

UNIFIED MULTI-MODAL INTERLEAVED DOCUMENT REPRESENTATION FOR INFORMATION RETRIEVAL

Anonymous authors

Paper under double-blind review

ABSTRACT

Information Retrieval (IR) methods aim to identify relevant documents in response to a given query, which have gained remarkable attention due to their successful application in various natural language tasks. However, existing approaches typically consider only the textual information within the documents, which overlooks the fact that documents can contain multiple modalities, including texts, images, and tables. Further, they often segment each long document into multiple discrete passages for embedding, preventing them from capturing the overall document context and interactions between paragraphs. We argue that these two limitations lead to suboptimal document representations for retrieval. In this work, to address them, we aim to produce more comprehensive and nuanced document representations by holistically embedding documents interleaved with different modalities. Specifically, we achieve this by leveraging the capability of recent vision-language models that enable the processing and integration of text, images, and tables into a unified format and representation. Moreover, to mitigate the information loss from segmenting documents into passages, instead of representing and retrieving passages individually, we further merge the representations of segmented passages into one single document representation, while we additionally introduce a reranking strategy to decouple and identify the relevant passage within the document if necessary. Then, through extensive experiments on diverse information retrieval scenarios considering both the textual and multimodal queries, we show that our approach substantially outperforms relevant baselines, thanks to the consideration of the multimodal information interleaved within the documents in a unified way.

1 INTRODUCTION

Information Retrieval (IR) is the task of fetching relevant documents from a large corpus in response to an input query, which becomes a fundamental process to various real-world applications including web search engines and question-answering systems (Shah et al., 2019; Lewis et al., 2020; Guu et al., 2020). Specifically, to retrieve documents for the query, traditional approaches have focused on their textual representations, utilizing either sparse retrieval methods such as TF-IDF and BM25 (Robertson et al., 1994; Jones, 2004), which rely on exact term matching between the query and document, or dense retrieval methods such as DPR and ANCE (Karpukhin et al., 2020; Xiong et al., 2021), which leverage neural embeddings of the query and document text to capture semantic similarities between them over a continuous vector space. Recently, dense retrieval methods have gained more popularity over sparse methods due to their capability to capture semantic nuances and context beyond simple keyword matching, leading to multiple successes with improved performance.

Despite their huge successes, existing (dense) retrieval methods face a couple of severe challenges. First, they primarily rely on the textual data for document embedding and retrieval, overlooking the fact that modern documents often contain multimodal content, such as images and tables (beyond the plain text), which can carry critical information that may be essential for accurately understanding and retrieving the relevant documents. To be specific, a diagram within a medical article can more effectively represent the structure of a molecule or the progression of a disease, offering more clarity that would be difficult to achieve with text alone, and omitting such multimodal content can lead to an incomplete understanding (and potentially inaccurate retrieval) of the documents. Also, the segmentation of long documents into discrete passages, which is commonly employed by retrieval models to handle the length limitation for embeddings, may prevent models from capturing the full

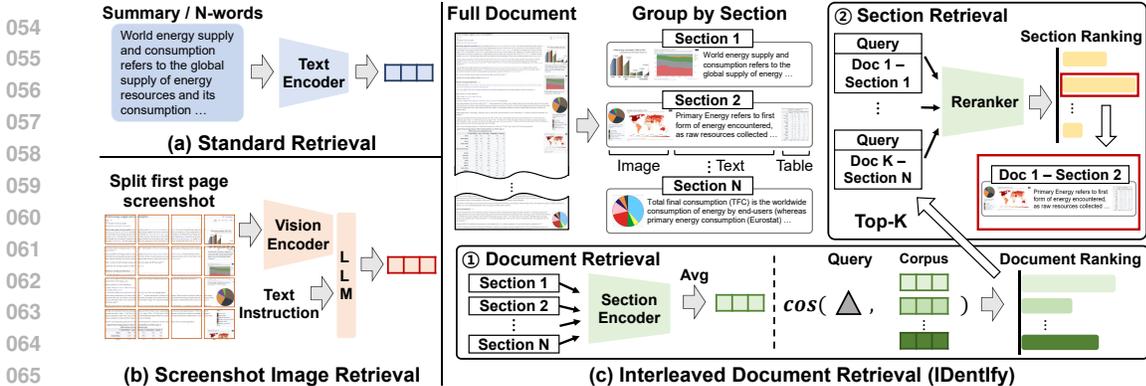


Figure 1: Comparison of different IR approaches. (a): Conventional methods use a small portion of the text within the document for its representation. (b): Recent methods use first-page screenshot images to represent the document. (c): Our approach leverages the full contextual information within documents interleaved with multiple modalities by considering them in their original format, and is capable of pinpointing relevant sections.

context and the intricate relationships between different parts of the document, ultimately leading to suboptimal retrieval performance. It is worthwhile noting that, concurrent to our work, while there has been recent work that screen captures the document and then embeds its screenshots (to consider different modalities in a unified format) (Faysse et al., 2024; Ma et al., 2024a), not only its content (such as paragraphs, images, and tables) can be fragmented into different sub-images, leading to the loss of contextual coherence across the entire document, but also the visual representation of text may hinder the model’s ability to capture the semantic relationships present in the original textual data, while increasing the image resolution leads to the critical concern on the memory requirements.

In this work, we introduce a novel approach to holistic document embedding for IR, which addresses the aforementioned challenges by representing and retrieving the documents interleaved with different modalities in a unified manner. Specifically, our method revolves around the recent advance of Vision-Language Models (VLMs), which enables the processing and integration of multimodal content (such as text, images, and tables) directly into a single token sequence, thereby preserving the context and relationships between various parts of the document, unlike the previous approaches that rely on the fragmented visual representations. Furthermore, in cases where the number of tokens in a document is large and exceeds the capacity of a single context window of VLMs, we propose a strategy to segment the document into manageable passages, each represented within the token limit, and combine these passage representations into a unified document representation, which differs from existing IR approaches that independently represent and retrieve at the passage level, potentially losing the overall document context. Lastly, to accurately identify only the relevant sections within the lengthy documents, we introduce a reranking mechanism that is trained to pinpoint the passage most pertinent to the query (among all the other passages within the document), effectively allowing for both the coarse-grained document-level matching and the fine-grained passage-level retrieval. We provide the visual illustrations of the overall pipeline of IDentify against prior work in Figure 1. We refer to our overall method as **Interleaved Document Information Retrieval System (IDentify)**.

We experimentally validate the effectiveness of IDentify on four different benchmark datasets, considering both the text-only and multimodal queries. On a battery of tests conducted, we then observe that our approach substantially outperforms relevant baselines that consider only the uni-modality for document representations, thanks to the consideration of multimodal content. Further, we find that the strategy to represent the whole document with its single representation (by merging embeddings of its splits if necessary) is superior to the approach of individually representing them for document retrieval, but also performing reranking over the sections of the retrieved document is superior to the approach of directly retrieving those sections, which confirm the efficacy of the proposed retrieval and reranking pipeline for document and passage retrieval, respectively.

2 RELATED WORK

Information Retrieval Information Retrieval (IR) is the task of accurately finding documents relevant to a given query from a large corpus, such as Wikipedia, which has been a crucial component for a variety of applications, including search engines, question-answering systems, and conversational agents (Zhu et al., 2023; Gao et al., 2023; Ram et al., 2023; Shi et al., 2024; Jeong et al., 2024a).

Specifically, to retrieve the relevant documents, earlier IR approaches measured similarity between queries and documents based on their lexical term matching, such as BM25 and TF-IDF (Robertson et al., 1994; Jones, 2004). Yet, these methods often struggled to capture the semantic nuances beyond surface-level term overlaps. To overcome this, along with advancements in language models (Devlin et al., 2019; Liu et al., 2019), there has been dense retrieval approaches that embed both the queries and documents into a shared dense vector space (Karpukhin et al., 2020; Xiong et al., 2021), enabling the calculation of semantic similarity between them more effectively by capturing the deeper contextual information. However, previous IR studies have mainly focused on enhancing the textual representations of queries and documents, while overlooking the fact that documents often consist of diverse modalities (such as images and tables) beyond text, which can potentially provide richer context and aid in more accurate retrieval (Liu et al., 2021; Jeong et al., 2024b).

Multimodal Information Retrieval Recent studies in IR have expanded the focus from purely text-based retrieval models to those that consider other modalities, such as images (Radford et al., 2021; Xiao et al., 2024), tables (Herzig et al., 2021; Chen et al., 2024) and graphs (Baek et al., 2023); however, the majority of these approaches (Zhou et al., 2024; Long et al., 2024; Lerner et al., 2024; Nowak et al., 2024; Caffagni et al., 2024) have primarily explored how to process the multimodal *queries*, meanwhile, they often overlook the equally important multimodal characteristics of the *documents* being retrieved. Specifically, we argue that, while incorporating multimodal elements in queries has expanded the range and diversity of query types that IR systems can handle, considering the multimodal nature of the documents can lead to a more holistic representation of retrieval targets, which can ultimately lead to enhancing the overall retrieval performance. In efforts to handle diverse multimodal elements within documents, there are concurrent studies that have proposed to capture screenshots of documents, such as PDFs (Faysse et al., 2024) or Wikipedia web pages (Ma et al., 2024a), and subsequently encoding them through vision models (Ding et al., 2024). However, these methods are not only limited by factors, such as image resolution and computational memory, constraining their application to documents longer than a single page¹, but also fall short by treating the diverse modalities within a document as a single visual entity, leading to suboptimal document representations that fail to effectively capture the nuanced interdependence between text and images. Furthermore, they do not address the critical issue of splitting documents into smaller fragments (e.g., sub-images), which may disrupt the holistic contextual view of the entire document.

Vision-Language Models Recent Vision-Language Models (VLMs) have emerged as a powerful tool for jointly processing visual and textual data, combining the image understanding capabilities of visual encoders (Radford et al., 2021; Zhai et al., 2023) with the advanced reasoning abilities of language models (OpenAI, 2022; 2023a). These models have achieved remarkable performance across diverse vision-language (VL) tasks (such as image captioning and visual question answering) (Dai et al., 2023; OpenAI, 2023b), with the substantially limited attention on their applications to IR. We note that the latest developments in this field have particularly focused on enabling VLMs to handle interleaved, multimodal content, which involves a mixed sequence of images and text (Zhang et al., 2023; Li et al., 2024). In particular, LLaVA-NeXT-Interleave (Li et al., 2024) introduces a fine-tuning approach that specifically enhances the VLMs’ capacity to understand complex interleavings of multiple images and text within a single context. Drawing inspiration from these advances, in this work, we propose to harness the capabilities of VLMs to create unified embeddings for documents interleaved with text and images (as well as tables) for IR, which is a big shift from even the recent IR approach (Ma et al., 2024b) that still embeds the documents with the recent but text-based models like Llama (Touvron et al., 2023a;b), failing to fully capture the diverse multimodal content.

3 METHOD

We present IDentify to holistically represent documents interleaved with multimodal elements.

3.1 PRELIMINARY

We begin with preliminaries, formally explaining information retrieval and vision-language models.

Information Retrieval Recall that Information Retrieval (IR) is the task of searching for relevant documents from a large corpus in response to a given query. Formally, let q denote a query, d denote

¹For instance, Ma et al. (2024a) requires processing 9.8k image tokens just to process a single-page document, and it results in 2TB of memory for handling the entire Wikipedia corpus, which is not much practical.

a document, and \mathcal{D} denote a collection of documents ($\mathbf{d} \in \mathcal{D}$), where each query and document can be represented as a sequence of tokens: $\mathbf{q} = [q_1, q_2, \dots, q_n]$ and $\mathbf{d} = [d_1, d_2, \dots, d_m]$ where $[\cdot]$ indicates a concatenation operation in a specific order. We note that traditional IR approaches typically consider these tokens as purely textual elements; however, in this work, we propose to extend this assumption to have the tokens of both the textual and visual content, to capture the multimodal nature of many real-world documents. Then, this new extension raises important questions of how can both the textual and visual content be represented within a unified token framework, and how can these multimodal tokens be seamlessly integrated and encoded for document representations. To answer those two questions, we harness the power of recent vision-language models below.

Vision-Language Models We now turn to describing Vision-Language Models (VLMs), which are designed to jointly encode the textual and visual information in a unified token framework. We note that these models are generally comprised of two main components: a visual encoder and a language model, interconnected through a projection layer. Specifically, given an input document that may contain interleaved modalities (e.g., text and images), the visual encoder extracts high-level visual features from (multiple) images embedded within the document, mapping them into a latent space. Then, these visual features are transformed into a sequence of visual tokens via the projection layer, represented as follows: $\mathbf{V} \in \mathbb{R}^{V \times d_{\text{emb}}}$ where V denotes the visual token length and d_{emb} is the token dimension size. Similarly, for the textual content embedded within the document alongside images, the language model uses a word embedding layer to convert the input text into a sequence of text tokens, represented as follows: $\mathbf{L} \in \mathbb{R}^{L \times d_{\text{emb}}}$ where L denotes the token length of text.

In this work, we also propose to account for tables that are an integral modality for holistically representing the full content of documents. However, in contrast to text and images that have dedicated processing layers within the VLM architectures, tables do not have a specific representation layer. Nevertheless, we argue that recent VLMs are pre-trained on diverse web data, and subsequently they are implicitly learned to handle the table structures formatted in HTML. Consequently, we treat HTML-format table data as a linearized sequence of HTML words, applying the same word embedding layer as is used for plain text. To be formal, this process converts the table content into table tokens, as follows: $\mathbf{T} \in \mathbb{R}^{T \times d_{\text{emb}}}$ where T is the token length of the table. Lastly, once extracted, the visual tokens, text tokens, and table tokens are concatenated (to form a unified token sequence) and then passed through the remaining layers of VLMs, to capture both uni- and cross-modal relationships across different modalities, enabling the comprehensive understanding of the input document.

3.2 RETRIEVER

We now turn to explaining how we design a retriever specifically tailored for multimodal interleaved document retrieval. In particular, to effectively retrieve documents that contain multiple modalities, our approach leverages a VLM capable of processing text, images, and tables within a single document. Further, following the standard practice of existing retrieval architectures (Karpukhin et al., 2020; Xiong et al., 2021), we use a dual-encoder structure, which consists of a query encoder and a section encoder, both are based on the VLM, which is illustrated in Figure 2 (a).

Specifically, thanks to the use of the VLM, our query encoder can take either purely textual queries $\mathbf{q} = \mathbf{L}_Q$ or multimodal queries consisting of text and corresponding visual elements $\mathbf{q} = [\mathbf{V}_Q, \mathbf{L}_Q]$. Also, to obtain the final query representation, we introduce a learnable token called ‘End of Query’, $[\text{E} \circ \text{Q}] \in \mathbb{R}^{d_{\text{emb}}}$. This token is appended to the end of the sequence of query tokens \mathbf{q} , and the final concatenated tokens $[\mathbf{q}, [\text{E} \circ \text{Q}]]$ are then passed through the query encoder. Then, the model output corresponding to $[\text{E} \circ \text{Q}]$ is used as the final query representation, as follows: $\mathbf{Z}_Q \in \mathbb{R}^{d_{\text{emb}}}$.

For documents, we first represent each document \mathbf{d} as a sequence of sections $\mathbf{d} = [\mathbf{s}_i]_{i=1}^S$ (with a total of S sections), where each section \mathbf{s}_i is derived by dividing the document according to the subtitles in the document. \mathbf{s}_i can contain a combination of text tokens \mathbf{L}_{S_i} , visual tokens from embedded images \mathbf{V}_{S_i} , and table tokens \mathbf{T}_{S_i} , denoted as follows: $\mathbf{s}_i = [\mathbf{V}_{S_i}, \mathbf{L}_{S_i}, \mathbf{T}_{S_i}]$. Then, to obtain a section-level representation, similar to the query representation, we introduce a learnable token, called ‘End of Section’: $[\text{E} \circ \text{S}] \in \mathbb{R}^{d_{\text{emb}}}$, which is similarly appended at the end of each section. We then forward the concatenated tokens $[\mathbf{s}_i, [\text{E} \circ \text{S}]]$ to the section encoder, and, after that, the output corresponding to $[\text{E} \circ \text{S}]$ is used to form the section representation, as follows: $\mathbf{Z}_{S_i} \in \mathbb{R}^{d_{\text{emb}}}$. Additionally, the overall document representation is obtained by averaging the representations of all sections within the document, defined as follows: $\mathbf{Z}_D = \frac{1}{S} \sum_{i=1}^S \mathbf{Z}_{S_i}$.

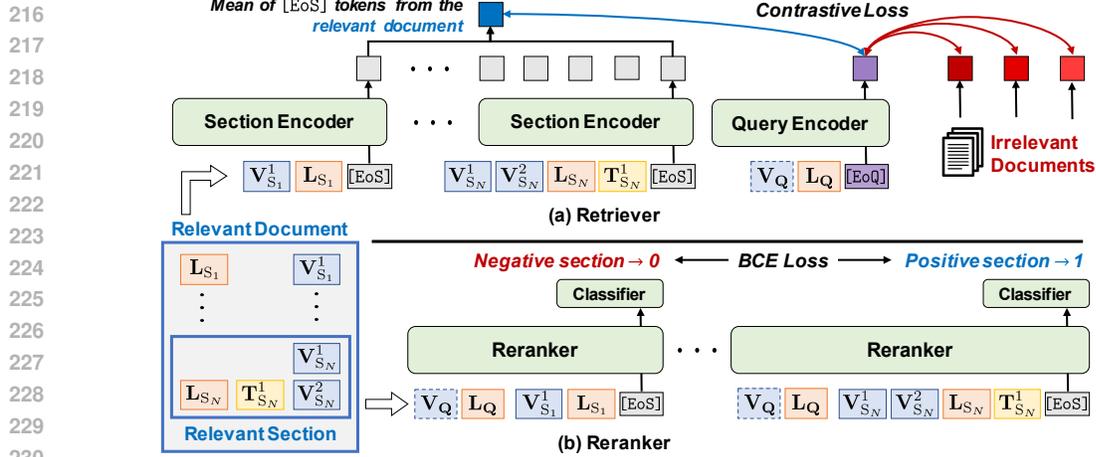


Figure 2: Overview of IDentify. (a): In our document retriever, a query encoder represents a query (purple), and sections are encoded with a section encoder whose embeddings averaged to form a document representation (blue). Contrastive learning loss (red) is used for training the document retriever. (b): Reranker scores query-section relevance with the concatenation of the query and section, trained using Binary Cross-Entropy loss.

The remaining step to discuss here is how to train those two query and document retrievers for IR. Recall that the goal of the retriever is to assess a relevance score between the query and the document. To achieve this goal, we use a contrastive learning loss based upon the query and document representations, whose objective is to assign higher similarity scores to relevant documents (positive samples) and lower scores to irrelevant ones (negative samples) for the query, formulated as follows:

$$\mathcal{L}_{\text{retriever}} = -\frac{1}{B} \sum_{i=1}^B \log \left(\frac{\text{sim}(\mathbf{Z}_{Q_i}, \mathbf{Z}_{D_i})}{\sum_{j=1}^B \text{sim}(\mathbf{Z}_{Q_i}, \mathbf{Z}_{D_j})} \right), \quad \text{sim}(\mathbf{Z}_Q, \mathbf{Z}_D) = \frac{\mathbf{Z}_Q^\top \mathbf{Z}_D}{\|\mathbf{Z}_Q\| \|\mathbf{Z}_D\|},$$

where B is the batch size during the training phase. Here, by minimizing $\mathcal{L}_{\text{retriever}}$, the retriever learns to optimize the similarity between queries and their relevant documents, enabling the retrieval of the most pertinent documents (among all) for the given input query during inference.

3.3 RERANKER

To enable fine-grained retrieval within documents beyond the retrieval of documents themselves, we introduce a section-level reranking mechanism that identifies the section most relevant to the input query. In particular, once the document is retrieved, the objective of the reranker f_R is to pinpoint the specific sections within the document that best match the query. We also note that this reranker is similarly operationalized with a single VLM along with a binary classifier on top of it, which directly measures the relevance of each query-section pair, illustrated in Figure 2 (b).

Formally, for a retrieved document, we take each of its sections s_i and concatenate it with the query q and a learnable token for section embedding $[\text{EoS}]$, forming the input sequence of $[q, s_i, [\text{EoS}]]$. The concatenated tokens are then processed through the reranker, and the model output corresponding to $[\text{EoS}]$ captures the relevant between the query and section, which is further subsequently passed to a binary classifier consisting of a linear layer followed by a Sigmoid function. Through this process, the classifier outputs a probability score indicating the likelihood of the section being relevant to the query, i.e., a score close to one denotes a high relevance (positive section), meanwhile, a score near zero indicates irrelevance (negative section).

To train this reranker, we use the Binary Cross-Entropy (BCE) loss, formalized as follows:

$$\mathcal{L}_{\text{reranker}} = \sum_{i=1}^B \sum_{j=1}^{S_i} \frac{1}{BS_i} \ell(y_{(S_{i,j})}, f_R([q, s_{i,j}, [\text{EoS}]])], \quad \ell(y, \hat{y}) = -[y \log \hat{y} + (1-y) \log(1-\hat{y})],$$

where S_i is the number of sections in the i -th document, $y_{(S_{i,j})}$ is the label for the j -th section of the i -th document $s_{i,j}$ (with its value of one if relevant to the query q and zero otherwise), and B is the batch size during training. Also, in this training process, the sections not labeled as relevant to the query are considered negative samples. Then, by minimizing $\mathcal{L}_{\text{reranker}}$, the reranker learns to predict section relevance for any query, thus refining our overall retrieval process by allowing the retrieval of not just whole documents but also their most relevant sections, for multiple use cases of IR.

4 EXPERIMENTS

4.1 DATASETS

To evaluate the effectiveness of IDentIfy, we focus on multimodal IR tasks that require understanding of both the textual and visual cues within queries and documents, which align well with our goal of enhancing retrieval of multimodal interleaved documents. The datasets considered are as follows:

Encyclopedic-VQA (Mensink et al., 2023) is a large-scale visual question-answering (VQA) benchmark dataset, widely used for measuring the performance of multimodal IR models. Each query is linked to a specific section of a Wikipedia document (containing an answer for it) and is manually annotated by humans. Also, this dataset offers both text-only and multimodal queries. In addition to this, the queries are related to fine-grained properties of species and landmarks. Our experiments focus on the single-hop category where questions that can be answered in a single retrieval step.

InfoSeek (Chen et al., 2023) is a dataset designed for knowledge-intensive VQA, covering a wide range of entities (such as landmarks, animals, and food). Questions are generated by filling human-written templates with knowledge triples (subject, relation, object) available from Wikidata, which involve only the multimodal queries. As the test dataset is not available, we use the validation set as our test set, and split the training set into training and validation subsets with a 9:1 ratio.

ViQuAE (Lerner et al., 2022) is a dataset focused about human entities. It provides both text-based and multimodal queries, with each query linked to a specific section of a Wikipedia document that contains an answer (annotated by humans), which makes it an idea benchmark for section retrieval.

Open-WikiTable (Kweon et al., 2023) is an extension of WikiSQL (Zhong et al., 2017) and WikiTableQuestions (Pasupat & Liang, 2015), designed for open-domain table question answering that requires retrieval of the most relevant table from a broader corpus. For our experiments, we adapt this dataset, aiming at identifying the document or document section containing the target table, and correspondingly, utilize the WikiTableQuestions subset of Open-WikiTable that has labels for it.

4.2 IMPLEMENTATION DETAILS

Model Training and Evaluation We use LLaVA-NeXT-Interleave (Li et al., 2024) of 0.5B parameters as the basis VLM, for both the retriever and reranker. To take the advantage of larger batch sizes (while reducing GPU memory usage), we apply LoRA (Hu et al., 2022). Also, to further optimize the GPU usage, we combine four images into one, scaling each down to half of its original height and width. During retriever and reranker training, we consider four sections per document in representing documents and selecting negative samples. In contrast, during inference, we consider all sections within each document. For section retrieval, the top 25 documents retrieved are split into sections and passed to the rerankers. All experiments are conducted using a single H100 GPU.

Baselines We compare our approach against a variety of IR baselines designed to capture different document representations. We start with Entity and Summary baselines, which are trained to retrieve documents based on their titles and summary sections. Next, we consider the Text-document retriever, which retrieves documents based on their textual content. Additionally, to consider a visual component, we consider the Single-image baseline, incorporating the first document image.

Evaluation Metrics We evaluate the performance of the retriever and reranker with standard metrics: Recall@K (R@K) and Mean Reciprocal Rank@K (MRR@K). First, R@K measures whether the relevant document or section is retrieved within the top-K results. MRR@K evaluates the ranking quality by measuring the position of the first relevant item among the top-K retrieved results.

4.3 RESULTS AND DISCUSSION

Interleaved format improves document retrieval. First, we report the retrieval performance on the Encyclopedic-VQA dataset in Table 1, where each query consists of both image and text content. From this, we observe that our approach achieves the best performance, with R@1 score improvements of 53.0%, 64.0%, and 25.0% compared to the Summary, Text-document, and Single-image retrieval baselines, respectively. The MRR@10 score similarly shows significant gains, with improvements of 36.1%, 48.5%, and 16.2% over the same baselines. This demonstrates the effectiveness of our approach in incorporating the interleaved multimodal format for document representations.

Table 1: Comparison results of the document formats for document retrieval with multimodal queries. Entity uses the document title, Summary uses its first summary section, and Text-document uses only textual content. Based on the Text-document, we add a single image (+ Single-image) or interleaved multimodal content (+ Interleaved).

Format	R@1	R@10	R@100	MRR@10
Entity	3.1	15.5	39.7	6.1
Summary	13.4	41.3	66.5	21.6
Text-document	12.5	37.8	68.7	19.8
+ Single-image	16.4	45.4	77.1	25.3
+ Interleaved (Ours)	20.5	50.0	78.0	29.4

Table 2: Investigation of retrieval granularity in information retrieval. Two granularities are tested on section retrieval with multimodal queries: Passage splits a document by section boundaries, and, using the split passages as retrieval targets, trains the retriever; meanwhile, Document uses documents as retrieval targets. The same reranker is applied to items retrieved by each method. * indicates a result that does not leverage reranking.

Granularity	R@1	R@10	R@20	MRR@10
Passage*	3.9	16.9	22.0	7.5
Passage	28.6	36.4	37.8	31.2
Document (Ours)	35.1	50.8	53.6	40.3

Table 3: Performance in document retrievals. (a): Results of retrieval for multimodal queries on InfoSeek and ViQuAE. (b): Results of retrieval for textual queries on Encyclopedic-VQA (Enc-VQA) and ViQuAE.

(a) Document Retrieval with Multimodal Queries						(b) Document Retrieval with Textual Queries					
Format	Dataset	R@1	R@10	R@100	MRR@10	Format	Dataset	R@1	R@10	R@100	MRR@10
Text-document	InfoSeek	6.8	23.6	52.5	11.2	Text-document	Enc-VQA	62.7	76.3	87.4	67.0
+ Interleaved		10.2	30.4	57.3	15.7	+ Interleaved		65.4	76.8	87.8	69.0
Text-document	ViQuAE	13.5	40.4	67.4	20.9	Text-document	ViQuAE	55.8	71.5	83.0	60.9
+ Interleaved		17.5	46.0	69.4	26.3	+ Interleaved		56.5	72.2	83.0	61.6

To further understand the source of these performance gains, we explore two levels of retrieval granularity: passages and documents. Specifically, the passage retriever uses individual sections of documents as retrieval units, while the document retriever treats entire documents as single units. Both models are trained on the Encyclopedic-VQA dataset for multimodal retrieval. Then, we use the same reranker to both sets of results from passage and document retrievers, to directly compare their performance. In Table 2, we observe that relying solely on the passage retriever (Passage*) results in suboptimal retrieval performance, highlighting the challenge in pinpointing the most relevant section within a document using traditional retrieval methods. In contrast, when the reranker is used alongside the document retriever, the performance significantly surpasses that of the passage retrieval, achieving a 22.7% improvement in R@1 and a 29.2% improvement in MRR@10, even though the document retriever provides eight times fewer retrieval units to the reranker. These results confirm the importance of leveraging holistic context from multiple, interrelated sections within documents. In addition to this, these findings also demonstrate the notable advantages of using the interleaved multimodal elements within documents, emphasizing the potential of this direction.

Interleaved format enhances document retrieval across modalities. We further expand our experiments to two additional IR datasets, the InfoSeek and ViQuAE. As shown in Table 3 (a), our proposed retriever consistently surpasses the Text-document baseline in document retrieval with multimodal queries. Specifically, this leads to 50.0% and 29.6% improvements in the R@1 score, and 40.2% and 25.7% improvements in the MRR@10 score for the InfoSeek and ViQuAE, respectively. We also examine the impact of interleaved documents on textual retrieval tasks, where queries consist solely of text, and report the results in Table 3 (b). Then, the results demonstrate that the interleaved format offers advantages in retrieval of textual queries as well, resulting in 4.3% and 1.3% improvements in the R@1 score and 3.0% and 1.1% improvements in the MRR@10 score for the Encyclopedic-VQA and ViQuAE, respectively. We attribute these gains to the integration of multimodal content within documents, enabling the VLM to capture the multimodal alignment and to exploit its pre-existing knowledge for more effective document representations (Xu et al., 2024).

Interleaved format is also beneficial in section retrieval. Similarly, we evaluate the efficacy of our approach in section retrieval across both multimodal and textual queries, using the Encyclopedic-VQA and ViQuAE datasets. First, in section retrieval with multimodal queries, our model attains 4.2% improvement in the R@1 score and 3.3% improvement in the MRR@10 score for the Encyclopedic VQA, as shown in Table 4 (a). Similarly, in section retrieval with textual queries, our model achieves 2.3% and 7.5% improvements in the R@1 score and 1.8% and 4.9% improvements in the

Table 4: Performance of section retrievals. (a): Results of retrieval for multimodal queries on Encyclopedic-VQA (Enc-VQA) and ViQuAE. (b): Results of retrieval for textual queries on Enc-VQA and ViQuAE. The same document retrieval outcomes are used in section retrieval to solely measure the reranker’s performance.

(a) Section Retrieval with Multimodal Queries						(b) Section Retrieval with Textual Queries					
Format	Dataset	R@1	R@10	R@20	MRR@10	Format	Dataset	R@1	R@10	R@20	MRR@10
Text-document	Enc-VQA	40.7	52.8	55.5	44.8	Text-document	Enc-VQA	68.1	79.4	80.2	72.3
+ Interleaved		42.4	53.6	55.7	46.3	+ Interleaved		69.7	80.1	80.6	73.6
Text-document	ViQuAE	12.6	31.7	37.7	18.2	Text-document	ViQuAE	27.8	50.2	57.7	35.0
+ Interleaved		11.4	32.1	39.2	17.5	+ Interleaved		29.9	50.9	59.8	36.7

Table 5: Document and section retrieval results for tables, where Zero-shot denotes a model finetuned on Encyclopedic-VQA but not trained on the target dataset. Finetuned refers to additional training of the model on the target dataset. (a): Results for tabular document retrieval on the Open-WikiTable (OWT) dataset. (b): Textual section retrieval results on the ViQuAE dataset and tabular section retrieval results on the OWT dataset. (c): Reranker accuracy (Acc@1) of a classification task that identifies the section containing the query-associated table within a gold document containing multiple tables. Random indicates random selection in the task.

(a) Document Retrieval for Tables					(b) Section Retrieval for Tables						
Model	R@1	R@10	R@100	MRR@10	Model	Modality	Dataset	R@1	R@10	R@20	MRR@10
Zero-shot	29.4	58.0	86.0	38.1	Zero-shot	Text	ViQuAE	20.3	49.0	57.7	28.9
Finetuned	55.8	84.1	93.5	66.1	Finetuned			29.9	50.9	59.8	36.7
(c) Tabular Classification					Zero-shot	Table	OWT	5.9	20.5	29.4	9.1
Model	Random	Zero-shot	Finetuned	Finetuned	Finetuned			8.4	36.7	52.8	15.2
Acc@1	11.9	9.3	56.5								

MRR@10 score for the Encyclopedic and ViQuAE datasets, as shown in Table 4 (b). Overall, the design of our Interleaved rerankers exhibit superior or comparable performance to the Text-document rerankers. However, since the rerankers assess query relevance using a single section, they may lack access to broader contextual information from a document, which limits the potential performance gain compared to the retrievers. Nonetheless, the multimodal content interleaved within documents improves the reranker’s ability to evaluate the relevance of the query to individual sections.

Information retrieval of tabular contents in interleaved documents is challenging. We explore a retrieval task for tabular data, whose goal is to select the document or section containing the target table relevant to the input query. Specifically, we use the Open-WikiTable dataset to train the retriever and reranker, and then compare these trained models (Finetuned) with the models trained on the Encyclopedic-VQA dataset (Zero-shot). Then, as shown in Table 5 (a), despite Open-WikiTable consisting of only 3.2k training samples, the Finetuned retriever achieves strong retrieval performance. Meanwhile, the Zero-shot retriever demonstrates only about half of the R@1 score and the MRR@10 score of the Finetuned retriever, though it remains competitive in R@100.

In contrast, the performance trends for the rerankers exhibit notable differences. The discrepancies in R@10 and R@20 scores between the Zero-shot and Finetuned retrievers, denoted in color red in Table 5 (b), are much more pronounced in the Open-WikiTable (table retrieval) experiments than the ones in the ViQuAE (text retrieval) experiments. This highlights a substantial difference between textual and tabular modalities, despite both being represented using word tokens. This suggests that these two modalities may require different handling for retrieval, which we leave as future work.

Notably, the R@1 scores for tabular section retrieval are significantly lower than those for textual section retrieval (Table 5 (b)). To better understand the difficulty of identifying the query-relevant table, we use the Open-WikiTable dataset and design a classification task. In this task, both Zero-shot and Finetuned rerankers are provided with a golden document — a document containing the target of the input query — and should identify the section that contains the target table. Notably, this setup isolates the reranker’s ability to locate the target within a golden document. To ensure that the difficulty of the task is accurately assessed, we focus on documents containing multiple tables. Then, as shown in Table 5 (c), the Zero-shot reranker performs similarly to a random selection, failing to find the correct section. This accounts for its low R@10 and R@20 scores in the tabular retrieval task (Table 5 (b)). In aggregate, while the Finetuned reranker shows improved performance, it still misclassifies nearly half of the tables due to the high similarity between tables within the same

Figure 3: Trade-off between retrieval performance and training cost.

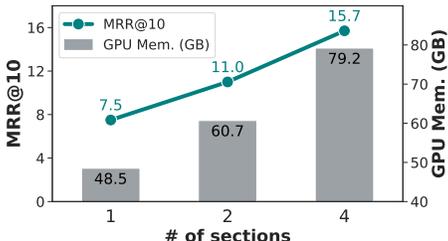


Table 6: Impact of negative sample selection in the reranker training. Top-K selects the top-k retrieved sections from the retriever as negatives. In-batch uses negatives from other sections in the same batch. In-document selects negatives from sections within the document containing the positive section.

Negative	R@1	R@10	R@20	MRR@10
Top-K	38.1	53.7	55.3	44.4
In-batch	39.5	53.8	55.4	45.0
In-document (Ours)	42.4	53.6	55.7	46.3

Table 7: Performance of different reranker designs. Contrastive follows the same training pipeline used for retrievers but here it uses sections for retrieval. Document+BCE concatenates the input query with multiple sections gathered from the same document and uses BCE loss to train the re-ranker. Section+BCE concatenates the query with each section of the document, and the re-ranker is trained with the BCE loss.

(a) Section Retrieval for Multimodal Queries					(b) Section Retrieval for Textual Queries				
Train Loss	R@1	R@10	R@20	MRR@10	Train Loss	R@1	R@10	R@20	MRR@10
Contrastive	3.6	15.0	21.3	6.5	Contrastive	13.6	37.7	45.1	20.6
Document+BCE	13.6	29.6	32.9	24.1	Document+BCE	23.8	43.4	47.2	39.1
Section+BCE (Ours)	42.4	53.6	55.7	46.3	Section+BCE (Ours)	69.7	80.1	80.6	73.6

document. When combined with tables from other documents, this further complicates the task of identifying the exact query-relevant table, as shown in Finetuned reranker’s low score in Table 5 (b).

4.4 FURTHER ANALYSIS AND ABLATION

More sections enhance document retrieval performance but raise computational costs. In Table 1 and Table 2, we observe that using the comprehensive multimodal content and enriched contextual information significantly improves document retrieval performance. Accordingly, we anticipate further improvements as more sections are gathered to represent the document, during training. To validate this, we measure the document retrieval performance with varying the number of sections per document on the InfoSeek dataset for training. The results shown in Figure 3 then indicate that incorporating more sections raises the MRR@10 score from 7.5 to 15.7. However, this performance boost comes with a clear trade-off; as the number of sections increases, the retriever must process additional end-of-section tokens, leading to higher GPU memory consumption. To balance resource limitations and performance gains, we select four sections per document for all experiments.

Sections from the same document act as effective negatives to enhance reranker performance. We explore another method to improve IR effectiveness by leveraging the entire document. Specifically, we investigate the use of sections from the same document as negatives for reranker training, namely In-document. We compare this approach with traditional methods, including Top-K, which selects the top-K retrieved sections as negatives, and In-batch, which uses the positive sections for other samples in the same batch as negatives. After training rerankers with each method, we evaluate section retrieval on the Encyclopedic-VQA dataset. The results shown in Table 6 demonstrate that our In-document approach achieves superior R@1 and MRR@10 scores. This suggests that the use of sections from the same document as negatives provides natural, cost-effective advantages thanks to their high similarity to the positive section. However, it does not consistently outperform the other methods on the R@10 score. We hypothesize that this inconsistency may arise from the strengths of each method: the In-document approach excels at distinguishing sections from the same document, while Top-K and In-batch methods better differentiate sections from different documents.

BCE loss applied to each section produces the best reranker performance. In our reranker design, we apply BCE loss using the query concatenated with each document section (Section+BCE). We also explore alternative training objectives to identify the most effective approach for section retrieval in interleaved documents. One such objective is contrastive loss (Contrastive). This approach is similar to the retriever, but the retrieval unit is a section. Additionally, we also explore a variant of the BCE loss (Document+BCE), where, unlike Section+BCE, the query is concatenated with multi-

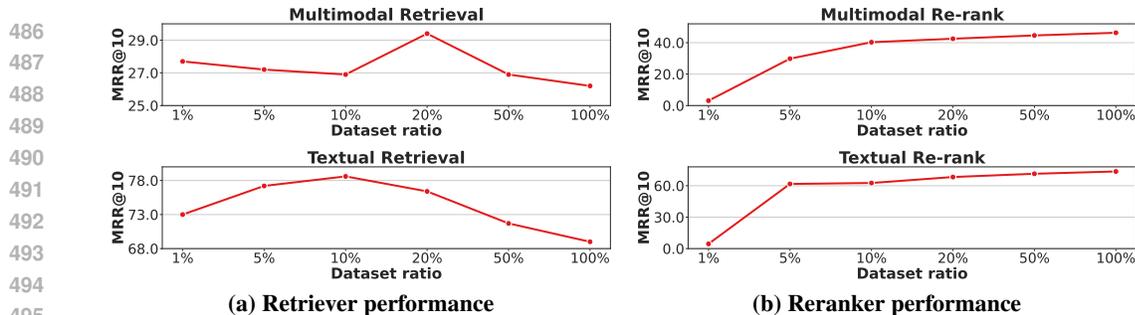


Figure 4: Retrieval performance with different dataset sizes for training. (a): When training a retriever, large datasets rather deteriorate the retrieval performance as it may be overfitted, resulting in low generalization. (b): On the other hand, a larger dataset size is beneficial to training a re-ranker.

ple sections from the same document, including both positive and negative sections. An $[\text{EoS}]$ token is appended to each section, and the Document+BCE follows the same BCE loss calculation of the Section+BCE using its $[\text{EoS}]$ outputs. This design allows the Document+BCE reranker to leverage the long-context understanding of VLMs to improve section retrieval in interleaved documents.

Then, in Table 7, we compare the section retrieval performance of different reranker designs on the Encyclopedic-VQA dataset. We find that the Contrastive reranker performs the worst, indicating that directly concatenating the query with the section at the input level provides more effective clues for query-section relevance assessment. Notably, this observation is consistent with conventional reranker approaches. Further, the Document+BCE reranker underperforms compared to the Section+BCE reranker, likely due to training constraints. Specifically, while the evaluation phase uses all sections within each document, the training phase is limited to a maximum of four sections per document, with an average of eight sections per document. Such a mismatch may degrade the model’s performance. Building on concurrent discoveries (Jiang et al., 2024; Lee et al., 2024), addressing these training constraints will potentially open up new reranker designs that can better handle long, interleaved documents using VLMs, ultimately improving section retrieval performance.

Rerankers require much larger datasets than retrievers. We analyze the effect of different dataset sizes for training on retriever and reranker performance. To achieve this, we randomly prune samples in the Encyclopedic-VQA dataset at various ratios and report the performance of models trained on these subsets. In Figure 4 (a), we observe that too many samples can degrade retrieval performance. Also, retrieval of textual queries requires fewer samples to reach its optimal performance compared to multimodal retrieval. Similarly, in Figure 4 (b), section retrieval for multimodal queries requires 10% of the dataset to achieve 80% of the full-dataset performance, while section retrieval for textual queries needs only 5%. These observations suggest that additional modalities increase the need for more data. This accounts for the inferior performance of the interleaved format in the ViQuAE experiments (Table 4 (a)). The ViQuAE dataset, at only 2.2% of the size of Encyclopedic-VQA, may be small for the reranker to effectively learn multimodal query-section alignments. We also observe that section retrieval is more challenging, with more samples improving the reranker’s performance. This explains why the ViQuAE reranker has much lower section retrieval scores compared to the one trained on the Encyclopedic-VQA (Table 4 (b)). Given the challenge of obtaining large query-section pair samples, exploring more effective reranker training pipelines is necessary.

5 CONCLUSION

In this paper, we introduced IDentify, a novel IR framework designed to address the limitations of conventional methods that rely on solely textual content of documents and their segmented passages. Specifically, our approach sits on top of recent VLMs, which enables integration and representation of diverse multimodal content (including text, images, and tables) into a unified document representation. Also, unlike previous strategies that segment documents at the passage level, our method merges these segments to maintain the document’s structural coherence, while further introducing a reranking strategy for precise identification of relevant sections. Extensive experiments across various IR datasets demonstrated that IDentify consistently outperforms existing baselines, confirming that the interleaved multimodal representation significantly enhances the quality of the document retrieval. We believe IDentify represents a crucial step toward more comprehensive and contextually aware IR systems, capable of handling the increasing multimodality of modern information sources.

540 REPRODUCIBILITY STATEMENT

541
542 Our codes are based on publicly available LLaVA-NeXT (Li et al., 2024). The experimental setup
543 and details can be found in §4 and Appendix A. The experiments are conducted with publicly avail-
544 able datasets (Mensink et al., 2023; Chen et al., 2023; Lerner et al., 2022; Kweon et al., 2023). We
545 have included our codes in the supplementary material and will publicly release our code.

546
547 REFERENCES

548
549 Jinheon Baek, Alham Fikri Aji, Jens Lehmann, and Sung Ju Hwang. Direct fact retrieval from
550 knowledge graphs without entity linking. In *Proceedings of the Association for Computational*
551 *Linguistics (ACL)*, 2023.

552 Davide Caffagni, Federico Cocchi, Nicholas Moratelli, Sara Sarto, Marcella Cornia, Lorenzo
553 Baraldi, and Rita Cucchiara. Wiki-llava: Hierarchical retrieval-augmented generation for multi-
554 modal llms. *arXiv preprint arXiv:2404.15406*, 2024. URL [https://doi.org/10.48550/
555 arXiv.2404.15406](https://doi.org/10.48550/arXiv.2404.15406).

556 Peter Baile Chen, Yi Zhang, and Dan Roth. Is table retrieval a solved problem? exploring join-aware
557 multi-table retrieval. In *Proceedings of the Association for Computational Linguistics (ACL)*,
558 2024.

559 Yang Chen, Hexiang Hu, Yi Luan, Haitian Sun, Soravit Changpinyo, Alan Ritter, and Ming-Wei
560 Chang. Can pre-trained vision and language models answer visual information-seeking ques-
561 tions? In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*
562 *(EMNLP)*, 2023.

563 Wenliang Dai, Junnan Li, Dongxu Li, Anthony Meng Huat Tiong, Junqi Zhao, Weisheng Wang,
564 Boyang Li, Pascale Fung, and Steven C. H. Hoi. Instructblip: Towards general-purpose vision-
565 language models with instruction tuning. In *Advances in Neural Information Processing Systems*
566 *(NeurIPS)*, 2023.

567 Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: pre-training of deep
568 bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of*
569 *the North American Chapter of the Association for Computational Linguistics: Human Language*
570 *Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and*
571 *Short Papers)*, 2019.

572 Yihao Ding, Kaixuan Ren, Jiabin Huang, Siwen Luo, and Soyeon Caren Han. PDF-MVQA:
573 A dataset for multimodal information retrieval in pdf-based visual question answering. *arXiv*
574 *preprint arXiv:2404.12720*, 2024. URL [https://doi.org/10.48550/
575 arXiv.2404.12720](https://doi.org/10.48550/arXiv.2404.12720).

576 Manuel Faysse, Hugues Sibille, Tony Wu, Bilel Omrani, Gautier Viaud, Céline Hudelot, and Pierre
577 Colombo. Colpali: Efficient document retrieval with vision language models. *arXiv preprint*
578 *arXiv:2407.01449*, 2024. URL [https://doi.org/10.48550/
579 arXiv.2407.01449](https://doi.org/10.48550/arXiv.2407.01449).

580 Yunfan Gao, Yun Xiong, Xinyu Gao, Kangxiang Jia, Jinliu Pan, Yuxi Bi, Yi Dai, Jiawei Sun, Qianyu
581 Guo, Meng Wang, and Haofen Wang. Retrieval-augmented generation for large language models:
582 A survey. *arXiv preprint arXiv:2312.10997*, 2023. URL [https://doi.org/10.48550/
583 arXiv.2312.10997](https://doi.org/10.48550/arXiv.2312.10997).

584 Kelvin Guu, Kenton Lee, Zora Tung, Panupong Pasupat, and Ming-Wei Chang. REALM: retrieval-
585 augmented language model pre-training. *arXiv preprint arXiv:2002.08909*, 2020. URL [https://
586 arxiv.org/abs/2002.08909](https://arxiv.org/abs/2002.08909).

587 Jonathan Herzig, Thomas Müller, Syrine Krichene, and Julian Martin Eisenschlos. Open domain
588 question answering over tables via dense retrieval. In *Proceedings of the North American Chapter*
589 *of the Association for Computational Linguistics (NAACL)*, 2021.

590 Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang,
591 and Weizhu Chen. Lora: Low-rank adaptation of large language models. In *Proceedings of the*
592 *International Conference on Learning Representations (ICLR)*, 2022.

- 594 Soyeong Jeong, Jinheon Baek, Sukmin Cho, Sung Ju Hwang, and Jong Park. Adaptive-rag: Learn-
595 ing to adapt retrieval-augmented large language models through question complexity. In *Proceed-*
596 *ings of the North American Chapter of the Association for Computational Linguistics (NAACL)*,
597 2024a.
- 598 Soyeong Jeong, Jinheon Baek, Sukmin Cho, Sung Ju Hwang, and Jong C. Park. Database-
599 augmented query representation for information retrieval. *arXiv preprint arXiv:2406.16013*,
600 2024b. URL <https://doi.org/10.48550/arXiv.2406.16013>.
- 601
- 602 Ziyang Jiang, Xueguang Ma, and Wenhua Chen. Longrag: Enhancing retrieval-augmented generation
603 with long-context llms. *arXiv preprint arXiv:2406.15319*, 2024. URL [https://doi.org/](https://doi.org/10.48550/arXiv.2406.15319)
604 [10.48550/arXiv.2406.15319](https://doi.org/10.48550/arXiv.2406.15319).
- 605 Karen Spärck Jones. A statistical interpretation of term specificity and its application in re-
606 trieval. *J. Documentation*, 60(5):493–502, 2004. URL [https://doi.org/10.1108/](https://doi.org/10.1108/00220410410560573)
607 [00220410410560573](https://doi.org/10.1108/00220410410560573).
- 608 Vladimir Karpukhin, Barlas Oguz, Sewon Min, Patrick S. H. Lewis, Ledell Wu, Sergey Edunov,
609 Danqi Chen, and Wen-tau Yih. Dense passage retrieval for open-domain question answering. In
610 *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*,
611 2020.
- 612
- 613 Sunjun Kweon, Yeonsu Kwon, Seonhee Cho, Yohan Jo, and Edward Choi. Open-wikitable : Dataset
614 for open domain question answering with complex reasoning over table. In *Findings of the Asso-*
615 *ciation for Computational Linguistics (ACL)*, 2023.
- 616 Jinhyuk Lee, Anthony Chen, Zhuyun Dai, Dheeru Dua, Devendra Singh Sachan, Michael Boratko,
617 Yi Luan, Sébastien M. R. Arnold, Vincent Perot, Siddharth Dalmia, Hexiang Hu, Xudong Lin,
618 Panupong Pasupat, Aida Amini, Jeremy R. Cole, Sebastian Riedel, Iftekhhar Naim, Ming-Wei
619 Chang, and Kelvin Guu. Can long-context language models subsume retrieval, rag, sql, and
620 more? *arXiv preprint arXiv:2406.13121*, 2024. URL [https://doi.org/10.48550/](https://doi.org/10.48550/arXiv.2406.13121)
621 [arXiv.2406.13121](https://doi.org/10.48550/arXiv.2406.13121).
- 622 Paul Lerner, Olivier Ferret, Camille Guinaudeau, Hervé Le Borgne, Romaric Besançon, José G.
623 Moreno, and Jesús Lovón-Melgarejo. Viquae, a dataset for knowledge-based visual question
624 answering about named entities. In *SIGIR '22: The 45th International ACM SIGIR Conference*
625 *on Research and Development in Information Retrieval, Madrid, Spain, July 11 - 15, 2022*, 2022.
- 626 Paul Lerner, Olivier Ferret, and Camille Guinaudeau. Cross-modal retrieval for knowledge-based
627 visual question answering. In *Advances in Information Retrieval - 46th European Conference on*
628 *Information Retrieval*, 2024.
- 629
- 630 Patrick S. H. Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman
631 Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, Sebastian Riedel, and
632 Douwe Kiela. Retrieval-augmented generation for knowledge-intensive NLP tasks. In *Advances*
633 *in Neural Information Processing Systems (NeurIPS)*, 2020.
- 634 Feng Li, Renrui Zhang, Hao Zhang, Yuanhan Zhang, Bo Li, Wei Li, Zejun Ma, and Chunyuan Li.
635 Llava-next-interleave: Tackling multi-image, video, and 3d in large multimodal models. *arXiv*
636 *preprint arXiv:2407.07895*, 2024. URL [https://doi.org/10.48550/arXiv.2407.](https://doi.org/10.48550/arXiv.2407.07895)
637 [07895](https://doi.org/10.48550/arXiv.2407.07895).
- 638 Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike
639 Lewis, Luke Zettlemoyer, and Veselin Stoyanov. Roberta: A robustly optimized BERT pretraining
640 approach. *arXiv preprint arXiv:1907.11692*, 2019. URL [http://arxiv.org/abs/1907.](http://arxiv.org/abs/1907.11692)
641 [11692](http://arxiv.org/abs/1907.11692).
- 642 Zheyuan Liu, Cristian Rodriguez Opazo, Damien Teney, and Stephen Gould. Image retrieval on
643 real-life images with pre-trained vision-and-language models. In *Proceedings of the International*
644 *Conference on Computer Vision (ICCV)*, 2021.
- 645
- 646 Xinwei Long, Jiali Zeng, Fandong Meng, Zhiyuan Ma, Kaiyan Zhang, Bowen Zhou, and Jie Zhou.
647 Generative multi-modal knowledge retrieval with large language models. In *Proceedings of the*
AAAI National Conference on Artificial Intelligence (AAAI), 2024.

- 648 Xueguang Ma, Sheng-Chieh Lin, Minghan Li, Wenhu Chen, and Jimmy Lin. Unifying multimodal
649 retrieval via document screenshot embedding. *arXiv preprint arXiv:2406.11251*, 2024a. URL
650 <https://doi.org/10.48550/arXiv.2406.11251>.
651
- 652 Xueguang Ma, Liang Wang, Nan Yang, Furu Wei, and Jimmy Lin. Fine-tuning llama for multi-stage
653 text retrieval. In *Proceedings of the 47th International ACM SIGIR Conference on Research and*
654 *Development in Information Retrieval*, 2024b.
- 655 Thomas Mensink, Jasper R. R. Uijlings, Lluís Castrejón, Arushi Goel, Felipe Cadar, Howard Zhou,
656 Fei Sha, André Araújo, and Vittorio Ferrari. Encyclopedic VQA: visual questions about detailed
657 properties of fine-grained categories. In *Proceedings of the International Conference on Computer*
658 *Vision (ICCV)*, 2023.
- 659 Averi Nowak, Francesco Piccinno, and Yasemin Altun. Multimodal chart retrieval: A comparison
660 of text, table and image based approaches. In *Proceedings of the North American Chapter of the*
661 *Association for Computational Linguistics (NAACL)*, 2024.
662
- 663 OpenAI. Introducing chatgpt. <https://openai.com/blog/chatgpt>, 2022.
664
- 665 OpenAI. GPT-4 technical report. *arXiv preprint arXiv:2303.08774*, 2023a. URL [https://doi.](https://doi.org/10.48550/arXiv.2303.08774)
666 [org/10.48550/arXiv.2303.08774](https://doi.org/10.48550/arXiv.2303.08774).
- 667 OpenAI. GPT-4V(ision) system card. [https://openai.com/index/
668 gpt-4v-system-card/](https://openai.com/index/gpt-4v-system-card/), 2023b.
669
- 670 Panupong Pasupat and Percy Liang. Compositional semantic parsing on semi-structured tables. In
671 *Proceedings of the Association for Computational Linguistics (ACL)*, 2015.
- 672 Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agar-
673 wal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya
674 Sutskever. Learning transferable visual models from natural language supervision. In *Proceed-*
675 *ings of the International Conference on Machine Learning (ICML)*, 2021.
- 676 Ori Ram, Yoav Levine, Itay Dalmedigos, Dor Muhlgay, Amnon Shashua, Kevin Leyton-Brown,
677 and Yoav Shoham. In-context retrieval-augmented language models. *Trans. Assoc. Comput.*
678 *Linguistics*, 11:1316–1331, 2023. doi: 10.1162/TACL_A_00605. URL [https://doi.org/
679 10.1162/tac1_a_00605](https://doi.org/10.1162/tac1_a_00605).
680
- 681 Stephen E. Robertson, Steve Walker, Susan Jones, Micheline Hancock-Beaulieu, and Mike Gat-
682 ford. Okapi at TREC-3. In *Proceedings of The Third Text REtrieval Conference, TREC 1994,*
683 *Gaithersburg, Maryland, USA, November 2-4, 1994*, 1994.
- 684 Sanket Shah, Anand Mishra, Naganand Yadati, and Partha Pratim Talukdar. KVQA: knowledge-
685 aware visual question answering. In *Proceedings of the AAAI National Conference on Artificial*
686 *Intelligence (AAAI)*, 2019.
687
- 688 Weijia Shi, Sewon Min, Michihiro Yasunaga, Minjoon Seo, Richard James, Mike Lewis, Luke
689 Zettlemoyer, and Wen-tau Yih. REPLUG: retrieval-augmented black-box language models. In
690 *Proceedings of the North American Chapter of the Association for Computational Linguistics*
691 *(NAACL)*, 2024.
- 692 Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée
693 Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, Aurélien Rodriguez, Ar-
694 mand Joulin, Edouard Grave, and Guillaume Lample. Llama: Open and efficient foundation
695 language models. *arXiv preprint arXiv:2302.13971*, 2023a. URL [https://doi.org/10.
696 48550/arXiv.2302.13971](https://doi.org/10.48550/arXiv.2302.13971).
- 697 Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Niko-
698 lay Bashlykov, Soumya Batra, Prajwal Bhargava, Shruti Bhosale, Dan Bikel, Lukas Blecher,
699 Cristian Canton-Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy
700 Fu, Wenyin Fu, Brian Fuller, Cynthia Gao, Vedanuj Goswami, Naman Goyal, Anthony Hartshorn,
701 Saghar Hosseini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor Kerkez, Madian Khabsa, Isabel
Kloumann, Artem Korenev, Punit Singh Koura, Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee,

- 702 Diana Liskovich, Yinghai Lu, Yuning Mao, Xavier Martinet, Todor Mihaylov, Pushkar Mishra,
703 Igor Molybog, Yixin Nie, Andrew Poulton, Jeremy Reizenstein, Rashi Rungta, Kalyan Saladi,
704 Alan Schelten, Ruan Silva, Eric Michael Smith, Ranjan Subramanian, Xiaoqing Ellen Tan, Binh
705 Tang, Ross Taylor, Adina Williams, Jian Xiang Kuan, Puxin Xu, Zheng Yan, Iliyan Zarov, Yuchen
706 Zhang, Angela Fan, Melanie Kambadur, Sharan Narang, Aurélien Rodriguez, Robert Stojnic,
707 Sergey Edunov, and Thomas Scialom. Llama 2: Open foundation and fine-tuned chat models.
708 *arXiv preprint arXiv:2307.09288*, 2023b. URL [https://doi.org/10.48550/arXiv.
709 2307.09288](https://doi.org/10.48550/arXiv.2307.09288).
- 710 Zilin Xiao, Ming Gong, Paola Cascante-Bonilla, Xingyao Zhang, Jie Wu, and Vicente Ordonez.
711 Grounding language models for visual entity recognition. *arXiv preprint arXiv:2402.18695*, 2024.
712 URL <https://doi.org/10.48550/arXiv.2402.18695>.
- 713 Lee Xiong, Chenyan Xiong, Ye Li, Kwok-Fung Tang, Jialin Liu, Paul N. Bennett, Junaid Ahmed,
714 and Arnold Overwijk. Approximate nearest neighbor negative contrastive learning for dense text
715 retrieval. In *Proceedings of the International Conference on Learning Representations (ICLR)*,
716 2021.
- 717 Jialiang Xu, Michael Moor, and Jure Leskovec. Reverse image retrieval cues parametric memory
718 in multimodal llms. *arXiv preprint arXiv:2405.18740*, 2024. URL [https://doi.org/10.
719 48550/arXiv.2405.18740](https://doi.org/10.48550/arXiv.2405.18740).
- 720 Xiaohua Zhai, Basil Mustafa, Alexander Kolesnikov, and Lucas Beyer. Sigmoid loss for language
721 image pre-training. In *Proceedings of the International Conference on Computer Vision (ICCV)*,
722 2023.
- 723 Pan Zhang, Xiaoyi Dong, Bin Wang, Yuhang Cao, Chao Xu, Linke Ouyang, Zhiyuan Zhao, Shuang-
724 rui Ding, Songyang Zhang, Haodong Duan, Wenwei Zhang, Hang Yan, Xinyue Zhang, Wei
725 Li, Jingwen Li, Kai Chen, Conghui He, Xingcheng Zhang, Yu Qiao, Dahua Lin, and Jiaqi Wang.
726 Internlm-xcomposer: A vision-language large model for advanced text-image comprehension and
727 composition. *arXiv preprint arXiv:2309.15112*, 2023. URL [https://doi.org/10.48550/
728 arXiv.2309.15112](https://doi.org/10.48550/arXiv.2309.15112).
- 729 Victor Zhong, Caiming Xiong, and Richard Socher. Seq2sql: Generating structured queries from
730 natural language using reinforcement learning. *arXiv preprint arXiv:1709.00103*, 2017. URL
731 <http://arxiv.org/abs/1709.00103>.
- 732 Junjie Zhou, Zheng Liu, Shitao Xiao, Bo Zhao, and Yongping Xiong. VISTA: visualized text em-
733 bedding for universal multi-modal retrieval. In *Proceedings of the Association for Computational
734 Linguistics (ACL)*, 2024.
- 735 Yutao Zhu, Huaying Yuan, Shuting Wang, Jiongnan Liu, Wenhan Liu, Chenlong Deng, Zhicheng
736 Dou, and Ji-Rong Wen. Large language models for information retrieval: A survey. *arXiv preprint
737 arXiv:2308.07107*, 2023. URL <https://doi.org/10.48550/arXiv.2308.07107>.
- 741
742
743
744
745
746
747
748
749
750
751
752
753
754
755

Appendix

Organization The supplementary file is organized as follows: In [Appendix A](#), we explain the implementation details for our experiments. In [Appendix B](#), we outline the limitations of our study.

A IMPLEMENTATION DETAILS

Table 8: Information retrieval datasets summary.

Dataset	Query Modality	Target	Domain	Entities	Section ID	Train	Eval	Test	Corpus size
Encyclopedic-VQA	Text, Text-Image	Text	Species, Landmarks	17k	o	177k	2.2k	3.8k	100k
InfoSeek	Text-Image	Text	Diverse	11k	x	209k	23k	74k	500k
ViQuAE	Text, Text-Image	Text	Human	1k	o	1.2k	1.2k	1.2k	100k
Open-WikiTable	Text	Table	Table	-	o	3.3k	0.4k	0.4k	1.8k

Dataset configuration [Table 8](#) summarizes the key properties of the datasets used in our experiment, including query modality, target item, entity domain, number of entities, and whether a section ID is provided to indicate the section containing the answer. Additionally, we provide the number of samples in the training, evaluation, and test splits, as well as the size of the corpus.

Dataset pre-processing In our study, we leverage interleaved multimodal content from Wikipedia documents. However, the existing corpora associated with our IR datasets often lack this content, typically only including the first few words of each document. Therefore, we augment the corpora by downloading the HTML file of each Wikipedia document.

If the dataset provides Wikipedia URLs for its corpus, we use them to download the HTML files. Alternatively, if only entity names are provided, we generate Wikipedia URLs using those names. If a Wikipedia URL is deprecated, we remove the corresponding document from the corpus along with any associated queries. From the HTML files, we extract text, image URLs, and tables. We then split the contents by subtitles in the document where each chunk corresponds to a section. For the images, we use the image URLs to download the corresponding images, removing any invalid URLs. This process produces a dictionary that organizes text, images, and tables by section.

Since downloading the complete contents for all documents across datasets is time- and memory-intensive, we preprocess the subsets of each corpus, including documents relevant to queries in the training, evaluation, and test splits, as well as unrelated entity documents.

B LIMITATIONS

Due to the limitations of a single H100 GPU, we represent documents by selecting a limited number of sections and averaging their corresponding embeddings. While this reduces the computational demands, our findings suggest that capturing a broader document context leads to improved retrieval performance. Hence, leveraging the long context window of LVLMs could further enhance document retrieval by capturing more comprehensive information from the full document. Moreover, our reranker design follows the conventional approach of concatenating the input query with individual sections. However, we believe that providing the reranker with all the sections together would allow the model to better leverage the contextual information from the entire interleaved document, potentially resulting in improved performance. In order to fully leverage the interleaved format in the IR system, addressing the issues by reducing the GPU load when processing interleaved documents would greatly boost overall IR performance. We leave these explorations for future work.