.1109/CVPR52688.2022.01168). [Link](#).
45. ^Δ^ΔKoukounas A, Mastrapas G, G\u00f6nther M, Wang B, Martens S, Mohr I, Sturua S, Akram MK, Fontanals Mart\u00e9nez J, Ognawala S, Guzman S, Werk M, Wang N, Xiao H. Jina CLIP: Your CLIP Model Is Also Your Text Retriever. 2024. Available from: [arXiv:2405.20204](https://arxiv.org/abs/2405.20204).
46. ^Δ^ΔNussbaum Z, Duderstadt B, Mulyar A. "Nomic Embed Vision: Expanding the Latent Space". 2024. Available from: [arXiv:2406.18587](https://arxiv.org/abs/2406.18587).
47. ^ΔHu EJ, Shen Y, Wallis P, Allen-Zhu Z, Li Y, Wang S, Wang L, Chen W (2021). "LoRA: Low-Rank Adaptation of Large Language Models". *arXiv*. [arXiv:2106.09685](https://arxiv.org/abs/2106.09685).

Declarations

Funding: No specific funding was received for this work.

Potential competing interests: No potential competing interests to declare.