VISA: Retrieval Augmented Generation with Visual Source Attribution

Anonymous ACL submission

Abstract

Generation with source attribution is important for enhancing the verifiability of retrieval-augmented generation (RAG) systems. However, existing approaches in RAG primarily link generated content to document-level references, making it challenging for users to locate evidence among multiple content-rich retrieved documents. To address this challenge, we propose *Retrieval-Augmented Generation with Visual Source Attribution* (VISA), a novel approach that combines answer generation with visual source attribution. Leveraging large vision-language models (VLMs), VISA identifies the evidence and highlights the exact regions that support the generated answers with bounding boxes in the retrieved document screenshots. To evaluate its effectiveness, we curated two datasets: Wiki-VISA, based on crawled Wikipedia webpage screenshots, and Paper-VISA, derived from PubLayNet and tailored to the medical domain. Experimental results demonstrate the effectiveness of VISA for visual source attribution on documents' original look, as well as highlighting the challenges for improvement. Code, data, and model checkpoints will be released.

1 Introduction

Retrieval-augmented generation (RAG) has become a key technique for enhancing the reliability in information-seeking processes (Gao et al., 2024). Traditional RAG pipeline directly generates an answer to a user query from retrieved candidate documents (Chen et al., 2017; Lewis et al., 2020). Yet, it is hard for users to verify the sources and appropriately trust generated answers, given that models could produce hallucinated content (Min et al., 2023; Malaviya et al., 2024). Recent works have introduced the generation with citation paradigm (Gao et al., 2023; Ye et al., 2024), prompting the model to not only generate answers but also directly cite the identifiers of the source documents. Such source attribution approaches make it possible for users to check the reliability of the outputs (Asai et al., 2024).

However, text-based generation with source attribution faces several issues: First, citing the source at the document level could impose a heavy cognitive burden on users (Foster, 1979; Sweller, 2011), where users often struggle to locate the core evidence at the section or passage level within the dense and multi-page document. Despite such granularity mismatch could be addressed through passage-citation-based generation methods — linking answers to specific text chunks, it requires non-trivial extra engineering efforts to match the chunk in the document source. Moreover, visually highlighting text chunks in the source document is more intuitive for users, but it remains challenging as it requires control over document rendering, which is not always accessible, such as in PDF scenarios.

Inspired by the recent document screenshot embedding retrieval paradigm — dropping the document processing module and directly using VLM to preserve the content integrity and encoding document screenshots for retrieval (Ma et al., 2024), we ask whether source attribution can also be integrated into such a unified visual paradigm to establish a fully visual, end-to-end verifiable RAG pipeline that is both user-friendly and effective?

To this end, we propose *Retrieval Augmented Generation with Visual Source Attribution* (VISA). In our approach, a large vision-language model (VLM) processes single or multiple retrieved document images and not only generates an answer to the user query but also returns the bounding box of the relevant region within the evidence document. As illustrated in Figure 1, this method enables direct attribution by visually pinpointing the exact position within the document, allowing users to quickly check the supporting evidence within the original context for the generated answer. VLMs are not restricted by document format or rendering,

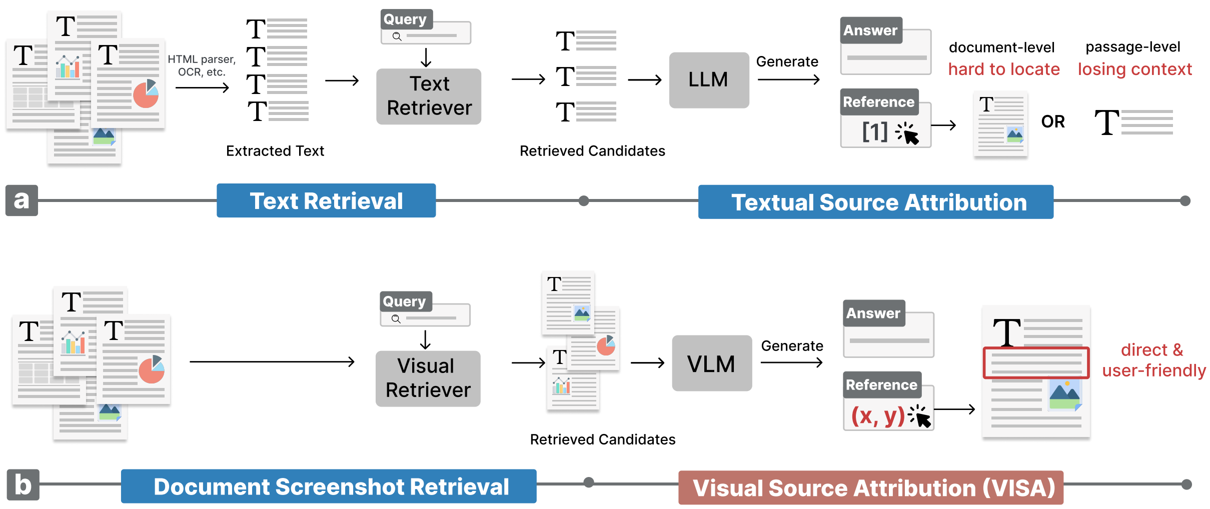

Figure 1: Comparison between (a) Text-based generation with source attribution in a RAG pipeline. and (b) Visual-based generation with source attribution in a V-RAG pipeline. VISA directly pinpoint the source evidence of the answer for user query in the original document with a bounding box.

making them more versatile for diverse use cases. Moreover, this task serves as a meaningful evaluation of VLMs, assessing their ability to provide self-explanations and accurately localize supporting information within their original input in an RAG paradigm.

To train and evaluate VISA, we curated two datasets: Wiki-VISA and Paper-VISA. Wiki-VISA is derived from the Natural Questions dataset (Kwiatkowski et al., 2019). It reconstructs the original Wikipedia webpages, using short answers as generation targets and corresponding long answer's HTML bounding box as source attribution targets. This dataset supports the test of model's ability to attribute sources across multi-document, multi-page, and multi-modal content. On the other hand, Paper-VISA, built from PubLayNet (Zhong et al., 2019) with synthetic query generation, focuses on the biomedical domain by evaluating performance on multi-modal scientific paper PDFs. Together, they provide diverse and challenging benchmarks for assessing the granularity and accuracy of source attribution in RAG systems. Our experiments, spanning both in-domain training and zero-shot evaluation, revealed existing state-of-the-art models like QWen2-VL-72B (Wang et al., 2024) struggle with precise visual source attribution in zero-shot prompting. Fine-tuning VISA on our curated datasets significantly improved model performance in visual attribution accuracy. Further analysis highlights key areas for improvement, such as enhancing bounding box precision for long image documents, multi-documents, and zero-shot generalization capabilities.

2 Related Work

2.1 RAG attribution

Open-domain question answering with LLMs often suffer from two key issues: hallucinations and outdated internal knowledge. Retrieval-Augmented Generation (RAG) has been recognized as an effective solution to these problems (Lewis et al., 2020; Gao et al., 2024; Ovadia et al., 2024). In RAG, relevant documents are first retrieved from an external database and then fed into LLMs alongside the question. This allows LLMs to reference the retrieved documents during answer generation. Furthermore, RAG can generate a list of citations attached to the generated answers, linking them to the retrieved documents so users can verify the accuracy of the output. This process is known as source attribution (Rashkin et al., 2023; Bohnet et al., 2023; Khalifa et al., 2024).

Typically, RAG with source attribution follows a text-only pipeline where all inputs and outputs, such as questions, retrieved documents, generated answers, and citations, are in textual form. Recently, vision-based RAG pipelines have emerged, where the retrieved documents are represented as screenshot images (Ma et al., 2024; Faysse et al., 2024), and VLMs process both textual questions and these document images to generate answers (Riedler and Langer, 2024; Xia et al., 2024; Yu et al., 2024). Compared to traditional text-only RAG, vision-based RAG can leverage structured and visual information from documents, such as tables, graphs, and images, which are often challenging to extract through text-only pipelines.

Our VISA attribution method proposed in this paper is a novel approach for vision-based RAG pipelines: directly drawing bounding boxes around the content in retrieved document screenshots that potentially supports the generated answers. This approach differs from existing attribution methods in two ways: (1) Granularity: Existing attribution methods often operate at the document level, requiring users to read entire documents to locate supportive content. In contrast, our method directly attributes the answer to specific content within the document, such as a passage, table, or image in the screenshot. (2) Presentation: Traditional attribution methods provide a list of textual citations, whereas our method uses bounding boxes, offering a visually-oriented form of attribution. This can help users quickly locate the relevant information.

2.2 Bounding Box Drawing with VLM

Bounding box-based object detection is a well-established task in computer vision (CV) (Zhao et al., 2019; Zou et al., 2023). Traditional approaches rely on convolutional neural networks (CNNs) (LeCun et al., 2015) or Vision Transformers (ViTs) (Dosovitskiy et al., 2021) to extract features and predict bounding boxes alongside object classification (Ren et al., 2015; Dai et al., 2016; Redmon et al., 2016; Carion et al., 2020).

Recent vision-language models (VLMs) like GPT4o (OpenAI et al., 2024), QWen2-VL (Wang et al., 2024), and PaliGemma (Steiner et al., 2024) have shown the ability to generate bounding box coordinates in an image-to-text manner, taking input images and generate the top-left and bottom-right coordinates of target objects. Unlike traditional object detection that focuses on natural images, our method applies bounding box drawing to text-intensive document screenshots.

Additionally, grounding elements on screenshots has been explored in GUI agent systems (Cheng et al., 2024; Lin et al., 2024), where bounding boxes are used to localize UI elements like buttons. While these approaches focus on GUI contexts, our work targets visual source attribution in vision-based RAG processes, grounding bounding boxes to locate evidence within document images.

3 Method

3.1 Task Definition

Our VISA is a novel source attribution method primarily designed for vision-based RAG systems. To formally define the task of RAG with VISA: given a textual user query q as the RAG system input, the retrieval component of the system needs to retrieve a set of candidate documents $D = \{d_1, ..., d_n\}$ from corpus $\mathcal{C}$. Then the generation component of the system needs to return three outputs: an answer a that answers the query q, the identifier i of the most relevant document d^* in D, and a bounding box coordinates $B_{d^*} = [(x_1, y_1), (x_2, y_2)]$ within d_* that highlight the content supporting the generated answer a.

In a vision-based RAG setup, user queries are textual, while all documents in the corpus $\mathcal{C}$ are screenshots of documents (e.g., webpages or PDF pages) provided as image inputs.

3.2 Generation with Visual Source Attribution

This paper focuses on VISA within the generation component of vision-based RAG systems. As discussed in the previous section, VISA must handle multimodal input. To achieve this, we leverage VLMs for implementing VISA. Specifically, for a given query and a set of retrieved candidate documents (i.e., screenshots of documents), the system processes the inputs as follows: query tokens are directly input into the language model, while document screenshots are first processed by the image encoder to extract image representations, which are then fed into the language model.

The language model subsequently generates the answer, the identifier of the relevant document, and the xy-coordinates of the bounding box's top-left and bottom-right corner on the content that supports the generated answer. Notably, this entire process can be framed as a next-token prediction task. Finally, the generated identifier and bounding box coordinates are used to draw the bounding box on the target document screenshot, which is presented to the user along with the generated answer.

Technically, existing instruction-tuned VLMs, such as Qwen2-VL-72B (Wang et al., 2024), can potentially be prompted to perform VISA in a zero-shot manner. However, we find that VISA remains a challenging task. Consequently, further supervised fine-tuning on a dedicated VISA task dataset is necessary. In the next section, we introduce the datasets we crafted specifically for training and evaluating VISA.

3.3 Dataset Acquisition

The training and evaluation data suitable for the VISA task needs to be formatted as follows: the input consists of a textual query and document screenshot images as multimodal inputs, while the target outputs include the textual short answer, the relevant document identifier, and the coordinates of the bounding box. To create datasets that meet these requirements, we craft existing publicly available datasets to support the training and evaluation of our proposed VISA method.

Wiki-VISA is derived from the Natural Questions (NQ) dataset (Kwiatkowski et al., 2019). The original NQ dataset provides natural questions, along with short and long answers sourced from Wikipedia webpages. We use the short answers as answer targets. However, the original dataset does not contain the original webpage screenshots. We use the Selenium Python toolkit[1] to access and render the webpage with the original URL with a history version stamp. And take a screenshot with 980 pixels width and up to 3920 pixels (4 pages) height. Using the long answer, we identify the corresponding element in the HTML from which the long answer is derived. We then draw a bounding box around this element to obtain thecoordinates. Notably, the answers in this dataset can come from various elements, such as passages, tables, lists, or images within the webpage. Since the questions and answers in Wiki-VISA are human-judged, we consider this dataset a high-quality, supervised dataset and evaluation for VISA on general knowledge, with Wikipedia webpage.

Paper-VISA is derived from PubLayNet (Zhong et al., 2019), a dataset originally designed for document layout analysis of single page PubMed PDF documents. PubLayNet provides bounding box coordinates and class labels (e.g., title, text, table, figure, etc.) for each element in a paper's PDF screenshot. However, the dataset does not include queries or answers associated with each document. To address this limitation, we leverage instruction-tuned VLMs (e.g. Qwen2-VL-72B) to synthetically generate queries and answers. Specifically, for each paper screenshot sample in the PubLayNet training data, we select a bounding box within the sample and overlay it on the screenshot. The modified screenshot is then input to the VLM with a prompt designed to instruct the model to generate a question and a short answer based on the content within the bounding box. See Appendix A.1 for the prompt details and generation example. By augmenting the original PubLayNet in this way, we create synthetic queries and answers, enabling it to support VISA training. We consider the resulting Paper-VISA dataset as synthetic training and evaluation for scientific paper PDFs in the medical domain.

FineWeb-VISA is based on the FineWeb-edu corpus (Penedo et al., 2024), a high-quality text corpus of crawled webpages. We sampled 60k webpage URLs and used Selenium to capture screenshots of diverse, content-rich webpages. A passage containing more than 50 words was randomly selected as the target source. A bounding box was drawn around the selected content, and a VLM was prompted to generate a query and short answer supported by the target content, similar as Paper-VISA. Although Fineweb-VISA provides diverse layout, it do not guaranteed to high quality data has human annotated in Wiki-VISA or Paper-VISA that assessing a specific domain, we only leverage Fineweb-VISA as training data to analysis zeroshot and data augmentation effectiveness.

Multi-Candidates By now, each query is paired with the triplet of a positive document, target short answer, and target evidence bounding box. To set up a RAG experimental environment for evaluating VISA, we in addition need to let the generator take multiple candidates as input, simulating the scenario that the generator is taking multiple retrieval candidates and attributing the evidence in most relevant documents. Given the query q, we use a retriever R to retrieve top-k candidates. And randomly sampled $m-1$ candidates that are not ground truth as hard negative candidates. The hard negative candidates are mixed with the one ground truth document together as the input for the multi-document VISA. The reason we did not directly take top-m documents as the retrieval candidate is that we do not want VISA biased on a specific retriever and position of the candidate docs. Generally, our VISA does not rely on the type of retriever. It can be either a traditional text-based retriever that indexes the document with extracted text or a recent document screenshot retriever that directly indexes the original document screenshot. However, integrating with those visual-based retrievers enables

[1]https://pypi.org/project/selenium/

Dataset	# Train	# Test
Wiki-VISA	87k	3,000
Paper-VISA	100k	2,160
Fineweb-VISA	60k	-

Table 1: Datasets statistics for train and test splits.

us to build an end-to-end RAG solution without the necessity of explicit document content processes such as HTML parsing or OCR. Thus, we leverage an off-the-shelf Document Screenshot Embedding (DSE) model (Ma et al., 2024) to serve as the retrieval component of the RAG system. When encoding queries and documents, the model directly encodes textual queries and document screenshot images into single vector embeddings and performs cosine similarity search during inference. In this work, we set $k = 20$ and $m = 3$.

Additionally, an RAG pipeline may have the chance of having no ground truth document returned from the retriever. We use a probability of 20% to randomly replace the ground truth document in the candidates, to access the model's capability to detect no-answer situations. After these operations, the data statistics are shown in Table 1.

4 Experiment Setup

4.1 Evaluation

Evaluation metrics assessed both the generated answers and bounding box predictions. For answer generation, relaxed exact match (EM) was used to measure accuracy. If the golden answer and predicted answer have a sub-sequence relationship and the difference in string length is within 20 characters. The predicted answer is considered as correct. For bounding boxes, Intersection over Union (IoU) was calculated to determine localization precision, with an IoU threshold of 0.5 indicating a correct prediction.

To analyze performance across varying content types, test samples were categorized by the modality and location of the evidence. For Wiki-VISA, categories included first-page passages, passages beyond the first page, and non-passage content such as tables and figures. For Paper-VISA, since it is a single-page document, categories were divided into passage and non-passage content. The overall accuracy for each dataset was computed as a macro average across these categories.

We evaluate the effectiveness of VISA in two different settings: *Single oracle candidate* and *Multi-candidate*. *Single oracle candidate* setting solely evaluates the generation and visual attribution component. We conduct controlled experiments by training and testing the VLMs using only a single ground truth relevant document screenshot as input. In this setup, it is guaranteed that the answer can be found within the input document. The VLMs do not need to predict the relevant document identifier and can focus exclusively on answer generation and bounding box prediction.

In a *Multi-candidate* setting, the model is evaluated on its ability to distinguish relevant documents from irrelevant ones, in addition to generating accurate answers and bounding boxes. This setup better reflects the RAG scenarios in which multiple candidate documents are retrieved, and the model must not only generate a correct response but also attribute it to the correct source document. For the *Multi-candidate* evaluation, we assess two configurations: *Multi-candidate, Oracle in Candidates* which has ground truth in candidates, this setting has the same query set as the single setting, hence directly comparable. *Multi-candidate, Full* contains the additional 20% cases where ground truth has no answer.

4.2 Training details

To train vision-language models (VLMs) for answer generation with VISA, we initialized the models using the open-source Qwen2-VL-2B and Qwen2-VL-7B (Wang et al., 2024), finetuning on the training datasets described in Section 3.3.

We first trained the models in a single-candidate setup, where the input was limited to a single oracle document image. In this setup, the model was trained to generate both the answer and its corresponding bounding box. We used the prompt template provided in Appendix A.2 to format the model's input and output.

Next, we trained the models in a multi-candidate setup. Here, the model received three document candidates and the task was to generate the identifier of the relevant document (if present), the answer, and the bounding box for the evidence. For cases where no relevant document was present (20% of the training samples), the model was trained to generate "No answer." We used the prompt template provided in Appendix A.3 to format the model's input and output.

The training objective for both setups was next-token prediction with cross-entropy loss. We fine-

Method	Wiki-VISA								Paper-VISA					
	Average		[<1] Passage		[>1] Passage		Non-Passage		Average		Passage		Non-Passage	
	bbx	ans	bbx	ans	bbx	ans	bbx	ans	bbx	ans	bbx	ans	bbx	ans
Zeroshot Prompt														
QWen2-VL-72B	1.5	60.4	3.4	58.5	0.1	54.9	0.9	67.9	1.5	43.1	0.5	40.2	2.5	45.9
Fine-tune, Single Oracle Candidates														
VISA-2B-single	37.5	57.1	70.0	61.1	18.7	44.9	23.8	65.3	63.0	38.3	50.6	34.4	75.3	42.1
VISA-7B-single	54.2	65.2	75.6	66.5	50.1	56.0	36.8	73.1	68.2	43.8	58.1	41.6	78.2	45.9
Fine-tune, Multi Candidates, Oracle in Candidates														
VISA-2B-multi	22.5	37.9	46.5	46.1	6.4	27.2	14.6	40.5	51.3	33.8	41.1	30.1	61.4	37.4
VISA-7B-multi	37.7	41.8	58.1	49.2	32.8	32.0	22.2	44.1	59.9	39.2	47.7	35.9	72.0	42.4
Fine-tune, Multi Candidates, Full														
VISA-2B-full	32.1	46.9	51.0	53.6	18.9	38.0	26.5	49.1	59.8	44.7	51.6	42.6	67.9	46.7
VISA-7B-full	41.6	51.1	56.6	57.1	34.4	43.2	33.9	53.1	66.8	50.3	57.1	47.5	76.5	53.0

Table 2: Effectiveness of VISA on Wiki-VISA and Paper-VISA datasets for bounding box accuracy (bbx) and answer accuracy (ans). Fine-tuned models are trained individually on in-domain data. The *Multi-Candidate, Oracle in Candidates* setting uses the same query set as the Single Oracle Candidates setting, allowing direct comparison. The full setting has an additional 20% queries with no ground truth documents in candidates.

tuned the models for two epochs in the single-candidate setting, using LoRA with a learning rate of 1e-4, a batch size of 64, and 4×H100 GPUs. For the multi-candidate setting, we initialized the models with weights from the single-candidate setup and trained for one epoch with the same learning rate. We froze the image encoder to reduce GPU memory usage in the multi-candidate setting.

During the training, random cropping was applied outside of the bounding box. This augmentation exposed the model to varying input sizes, which enhanced its zero-shot effectiveness on unseen document layouts. Bounding box targets were represented using absolute coordinate values. We also explored normalizing the scale of bounding box coordinates to values in the range[0-1]. Details can be found in Section 6.3.

5 Experimental Results

Table 2 presents the performance of VISA on the Wiki-VISA and Paper-VISA datasets across different experimental settings. Zero-shot prompting results reveal the difficulty of directly applying state-of-the-art VLMs to the visual source attribution task. QWen2-VL-72B achieves a reasonable answer generation accuracy of 60.4% on average on Wiki-VISA but fails to deliver effective bounding box predictions, with only 1.5% accuracy. This observation is consistent on Paper-VISA. These highlight the limitations of existing VLMs in pinpointing the source evidence in original documents with proper location and granularity.

Fine-tuning on our crafted training data enables the model to effectively execute the task. In the single-candidate setup, where the model processes only the relevant document, fine-tuned models demonstrate substantial gains compared to zero-shot prompting a much larger model. On Wiki-VISA, the 7B variant achieves 54.2% bounding box accuracy and 65.2% answer accuracy, while on Paper-VISA, the corresponding scores reach 68.2% and 43.8%. Performance in the multi-candidate setting, which more closely mirrors real-world retrieval-augmented generation (RAG) systems, shows similar trends. The 7B model achieves 41.6% bounding box accuracy and 51.1% answer accuracy when handling three candidate documents, including cases where no relevant document is present. This demonstrates the model's capability to identify relevant sources among multiple documents while enabling fine-grained attribution. However, when comparing the multi-candidates, oracle in candidates setting, We can see the model facing challenges when handling multiple candidates compared to just handling a single relevant document. E.g. on Wiki-VISA, bounding box accuracy for 7B model is 37.7% on average which is 17 points lower than the corresponding single candidate setting. Showing that visual source attribution among multi-candidates is much harder than just locating the source element in a single one.

It further demonstrates that the effectiveness of VISA is influenced by document characteristics, such as content location and modality. For Wiki-VISA, bounding box accuracy is significantly higher for passages on the first page ([<1] passage) compared to passages beyond the first page ([>1]

Train Data	Wiki-VISA								Paper-VISA					
	Average		[<1] Passage		[>1] Passage		Non-Passage		Average		Passage		Non-Passage	
	bbx	ans	bbx	ans	bbx	ans	bbx	ans	bbx	ans	bbx	ans	bbx	ans
Wiki	54.2	65.2	75.6	66.5	50.1	56.0	36.8	73.1	27.8	36.2	20.5	32.6	35.1	39.7
Paper	0.2	42.6	0	46.3	0.4	33.5	0.1	48.1	68.2	43.8	58.1	41.6	78.2	45.9
FineWeb	37.6	50.2	48.9	45.1	57.3	52.3	6.6	53.1	22.0	43.3	26.5	41.7	17.4	44.9
Wiki+Fineweb	58.2	65.3	68.7	66.6	61.7	57.1	44.1	72.1	21.0	43.1	18.5	42.2	23.4	43.9
Paper+Fineweb	36.1	48.7	51.8	49.6	49.6	44.2	6.8	52.4	66.5	44.6	56.1	42.2	76.9	47.0
Wiki+Paper+Fineweb	58.1	64.8	69.9	65.0	58.7	56.7	45.8	72.7	67.6	44.3	55.9	41.5	79.3	47.1

Table 3: Effectiveness of VISA trained on different combinations training data for bounding box accuracy (bbx) and answer accuracy (ans) in the single oracle candidate setting.

passage). For example, the 2B variant achieves 70.0% accuracy for [<1] passages but only 18.7% for [>1] passages, indicating the challenges posed by long, multi-page documents. The larger model, the 7B variant, narrows this gap, reflecting the better handling of long-context inputs. Non-passage content, such as tables and figures, also have obviously a different level of grounding effectiveness, indicating the difference of effectiveness in different visual elements.

6 Analysis

6.1 Out-of-Domain Zeroshot

Table 3 shows the effectiveness of VISA while trained with different data combinations in the single candidate setting. It enables us to study the effectiveness of out-of-domain transfer and augmentation. First, we highlight the challenges of zero-shot generalization in VISA. Training and evaluating on in-domain achieves an effective bounding box accuracy, e.g. 54.2% on average for Wiki-VISA. However, significant performance drops are observed when models are tested on out-of-domain datasets. For instance, a model trained on Wiki-VISA achieves only 27.8% bounding box accuracy on Paper-VISA, while a model trained on Paper-VISA achieves near-zero performance (0.2%) on Wiki-VISA. This gap underscores the difficulty of transferring visual source attribution capabilities across datasets with differing document structures, layouts, and content modalities. Interestingly, Wiki-VISA appears to transfer better to Paper-VISA compared to the reverse. This may be because of the multi-page nature of Wiki-VISA, which provides richer training signals that generalize better to simpler single-page setting in Paper-VISA.

FineWeb-VISA shows as a promising resource for training models with improved zero-shot capabilities. When trained on FineWeb-VISA alone, the model achieves 37.6% bounding box accuracy on Wiki-VISA and 22.0% on Paper-VISA. Notably, FineWeb-VISA outperforms Wiki-VISA training on [>1] passage bbx accuracy for Wiki-VISA (57.3% vs. 50.1%), suggesting its effectiveness in handling long and complex document structures. However, FineWeb-VISA does not perform as well on non-passage content, likely due to its training focus on passage-level targets.

6.2 Data Augmentation

The results also demonstrate the benefits of augmenting training data with FineWeb-VISA. On Wiki-VISA, combining Wiki and FineWeb training data improves bounding box accuracy from 54.2% to 58.2% and improves performance on [>1] passages from 50.1% to 61.7%, indicating that FineWeb complements Wiki by enhancing the model's ability to attribute evidence in multi-page contexts. For Paper-VISA, however, augmenting with FineWeb does not significantly improve in-domain performance. Training on Paper+FineWeb achieves a comparable bounding box accuracy to Paper alone, but it enhances zero-shot performance on Wiki-VISA (from 0.2% to 36.1%).

Training on the full combination of datasets (Wiki+Paper+FineWeb) yields strong results across both domains, with 58.1% bbx accuracy on Wiki-VISA and 67.6% on Paper-VISA. This shows the importance of diverse training data for building generalizable models capable of handling different document types, layouts, and evidence modalities. Future work should focus on expanding the dataset diversity to further improve generalization and enable robust visual source attribution for a wide range of document structures.

6.3 Bounding Box Target

Table 4 shows the impact of different bounding box target representations and cropping strategies during training. Training with random cropping

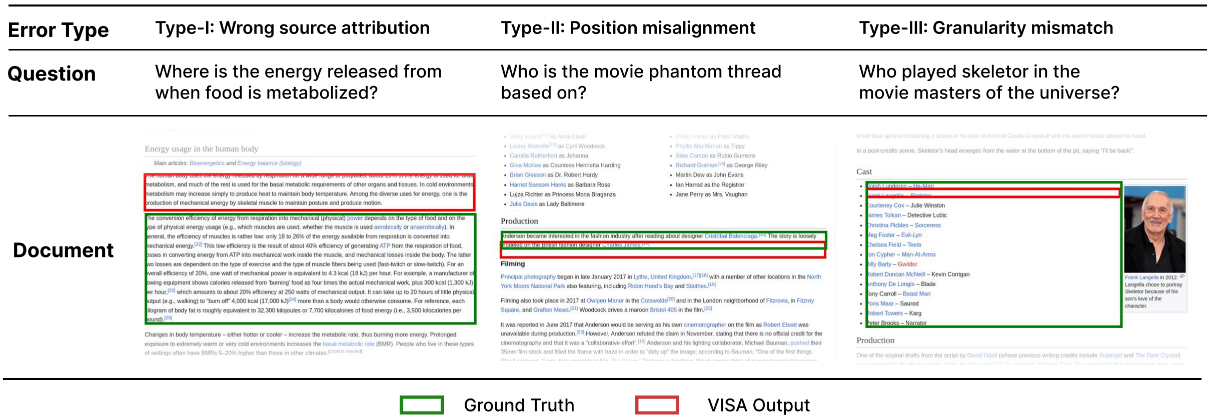

Figure 2: Type of errors in the evaluation of Wiki-VISA.

Train Data	Wiki-VISA bbx	Wiki-VISA ans	Paper-VISA bbx	Paper-VISA ans
Crop, Absolute	54.2	65.2	27.8	36.2
No Random Crop	58.8	65.6	1.7	36.9
Normalized Value	56.4	64.4	0.1	37.2
No Bounding Box	0	67.6	0	35.2

Table 4: Impact of bounding box target representation and cropping strategies during training on Wiki-VISA in the single oracle candidate setting.

and absolute coordinate values achieves a balance between in-domain performance on Wiki-VISA (54.2%) and zero-shot generalization to Paper-VISA (27.8%) in bounding box accuracy. Removing random cropping slightly improves Wiki performance but drastically reduces zero-shot generalization, indicating that random cropping enhances the model's robustness to varied input sizes. Normalizing coordinate values achieves moderate performance on Wiki-VISA but fails on Paper-VISA, suggesting that absolute bounding box values are better suited to our experiments.

The "No Bounding Box" row represents a vanilla visual retrieval-augmented generation setup without visual source attribution, where models generate answers without bounding box predictions. VISA enables visual source attribution capability while the effectiveness of answer generation is preserved at about the same level of effectiveness.

6.4 Error Analysis

We conducted an error analysis on 50 randomly sampled cases from Wiki-VISA to better understand the limitations of VISA. Errors were categorized into three main types as demonstrated in Figure 2. The first type, wrong source attribution, occurred in 43 cases where the model attributed the source to an incorrect section of the document, failing to identify the precise region containing the evidence. The second type, position misalignment, was observed in 4 cases where the model appeared to have the correct intent but drew the bounding box inaccurately, either slightly off position or incorrectly sized. The third type, granularity mismatch, appeared in 3 cases where the model's attributed source, such as a specific cell in a table or an item in a list, did not match the ground truth granularity. While these cases could potentially be considered false negatives, we leave it in error analysis to emphasize the challenge in real-world use cases where user preferences for granularity may differ from the model's output.

7 Conclusion

In this paper, we introduced VISA, a visual source attribution approach for retrieval-augmented generation pipeline. By leveraging vision-language models, VISA not only generates answers to user queries but also provides bounding boxes that visually attribute the supporting evidence within document screenshots. This capability enhances transparency and supports users in verifying the generated information effectively. Through the development of curated datasets, we demonstrated the effectiveness of VISA across diverse document types and layouts, including complex multi-page documents and multimodal content. Our experimental results highlight the potential of VISA to bridge the gap between information retrieval and answer generation by offering finer-grained, visually grounded evidence attribution. Moving forward, we hope VISA represents a pioneering step for more verifiable and user-friendly RAG systems.

8 Limitations

While VISA demonstrates promising results for answer generation and content grounding in vision-based RAG systems, it has several limitations. First, it focuses on generating short answers, which may not suffice for scenarios requiring detailed or explanatory responses, highlighting the need for enhancements in generating richer context. Second, it assumes answers are derived from a single, localized region within a document, which limits its effectiveness for cases where evidence spans multiple sections or modalities (e.g., combining text and tables). Third, while our evaluation spans web and medical scientific papers with various content modalities (e.g., passages, tables, figures), it does not fully capture the diversity of real-world documents such as scanned or handwritten content. Additionally, as VISA aims to make it intuitive for users to verify answers, conducting user studies could further confirm its practical utility.

References

Akari Asai, Zexuan Zhong, Danqi Chen, Pang Wei Koh, Luke Zettlemoyer, Hannaneh Hajishirzi, and Wen tau Yih. 2024. Reliable, adaptable, and attributable language models with retrieval. *Preprint*, arXiv:2403.03187.

Bernd Bohnet, Vinh Q. Tran, Pat Verga, Roee Aharoni, Daniel Andor, Livio Baldini Soares, Massimiliano Ciaramita, Jacob Eisenstein, Kuzman Ganchev, Jonathan Herzig, Kai Hui, Tom Kwiatkowski, Ji Ma, Jianmo Ni, Lierni Sestorain Saralegui, Tal Schuster, William W. Cohen, Michael Collins, Dipanjan Das, Donald Metzler, Slav Petrov, and Kellie Webster. 2023. Attributed question answering: Evaluation and modeling for attributed large language models. *Preprint*, arXiv:2212.08037.

Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. 2020. End-to-end object detection with transformers. In *Computer Vision – ECCV 2020*, pages 213–229, Cham. Springer International Publishing.

Danqi Chen, Adam Fisch, Jason Weston, and Antoine Bordes. 2017. Reading Wikipedia to answer open-domain questions. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1870–1879.

Kanzhi Cheng, Qiushi Sun, Yougang Chu, Fangzhi Xu, Li YanTao, Jianbing Zhang, and Zhiyong Wu. 2024. SeeClick: Harnessing GUI grounding for advanced visual GUI agents. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 9313–9332, Bangkok, Thailand. Association for Computational Linguistics.

Jifeng Dai, Yi Li, Kaiming He, and Jian Sun. 2016. R-fcn: Object detection via region-based fully convolutional networks. In *Advances in Neural Information Processing Systems*, volume 29. Curran Associates, Inc.

Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. 2021. An image is worth 16x16 words: Transformers for image recognition at scale. In *International Conference on Learning Representations*.

Manuel Faysse, Hugues Sibille, Tony Wu, Bilel Omrani, Gautier Viaud, Céline Hudelot, and Pierre Colombo. 2024. Colpali: Efficient document retrieval with vision language models. *Preprint*, arXiv:2407.01449.

Jeremy J. Foster. 1979. *The Use of Visual Cues in Text*, pages 189–203. Springer US, Boston, MA.

Tianyu Gao, Howard Yen, Jiatong Yu, and Danqi Chen. 2023. Enabling large language models to generate text with citations. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 6465–6488, Singapore. Association for Computational Linguistics.

Yunfan Gao, Yun Xiong, Xinyu Gao, Kangxiang Jia, Jinliu Pan, Yuxi Bi, Yi Dai, Jiawei Sun, Meng Wang, and Haofen Wang. 2024. Retrieval-augmented generation for large language models: A survey. *arXiv:2312.10997*.

Muhammad Khalifa, David Wadden, Emma Strubell, Honglak Lee, Lu Wang, Iz Beltagy, and Hao Peng. 2024. Source-aware training enables knowledge attribution in language models. In *First Conference on Language Modeling*.

Tom Kwiatkowski, Jennimaria Palomaki, Olivia Redfield, Michael Collins, Ankur Parikh, Chris Alberti, Danielle Epstein, Illia Polosukhin, Jacob Devlin, Kenton Lee, Kristina Toutanova, Llion Jones, Matthew Kelcey, Ming-Wei Chang, Andrew M. Dai, Jakob Uszkoreit, Quoc Le, and Slav Petrov. 2019. Natural Questions: A benchmark for question answering research. *Transactions of the Association for Computational Linguistics*, 7:452–466.

Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. 2015. Deep learning. *nature*, 521(7553):436–444.

Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, Sebastian Riedel, and Douwe Kiela. 2020.

Retrieval-augmented generation for knowledge-intensive nlp tasks. In *Proceedings of the 34th International Conference on Neural Information Processing Systems*, NIPS '20, Red Hook, NY, USA. Curran Associates Inc.

Kevin Qinghong Lin, Linjie Li, Difei Gao, Zhengyuan Yang, Zechen Bai, Weixian Lei, Lijuan Wang, and Mike Zheng Shou. 2024. Showui: One vision-language-action model for generalist gui. In *NeurIPS 2024 Workshop on Open-World Agents*.

Xueguang Ma, Sheng-Chieh Lin, Minghan Li, Wenhu Chen, and Jimmy Lin. 2024. Unifying multimodal retrieval via document screenshot embedding. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 6492–6505, Miami, Florida, USA. Association for Computational Linguistics.

Chaitanya Malaviya, Subin Lee, Sihao Chen, Elizabeth Sieber, Mark Yatskar, and Dan Roth. 2024. ExpertQA: Expert-curated questions and attributed answers. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 3025–3045, Mexico City, Mexico. Association for Computational Linguistics.

Sewon Min, Kalpesh Krishna, Xinxi Lyu, Mike Lewis, Wen-tau Yih, Pang Koh, Mohit Iyyer, Luke Zettlemoyer, and Hannaneh Hajishirzi. 2023. FActScore: Fine-grained atomic evaluation of factual precision in long form text generation. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 12076–12100, Singapore. Association for Computational Linguistics.

OpenAI, :, Aaron Hurst, Adam Lerer, Adam P. Goucher, Adam Perelman, Aditya Ramesh, Aidan Clark, AJ Ostrow, Akila Welihinda, Alan Hayes, Alec Radford, Aleksander Mądry, Alex Baker-Whitcomb, Alex Beutel, Alex Borzunov, Alex Carney, Alex Chow, Alex Kirillov, Alex Nichol, Alex Paino, Alex Renzin, Alex Tachard Passos, Alexander Kirillov, Alexi Christakis, Alexis Conneau, Ali Kamali, Allan Jabri, Allison Moyer, Allison Tam, Amadou Crookes, Amin Tootoochian, Amin Tootoonchian, Ananya Kumar, Andrea Vallone, Andrej Karpathy, Andrew Braunstein, Andrew Cann, Andrew Codispoti, Andrew Galu, Andrew Kondrich, Andrew Tulloch, Andrey Mishchenko, Angela Baek, Angela Jiang, Antoine Pelisse, Antonia Woodford, Anuj Gosalia, Arka Dhar, Ashley Pantuliano, Avi Nayak, Avital Oliver, Barret Zoph, Behrooz Ghorbani, Ben Leimberger, Ben Rossen, Ben Sokolowsky, Ben Wang, Benjamin Zweig, Beth Hoover, Blake Samic, Bob McGrew, Bobby Spero, Bogo Giertler, Bowen Cheng, Brad Lightcap, Brandon Walkin, Brendan Quinn, Brian Guarraci, Brian Hsu, Bright Kellogg, Brydon Eastman, Camillo Lugaresi, Carroll Wainwright, Cary Bassin, Cary Hudson, Casey Chu, Chad Nelson, Chak Li, Chan Jun Shern, Channing Conger, Charlotte Barette, Chelsea Voss, Chen Ding, Cheng Lu, Chong Zhang, Chris Beaumont, Chris Hallacy, Chris Koch, Christian Gibson, Christina Kim, Christine Choi, Christine McLeavey, Christopher Hesse, Claudia Fischer, Clemens Winter, Coley Czarnecki, Colin Jarvis, Colin Wei, Constantin Koumouzelis, Dane Sherburn, Daniel Kappler, Daniel Levin, Daniel Levy, David Carr, David Farhi, David Mely, David Robinson, David Sasaki, Denny Jin, Dev Valladares, Dimitris Tsipras, Doug Li, Duc Phong Nguyen, Duncan Findlay, Edede Oiwoh, Edmund Wong, Ehsan Asdar, Elizabeth Proehl, Elizabeth Yang, Eric Antonow, Eric Kramer, Eric Peterson, Eric Sigler, Eric Wallace, Eugene Brevdo, Evan Mays, Farzad Khorasani, Felipe Petroski Such, Filippo Raso, Francis Zhang, Fred von Lohmann, Freddie Sulit, Gabriel Goh, Gene Oden, Geoff Salmon, Giulio Starace, Greg Brockman, Hadi Salman, Haiming Bao, Haitang Hu, Hannah Wong, Haoyu Wang, Heather Schmidt, Heather Whitney, Heewoo Jun, Hendrik Kirchner, Henrique Ponde de Oliveira Pinto, Hongyu Ren, Huiwen Chang, Hyung Won Chung, Ian Kivlichan, Ian O'Connell, Ian O'Connell, Ian Osband, Ian Silber, Ian Sohl, Ibrahim Okuyucu, Ikai Lan, Ilya Kostrikov, Ilya Sutskever, Ingmar Kanitscheider, Ishaan Gulrajani, Jacob Coxon, Jacob Menick, Jakub Pachocki, James Aung, James Betker, James Crooks, James Lennon, Jamie Kiros, Jan Leike, Jane Park, Jason Kwon, Jason Phang, Jason Teplitz, Jason Wei, Jason Wolfe, Jay Chen, Jeff Harris, Jenia Varavva, Jessica Gan Lee, Jessica Shieh, Ji Lin, Jiahui Yu, Jiayi Weng, Jie Tang, Jieqi Yu, Joanne Jang, Joaquin Quinonero Candela, Joe Beutler, Joe Landers, Joel Parish, Johannes Heidecke, John Schulman, Jonathan Lachman, Jonathan McKay, Jonathan Uesato, Jonathan Ward, Jong Wook Kim, Joost Huizinga, Jordan Sitkin, Jos Kraaijeveld, Josh Gross, Josh Kaplan, Josh Snyder, Joshua Achiam, Joy Jiao, Joyce Lee, Juntang Zhuang, Justyn Harriman, Kai Fricke, Kai Hayashi, Karan Singhal, Katy Shi, Kavin Karthik, Kayla Wood, Kendra Rimbach, Kenny Hsu, Kenny Nguyen, Keren Gu-Lemberg, Kevin Button, Kevin Liu, Kiel Howe, Krithika Muthukumar, Kyle Luther, Lama Ahmad, Larry Kai, Lauren Itow, Lauren Workman, Leher Pathak, Leo Chen, Li Jing, Lia Guy, Liam Fedus, Liang Zhou, Lien Mamitsuka, Lilian Weng, Lindsay McCallum, Lindsey Held, Long Ouyang, Louis Feuvrier, Lu Zhang, Lukas Kondraciuk, Lukasz Kaiser, Luke Hewitt, Luke Metz, Lyric Doshi, Mada Aflak, Maddie Simens, Madelaine Boyd, Madeleine Thompson, Marat Dukhan, Mark Chen, Mark Gray, Mark Hudnall, Marvin Zhang, Marwan Aljubeh, Mateusz Litwin, Matthew Zeng, Max Johnson, Maya Shetty, Mayank Gupta, Meghan Shah, Mehmet Yatbaz, Meng Jia Yang, Mengchao Zhong, Mia Glaese, Mianna Chen, Michael Janner, Michael Lampe, Michael Petrov, Michael Wu, Michele Wang, Michelle Fradin, Michelle Pokrass, Miguel Castro, Miguel Oom Temudo de Castro, Mikhail Pavlov, Miles Brundage, Miles Wang, Minal Khan, Mira Murati, Mo Bavarian, Molly Lin, Murat Yesildal, Nacho Soto, Natalia Gimelshein, Natalie Cone, Natalie Staudacher, Natalie Summers, Natan LaFontaine, Neil Chowdhury, Nick Ryder, Nick Stathas, Nick Turley, Nik Tezak, Niko Felix, Nithanth Kudige, Nitish Keskar, Noah Deutsch, Noel Bundick, Nora Puckett, Ofir Nachum, Ola Okelola, Oleg Boiko, Oleg Murk, Oliver Jaffe, Olivia Watkins, Olivier Godement, Owen Campbell-Moore, Patrick Chao, Paul McMillan, Pavel Belov, Peng Su, Peter Bak, Peter Bakkum, Peter Deng, Peter Dolan, Peter Hoeschele, Peter Welinder, Phil Tillet, Philip Pronin, Philippe Tillet, Prafulla Dhariwal, Qiming Yuan, Rachel Dias, Rachel Lim, Rahul Arora, Rajan Troll, Randall Lin, Rapha Gontijo Lopes, Raul Puri, Reah Miyara, Reimar Leike, Renaud Gaubert, Reza Zamani, Ricky Wang, Rob Donnelly, Rob Honsby, Rocky Smith, Rohan Sahai, Rohit Ramchandani, Romain Huet, Rory Carmichael, Rowan Zellers, Roy Chen, Ruby Chen, Ruslan Nigmatullin, Ryan Cheu, Saachi Jain, Sam Altman, Sam Schoenholz, Sam Toizer, Samuel Miserendino, Sandhini Agarwal, Sara Culver, Scott Ethersmith, Scott Gray, Sean Grove, Sean Metzger, Shamez Hermani, Shantanu Jain, Shengjia Zhao, Sherwin Wu, Shino Jomoto, Shirong Wu, Shuaiqi, Xia, Sonia Phene, Spencer Papay, Srinivas Narayanan, Steve Coffey, Steve Lee, Stewart Hall, Suchir Balaji, Tal Broda, Tal Stramer, Tao Xu, Tarun Gogineni, Taya Christianson, Ted Sanders, Tejal Patwardhan, Thomas Cunninghman, Thomas Degry, Thomas Dimson, Thomas Raoux, Thomas Shadwell, Tianhao Zheng, Todd Underwood, Todor Markov, Toki Sherbakov, Tom Rubin, Tom Stasi, Tomer Kaftan, Tristan Heywood, Troy Peterson, Tyce Walters, Tyna Eloundou, Valerie Qi, Veit Moeller, Vinnie Monaco, Vishal Kuo, Vlad Fomenko, Wayne Chang, Weiyi Zheng, Wenda Zhou, Wesam Manassra, Will Sheu, Wojciech Zaremba, Yash Patil, Yilei Qian, Yongjik Kim, Youlong Cheng, Yu Zhang, Yuchen He, Yuchen Zhang, Yujia Jin, Yunxing Dai, and Yury Malkov. 2024. Gpt-4o system card. *arXiv:2410.21276*.

Oded Ovadia, Menachem Brief, Moshik Mishaeli, and Oren Elisha. 2024. Fine-tuning or retrieval? comparing knowledge injection in LLMs. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 237–250, Miami, Florida, USA. Association for Computational Linguistics.

Guilherme Penedo, Hynek Kydlíček, Loubna Ben allal, Anton Lozhkov, Margaret Mitchell, Colin Raffel, Leandro Von Werra, and Thomas Wolf. 2024. The fineweb datasets: Decanting the web for the finest text data at scale. *arXiv:2406.17557*.

Hannah Rashkin, Vitaly Nikolaev, Matthew Lamm, Lora Aroyo, Michael Collins, Dipanjan Das, Slav Petrov, Gaurav Singh Tomar, Iulia Turc, and David Reitter. 2023. Measuring attribution in natural language generation models. *Computational Linguistics*, 49(4):777–840.

Joseph Redmon, Santosh Divvala, Ross Girshick, and Ali Farhadi. 2016. You only look once: Unified, real-time object detection. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 779–788.

Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. 2015. Faster r-cnn: Towards real-time object detection with region proposal networks. In *Advances in Neural Information Processing Systems*, volume 28. Curran Associates, Inc.

Monica Riedler and Stefan Langer. 2024. Beyond text: Optimizing rag with multimodal inputs for industrial applications. *Preprint*, arXiv:2410.21943.

Andreas Steiner, André Susano Pinto, Michael Tschannen, Daniel Keysers, Xiao Wang, Yonatan Bitton, Alexey Gritsenko, Matthias Minderer, Anthony Sherbondy, Shangbang Long, Siyang Qin, Reeve Ingle, Emanuele Bugliarello, Sahar Kazemzadeh, Thomas Mesnard, Ibrahim Alabdulmohsin, Lucas Beyer, and Xiaohua Zhai. 2024. Paligemma 2: A family of versatile vlms for transfer. *arXiv:2412.03555*.

John Sweller. 2011. Chapter two - cognitive load theory. volume 55 of *Psychology of Learning and Motivation*, pages 37–76. Academic Press.

Peng Wang, Shuai Bai, Sinan Tan, Shijie Wang, Zhihao Fan, Jinze Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Yang Fan, Kai Dang, Mengfei Du, Xuancheng Ren, Rui Men, Dayiheng Liu, Chang Zhou, Jingren Zhou, and Junyang Lin. 2024. Qwen2-vl: Enhancing vision-language model's perception of the world at any resolution. *arXiv:2409.12191*.

Peng Xia, Kangyu Zhu, Haoran Li, Tianze Wang, Weijia Shi, Sheng Wang, Linjun Zhang, James Zou, and Huaxiu Yao. 2024. Mmed-rag: Versatile multimodal rag system for medical vision language models. *Preprint*, arXiv:2410.13085.

Xi Ye, Ruoxi Sun, Sercan Arik, and Tomas Pfister. 2024. Effective large language model adaptation for improved grounding and citation generation. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 6237–6251, Mexico City, Mexico. Association for Computational Linguistics.

Shi Yu, Chaoyue Tang, Bokai Xu, Junbo Cui, Junhao Ran, Yukun Yan, Zhenghao Liu, Shuo Wang, Xu Han, Zhiyuan Liu, and Maosong Sun. 2024. Visrag: Vision-based retrieval-augmented generation on multi-modality documents. *Preprint*, arXiv:2410.10594.

Zhong-Qiu Zhao, Peng Zheng, Shou tao Xu, and Xindong Wu. 2019. Object detection with deep learning: A review. *Preprint*, arXiv:1807.05511.

Xu Zhong, Jianbin Tang, and Antonio Jimeno Yepes. 2019. PubLayNet: Largest Dataset Ever for Document Layout Analysis . In *2019 International Conference on Document Analysis and Recognition (ICDAR)*, pages 1015–1022, Los Alamitos, CA, USA. IEEE Computer Society.

Zhengxia Zou, Keyan Chen, Zhenwei Shi, Yuhong Guo, and Jieping Ye. 2023. Object detection in 20 years: A survey. *Preprint*, arXiv:1905.05055.

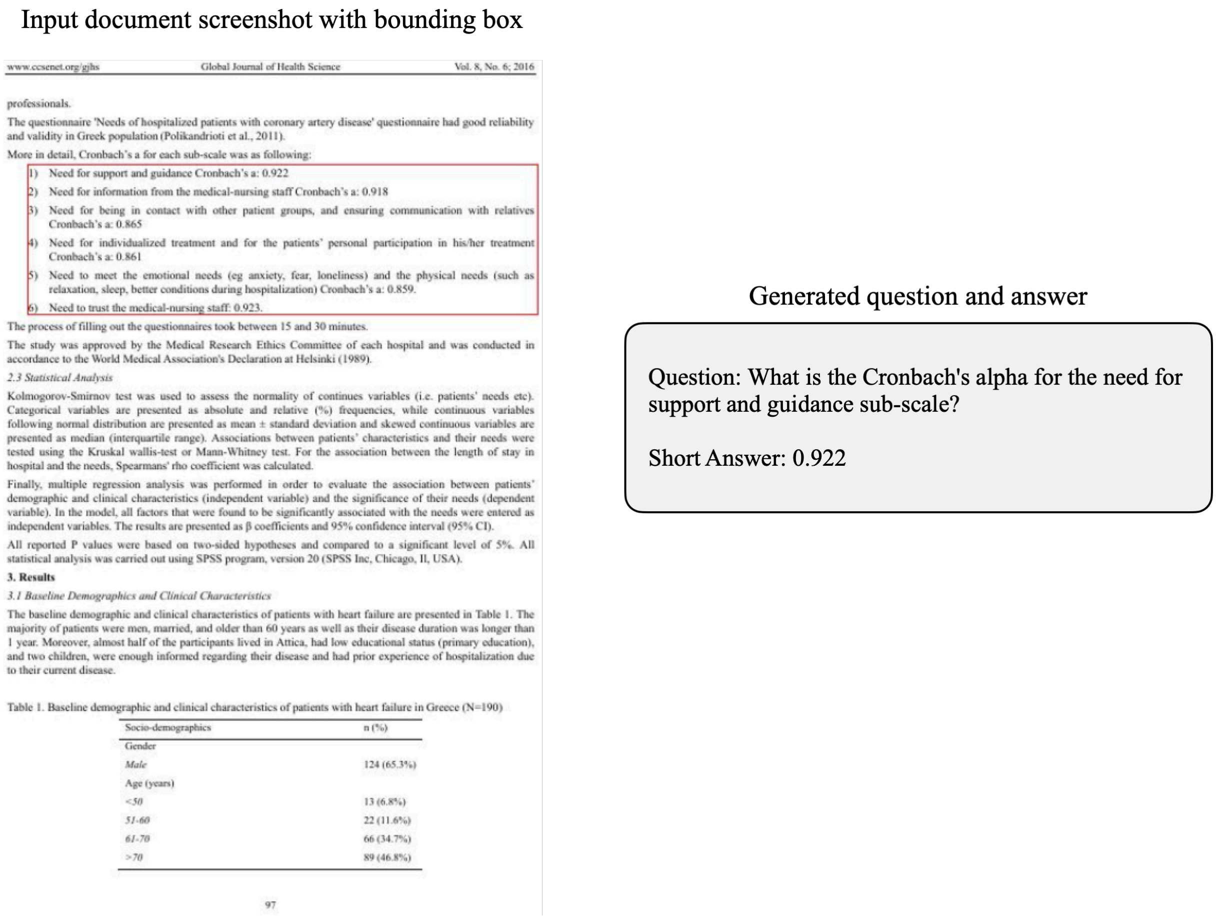

www.ccsenet.org/gjhs | Global Journal of Health Science | Vol. 8, No. 6; 2016

professionals.

The questionnaire 'Needs of hospitalized patients with coronary artery disease' questionnaire had good reliability and validity in Greek population (Polikandrioti et al., 2011).

More in detail, Cronbach's a for each sub-scale was as following:

1) Need for support and guidance Cronbach's a: 0.922
2) Need for information from the medical-nursing staff Cronbach's a: 0.918
3) Need for being in contact with other patient groups, and ensuring communication with relatives Cronbach's a: 0.865
4) Need for individualized treatment and for the patients' personal participation in his/her treatment Cronbach's a: 0.861
5) Need to meet the emotional needs (eg anxiety, fear, loneliness) and the physical needs (such as relaxation, sleep, better conditions during hospitalization) Cronbach's a: 0.859.
6) Need to trust the medical-nursing staff: 0.923.

The process of filling out the questionnaires took between 15 and 30 minutes.

The study was approved by the Medical Research Ethics Committee of each hospital and was conducted in accordance to the World Medical Association's Declaration at Helsinki (1989).

2.3 Statistical Analysis

Kolmogorov-Smirnov test was used to assess the normality of continues variables (i.e. patients' needs etc). Categorical variables are presented as absolute and relative (%) frequencies, while continuous variables following normal distribution are presented as mean ± standard deviation and skewed continuous variables are presented as median (interquartile range). Associations between patients' characteristics and their needs were tested using the Kruskal wallis-test or Mann-Whitney test. For the association between the length of stay in hospital and the needs, Spearmans' rho coefficient was calculated.

Finally, multiple regression analysis was performed in order to evaluate the association between patients' demographic and clinical characteristics (independent variable) and the significance of their needs (dependent variable). In the model, all factors that were found to be significantly associated with the needs were entered as independent variables. The results are presented as β coefficients and 95% confidence interval (95% CI).

All reported P values were based on two-sided hypotheses and compared to a significant level of 5%. All statistical analysis was carried out using SPSS program, version 20 (SPSS Inc, Chicago, Il, USA).

3. Results

3.1 Baseline Demographics and Clinical Characteristics

The baseline demographic and clinical characteristics of patients with heart failure are presented in Table 1. The majority of patients were men, married, and older than 60 years as well as their disease duration was longer than 1 year. Moreover, almost half of the participants lived in Attica, had low educational status (primary education), and two children, were enough informed regarding their disease and had prior experience of hospitalization due to their current disease.

Table 1. Baseline demographic and clinical characteristics of patients with heart failure in Greece (N=190)

Socio-demographics	n (%)
Gender	
Male	124 (65.3%)
Age (years)	
<50	13 (6.8%)
51-60	22 (11.6%)
61-70	66 (34.7%)
>70	89 (46.8%)

97

Figure 3: An example of synthetic data from Paper-VISA.

A Appendix

A.1 Prompt for synthetic data generation

The following prompt was used for prompting QWen2-VL-72B to generate synthetic questions and answers for Paper-VISA and Fineweb-VISA datasets.

System:
Ask a question that can be specifically answered by the content in the red bounding box area and give a short answer. The question can be a wh- question, a yes/no question, or a how question, that can be answered in a few words.
Output format:

Question: <question>
Short Answer: <short answer>

Or simply return 'Empty' if the bounding box area is not visible or informative.

User: {image}

Figure 3 shows an example of synthetic data from Paper-VISA.

A.2 Prompt for Single Oracle candidate VISA

The following prompt template was used to format the model's inputs and outputs for training the *Single Oracle Candidate* VISA.

```
Model Input:
System:
Given a document image, your task is to answer the question and locate the source of the answer
via a bounding box.

User:
{image} Image Size: {image.size}
Question: {question}

Model Output:
Assistant:
Answer: {answer}
Bounding Box: {bounding_box}
```

A.3 Prompt for Multi-candidate VISA

The following prompt template was used to format the model's inputs and outputs for training the *Multi-candidate* VISA.

```
Model Input:
System:
Given document images, your task is to answer the question and locate the source of the answer
via a bounding box.

User:
{image1} Image Size: {image1.size}
{image2} Image Size: {image2.size}
{image3} Image Size: {image3.size}
Question: {question}

Model Output:
Assistant:
Answer: {answer}
Evidence Document: {index}
Bounding Box: {bounding_box}
```

A.4 Dataset Licenses

- **NQ**: Apache License 2.0
- **Wikipedia**: Creative Commons Attribution Share Alike, GNU Free Documentation License family.
- **Fineweb-edu**: Open Data Commons License Attribution family.
- **PubLayNet**: Community Data License Agreement – Permissive, Version 1.0.
- **VISA Datasets**: Our crafted datasets follow the same license as the source of the documents.

A.5 Model Backbone Licenses

- **QWen2-VL-72B**: Qwen LICENSE AGREEMENT.
- **QWen2-VL-2B**: Apache License.
- **QWen2-VL-7B**: Apache License.
- **VISA Models**: Our fine-tuned models follow the same licenses as the original model backbone.

A.6 AI Assistant Usage

GPT4o is used during the writing to correct grammar errors and format tables.