

DOCLENS : A Tool-Augmented Multi-Agent Framework for Long Visual Document Understanding

Dawei Zhu[♣] Rui Meng[♣] Jiefeng Chen[♣] Sujian Li^{♣*}
Tomas Pfister[♣] Jinsung Yoon^{*♣}

[♣] School of Computer Science, Peking University

[♣] Google Cloud AI Research

{dwzhu, lisujian}@pku.edu.cn {rmeng, jiefengc, tpfister, jinsungyoon}@google.com

<https://dwzhu-pku.github.io/DocLens/>

Abstract

Comprehending long visual documents, where information is distributed across extensive pages of text and visual elements, is a critical but challenging task for modern Vision-Language Models (VLMs). Existing approaches falter on a fundamental challenge: evidence localization. They struggle to retrieve relevant pages and overlook fine-grained details within visual elements, leading to limited performance and model hallucination. To address this, we propose DOCLENS, a tool-augmented multi-agent framework that effectively “zooms in” on evidence like a lens. It first navigates from the full document to specific visual elements on relevant pages, then employs a sampling-adjudication mechanism to generate a single, reliable answer. Paired with Gemini-2.5-Pro, DOCLENS achieves state-of-the-art performance on MMLongBench-Doc and FinRAGBench-V, surpassing even human experts. The framework’s superiority is particularly evident on vision-centric and unanswerable queries, demonstrating the power of its enhanced localization capabilities.

1 Introduction

A vast repository of human knowledge is encapsulated in long visual documents such as financial reports, academic papers, and technical manuals (Liu et al., 2025). With information synthesized from various textual and visual elements (tables, charts, figures) distributed throughout the context, these long visual documents are formidably challenging to decipher, even for the most advanced Vision-Language Models (VLMs) (Comanici et al., 2025; ClaudeTeam, 2025; OpenAITeam, 2025; QwenTeam, 2025; Team et al., 2025; Guo et al., 2025).

This challenge stems from a fundamental problem: evidence localization. Existing efforts to localizing evidence from long visual documents

primarily operate at the page level, either feeding page screenshots to long-context VLMs (Ma et al., 2024b) or employing vector-based retrieval methods (Han et al., 2025; Cho et al., 2024). However, we observe that both approaches perform poorly in recalling evidence pages. On MMLongBench-Doc (Ma et al., 2024b), Gemini-2.5-Pro (Comanici et al., 2025) only recalls 68% of evidence pages, while vector-based methods using ColBERT (Santhanam et al., 2021) and ColPali (Faysse et al., 2024) achieve merely 55.3% Recall@10. This fundamental failure prevents models from producing accurate answers. Moreover, even on the correct page, crucial details within visual elements (e.g., charts, tables) remain obscured in a full-page view, akin to reading a map without a magnifying glass. This dual-level failure in evidence localization—at both the page and element scale—directly fuels model hallucination, causing models to invent responses for over half of unanswerable queries rather than admitting uncertainty on MMLongBench-Doc (Figure 1b).

In this paper, we propose DOCLENS, a multi-agent framework that overcomes these challenges by strategically leveraging document-parsing tools. Our core component is the *Lens Module*, which zooms into long visual documents like a lens to perform fine-grained evidence localization (Figure 1a). It includes a *Page Navigator* agent and an *Element Localizer* agent. The former uses OCR tools to augment VLMs for page-level retrieval, drastically improving recall of evidence pages; the latter employs layout detection and cropping tools to locate visual elements on these retrieved pages for detailed inspection. Following the Lens Module, the Reasoning Module synthesizes the extracted evidence—including page screenshots, text, and cropped visual elements—to formulate a final answer. To ensure both accuracy and reliability, this module employs a “sampling-adjudication” process that first proposes a set of potential answers

*Corresponding authors.

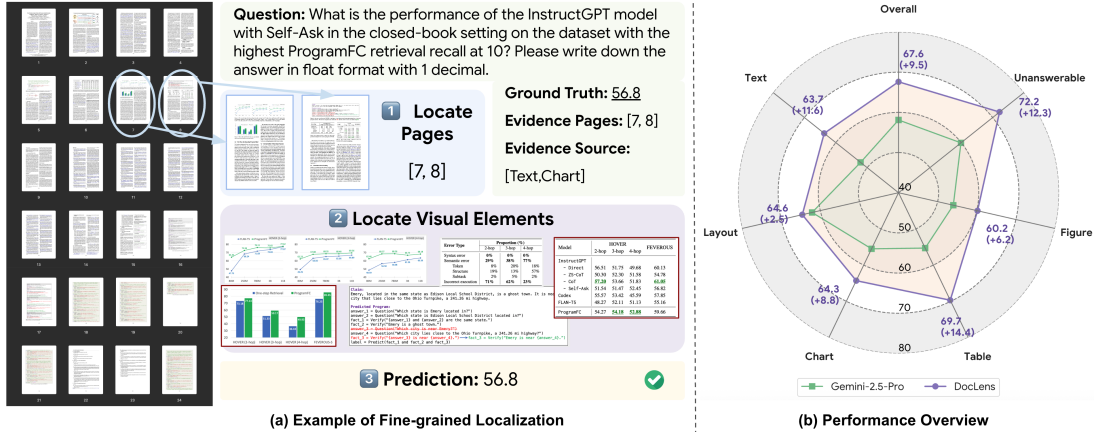


Figure 1: Workflow and performance of our proposed method, DOCLENS. (a) The workflow grounds its answer by navigating from the full document to visual elements (e.g., Text, Chart) within relevant pages. (b) It yields great improvement on MMLongBench-Doc, specifically for understanding visual elements and reducing hallucination.

using an *Answer Sampler* agent, and then critically assesses them using an *Adjudicator* agent to select the best candidate.

We evaluate DOCLENS on two challenging benchmarks, MMLongBench-Doc (Ma et al., 2024b) and FinRAGBench-V (Zhao et al., 2025). Our method achieves state-of-the-art performance, significantly reduces hallucination, and for the first time surpasses human experts. This breakthrough is driven by the efficacy of our core components: further analysis reveals that our Page Navigator achieves near-perfect evidence page recall (97.3%), while the Element Localizer dramatically enhances the comprehension of fine-grained visual details. Our main contributions are threefold:

- A novel, tool-augmented Lens Module that achieves near-perfect page recall and enables fine-grained inspection of visual elements, effectively solving evidence localization.
- A sampling-adjudication mechanism within the Reasoning Module that effectively mitigates hallucination and improves answer reliability.
- The establishment of a new state-of-the-art on MMLongBench-Doc and FinRAGBench-V, and for the first time, surpassing human experts.

2 Problem Formulation

We address the challenge of question answering over long visual documents. A document is a sequence of pages, $\mathcal{D} = \{P_i\}_{i=1}^N$, where each page P_i is a screenshot image. From each page, we can extract text T_i and a set of visual elements \mathcal{V}_i (e.g., tables, figures). Given a question Q , the goal is to

generate an accurate answer A that is grounded in a specific set of evidence pages $\mathcal{E} \subseteq \mathcal{D}$. This task can be abstractly formulated as learning a function f that maps the document and question to an answer:

$$A = f(\mathcal{D}, Q),$$

However, the sheer volume of information in a long document makes a direct mapping challenging to construct. We contend that a more principled approach is to decompose the problem into two stages: First, identifying a concise set of relevant evidence from the vast document, and second, generating the answer based on this evidence.

To formalize this, we factorize the function f into two components. First, an extraction function f_{extract} reads through the document to identify a concise evidence set \mathcal{S} relevant to the question:

$$\mathcal{S} = f_{\text{extract}}(\mathcal{D}, Q),$$

This evidence set \mathcal{S} contains the necessary pages from \mathcal{D} and the visual and textual elements within these pages. Second, an answer generation function f_{generate} infers the final answer exclusively from this condensed evidence:

$$A = f_{\text{generate}}(\mathcal{S}, Q),$$

We therefore model the composite function as:

$$A = f_{\text{generate}}(f_{\text{extract}}(\mathcal{D}, Q), Q),$$

The goal is to design and optimize both f_{extract} and f_{generate} to maximize the accuracy of the predicted answer A w.r.t. the ground-truth answer A^* .

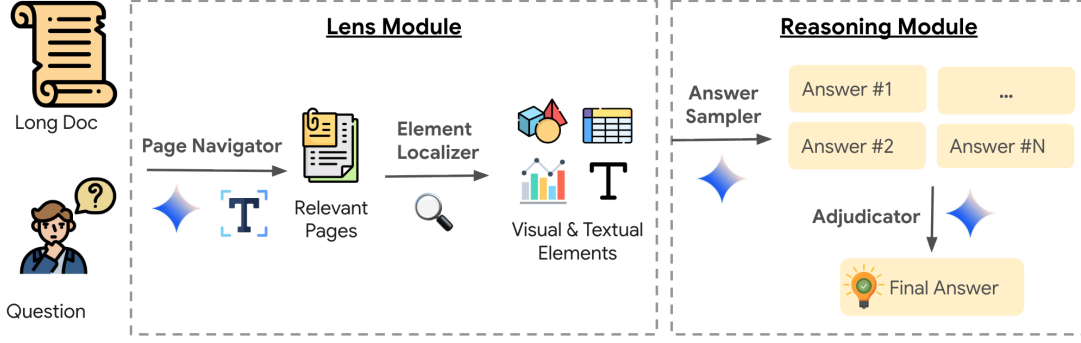


Figure 2: Overall workflow of DOCLENS. Given a long visual document and a corresponding question, we first apply a *Lens Module* to retrieve relevant pages and locate relevant visual&textual elements within these pages. We then use a *Reasoning Module* to do in-depth analysis of these elements and provide an accurate answer.

3 DOCLENS Framework

Figure 2 illustrates the overall workflow of DOCLENS. Our proposed framework consists of two primary components: a *Lens Module* and a *Reasoning Module*. Given a long visual document and an associated question, the *Lens Module* (f_{extract}) is responsible for identifying relevant pages and the key elements within them. Subsequently, the *Reasoning Module* (f_{generate}) conducts an in-depth analysis of this evidence to generate a precise answer. The prompt templates for all agents and the pseudocode for the entire workflow are presented in Appendices A and D, respectively.

3.1 Lens Module

Page Navigator. The *Lens Module* begins with the *Page Navigator* to identify a predicted set of evidence pages, $\mathcal{E}_{\text{pred}}$, from the full document $\mathcal{D} = \{P_i\}_{i=1}^N$. First, it uses an OCR tool to extract the text T_i from every page $P_i \in \mathcal{D}$.

$$T_i = \text{OCR}(P_i), \quad i = \{1, 2, \dots, N\},$$

To locate potential evidence pages $\mathcal{E} \subseteq \mathcal{D}$, the *Page Navigator* then prompts an LLM with the question Q , all page screenshots and their OCR text (as interleaved input). To ensure comprehensive coverage, this process is repeated T_e times with a temperature $\tau > 0$. Each sampling iteration j generates a candidate page set $\mathcal{E}^{(j)}$:

$$\mathcal{E}^{(j)} = \text{LLM}_{\text{PageNav}}(Q, \{(P_i, T_i)\}_{i=1}^N),$$

The final set of predicted pages is the union of all pages identified across these iterations:

$$\mathcal{E}_{\text{pred}} = \bigcup_{j=1}^{T_e} \mathcal{E}^{(j)},$$

In practice, an LLM’s finite context window may prevent processing all N pages simultaneously. In such cases, we divide pages into chunks, process them in parallel, and merge the resulting $\mathcal{E}_{\text{pred}}$ sets.

Element Localizer. Given the set of predicted pages $\mathcal{E}_{\text{pred}}$ identified by the *Page Navigator*, the *Element Localizer* enriches this set by parsing detailed visual and textual elements.

For each page $P_k \in \mathcal{E}_{\text{pred}}$, its corresponding textual content T_k is available from the prior step. Concurrently, a layout detection tool identifies the bounding boxes of key visual elements (such as figures, charts, and tables). These elements are then cropped from the page to form a set of focused visual inputs, denoted as \mathcal{V}_k :

$$\mathcal{V}_k = \{\text{Crop}(P_k, \text{bbox}) \mid \text{bbox} \in \text{LayoutDetect}(P_k)\},$$

where bbox is the bounding box of each visual element. Then with all predicted evidence pages, we construct the full evidence set \mathcal{S} by collecting tuples of the page screenshot (P_k), its extracted text (T_k), and its cropped visual elements (\mathcal{V}_k):

$$\mathcal{S} = \{(P_k, T_k, \mathcal{V}_k) \mid P_k \in \mathcal{E}_{\text{pred}}\}$$

3.2 Reasoning Module

Answer Sampler. The *Answer Sampler* agent receives the collection of evidence \mathcal{S} extracted by the *Lens Module*. It then integrates all this information to generate a reasoning process R (e.g., a chain-of-thought trace) and a corresponding answer A :

$$R, A = \text{LLM}_{\text{Sampler}}(Q, \mathcal{S}),$$

To generate a diverse set of candidate answers (Wang et al., 2023), we perform this reasoning process T_a times. The diversity is achieved

by strategy with a temperature $\tau > 0$. This encourages the model to explore different reasoning paths and wording, yielding T_a distinct reasoning-answer pairs: $\{R_i, A_i\}_{i=1}^{T_a}$.

Adjudicator. The final step is managed by the Adjudicator, whose goal is to synthesize the best answer from the T_a candidate answers. It carefully analyzes the reasoning path R_i of each candidate and cross-validates the different approaches to identify the most consistent and logical conclusion, which is then presented as the final answer A_{final} :

$$A_{final} = \text{LLM}_{\text{Adjud}}(\{(R_i, A_i)\}_{i=1}^{T_a}).$$

4 Experiments

In this section, we first introduce our experimental setup (§ 4.1), including benchmarks, metrics, tested models and baseline methods. We then demonstrate the overall effectiveness of our DOCLENS on the two selected benchmarks (§ 4.2 & § 4.3). For implementation details, please refer to Appendix D.

4.1 Experimental Setup

Benchmarks and Metrics. We evaluate our method on two challenging benchmarks: MMLongBench-Doc (Ma et al., 2024b) and FinRAGBench-V (Zhao et al., 2025). MMLongBench-Doc tests reasoning over lengthy, multi-domain documents (avg. 47.5 pages) that require integrating scattered information across diverse modalities. Crucially, its dedicated “Unanswerable” subset directly evaluates our model’s ability to mitigate hallucination. Performance of human experts on this benchmark is reported as 65.8. FinRAGBench-V is vital for our analysis due to two unique features: its use of documents with dense, newspaper-like layouts (See Figure 5), and its support for evaluating visual citation (pinpointing block-level evidence), which provides a direct assessment of our fine-grained localization strategy. We adhere to the original evaluation protocols: rule-based scoring for MMLongBench-Doc and an LLM-as-a-judge approach for FinRAGBench-V. Further statistics and evaluation details are provided in Appendices B, D.3, and D.4.

Models and Baselines. We evaluate our proposed agentic framework on three cutting-edge proprietary models: Gemini-2.5-Pro (Comanici et al., 2025), Gemini-2.5-Flash (Comanici et al., 2025), and Claude-4-Sonnet (ClaudeTeam, 2025). We benchmark its performance against three categories

of baselines. The first is the vanilla setting, which uses only page screenshots. The second augments screenshots with OCR text appended to each page, an approach we found particularly effective during our pilot study. The third category comprises existing agentic frameworks: MACT (Yu et al., 2025), M3DocRAG (Cho et al., 2024), MDocAgent (Han et al., 2025), and SimpleDoc (Jain et al., 2025). For MACT, M3DocRAG, and MDocAgent, we report the best scores from their original papers. For SimpleDoc, the most recent and best-performing training-free framework, we reproduce results across all proprietary models using our metrics to ensure fair comparison.

4.2 Main Results on MMLongBench-Doc

Table 1 presents our main experimental results. On MMLongBench-Doc, our approach yields substantial performance improvements across all three backbone models. These gains are particularly pronounced for comparatively weaker models, such as Claude-4-Sonnet and Gemini-2.5-Flash, compared to the more powerful Gemini-2.5-Pro. Notably, our DOCLENS framework enables both Claude-4-Sonnet and Gemini-2.5-Flash to achieve near-human performance. Furthermore, Gemini-2.5-Pro augmented with our method surpasses the human baseline by $\sim 2\%$. These results strongly demonstrate the effectiveness of DOCLENS.

Additionally, our method achieves significant improvements on the Unanswerable (UNA) subset, with absolute gains of +8.1%, +13.0%, and +13.8% for Claude-4-Sonnet, Gemini-2.5-Flash, and Gemini-2.5-Pro, respectively. This indicates that our agentic framework effectively mitigates model hallucination, a critical capability for real-world applications.

Finally, we observe that augmenting models with OCR text substantially improves performance across all three backbones compared to the vanilla setting. We attribute this improvement to OCR’s effectiveness in facilitating (implicit) evidence page retrieval, as further analyzed in Section 5.2.

4.3 Main Results on FinRAGBench-V

On the FinRAGBench-V benchmark, our framework’s superiority is even more pronounced. Compared to the strongest baseline, DOCLENS achieves substantial gains when paired with Claude-4-Sonnet (+3.1%), Gemini-2.5-Flash (+10.2%), and Gemini-2.5-Pro (+6.8%). We hypothesize these larger gains stem from FinRAGBench-V’s higher

Model	MMLongBench-Doc							FinRAGBench-V			
	TXT	LAY	CHA	TAB	FIG	UNA	ALL	TXT	TAB	CHA	ALL
<i>Vanilla VLMs</i>											
GPT-4o [†]	46.3	46.0	45.3	50.0	44.1	20.2	42.8	-	-	-	37.2
Claude-4-Sonnet	50.4	49.4	50.5	<u>57.3</u>	43.9	59.0	53.4	36.6	20.2	<u>51.9</u>	33.8
Gemini-2.5-Flash	44.0	53.2	46.0	<u>43.9</u>	48.2	56.7	49.6	49.0	41.6	41.0	43.0
Gemini-2.5-Pro	<u>52.1</u>	<u>62.1</u>	<u>55.5</u>	55.3	<u>54.0</u>	<u>59.9</u>	<u>58.1</u>	<u>62.2</u>	<u>55.3</u>	50.4	54.9
o4-mini [†]	-	-	-	-	-	-	-	-	-	-	<u>62.4</u>
<i>VLMs Augmented with OCR</i>											
Claude-4-Sonnet	52.7	51.6	50.0	58.1	45.3	<u>65.9</u>	56.0	58.7	21.6	54.3	41.0
Gemini-2.5-Flash	55.9	54.9	52.7	63.4	50.3	<u>60.8</u>	58.5	67.6	64.4	46.1	58.3
Gemini-2.5-Pro	<u>59.7</u>	65.3	<u>60.8</u>	<u>68.3</u>	<u>55.7</u>	58.4	<u>63.3</u>	<u>70.0</u>	<u>70.0</u>	<u>56.2</u>	<u>64.9</u>
<i>VLM-based Agentic Frameworks</i>											
MACT (w/ MiMo-VL-7B) [†]	-	-	-	-	-	-	47.4	-	-	-	-
M3DocRAG (w/ Qwen2-VL-7B) [†]	30.0	23.5	18.9	20.1	20.8	5.8	21.0	-	-	-	-
MDocAgent (w/ GPT-4o) [†]	-	-	-	-	-	-	42.0	-	-	-	-
SimpleDoc											
w/ Claude-4-Sonnet	52.1	53.3	58.3	62.4	46.9	66.5	58.6	59.6	68.9	54.9	61.7
w/ Gemini-2.5-Flash	45.5	57.4	49.0	51.6	45.2	66.5	53.3	70.2	56.2	53.6	58.3
w/ Gemini-2.5-Pro	48.4	54.8	55.7	56.1	52.5	59.7	56.6	67.5	64.0	60.9	63.6
DOCLENS (Ours)											
w/ Claude-4-Sonnet	59.9	58.2	54.4	63.9	55.3	74.0	63.3	70.2	66.0	60.3	64.8
w/ Gemini-2.5-Flash	59.5	61.5	54.8	66.9	59.0	73.8	64.7	69.9	71.3	64.5	68.5
w/ Gemini-2.5-Pro	63.7	<u>64.6</u>	64.3	69.7	60.2	72.2	67.6*	68.9	74.2	67.1	70.4

Table 1: Main Results on the MMLongBench-Doc and FinRAGBench-V benchmarks. We report the accuracy of five types of evidence sources including pure text (TXT), layout (LAY), chart (CHA), table (TAB), and figure (FIG), and on unanswerable (UNA) samples. **Bold** indicates the best score per column; underlined indicates the best per column within each block. [†] denotes results reported in the original paper, hence some results are unavailable. * Denotes results surpassing human experts (On MMLongBench-Doc, performance of human experts is 65.8).

proportion of documents with dense, complex visual layouts (e.g., newspapers).

A closer analysis confirms this hypothesis, revealing that the performance boost is primarily driven by our method’s superior handling of visual evidence. On chart-based questions, for instance, DOCLENS elevates the performance of Gemini-2.5-Pro and Gemini-2.5-Flash by absolute margins of +10.9% and +23.5% over the strong OCR-augmented baseline. This trend continues for table-based questions, with corresponding gains of +4.2% and +6.9%. Collectively, these results demonstrate that as visual complexity increases, the advantage of our fine-grained element localization becomes increasingly critical—a capability we analyze in further detail in Section 5.3.

5 Analysis

This section presents a comprehensive analysis of our framework. By default, all experiments are conducted with Gemini-2.5-Pro. We begin with an ablation study (§ 5.1), which confirms that the

Lens Module significantly boosts performance and the Reasoning Module can further reduce hallucination. We then delve deeper into the Lens Module (§ 5.2 and § 5.3), demonstrating through quantitative analysis and case studies how its Page Navigator improves page recall and its Element Localizer enhances visual comprehension by pinpointing specific elements. Next, we provide a detailed efficiency analysis (§ 5.4), showing that the sampling overhead is modest and our method achieves favorable latency-accuracy trade-offs compared to SimpleDoc. Finally, we demonstrate the framework’s cost efficiency via a hybrid-backbone variant that outperforms baseline with much lower cost (§ 5.5). We also discussed test-time scaling effect of our method in the Appendix E.1.

5.1 Ablation on Core Modules

Table 2 ablates the efficacy of our Lens Module and Reasoning Module. The ablation settings are as follows: to ablate the Lens Module, we provide the raw screenshot and OCR text directly to

Methods	MMLong			FinRAG
	ANS	UNA	ALL	
DOCLENS (Gemini-2.5-Pro)	66.4	<u>72.2</u>	67.6	70.4
w/o Lens Module	63.3	64.4	63.5	65.1
w/o Reasoning Module	66.6	68.2	67.0	69.9
DOCLENS (Gemini-2.5-Flash)	<u>62.4</u>	73.8	<u>64.7</u>	<u>68.5</u>
w/o Lens Module	58.4	68.4	60.4	60.9
w/o Reasoning Module	62.0	71.1	63.8	67.1

Table 2: Ablation study of key modules in our proposed method. *MMLong*, *FinRAG*, *ANS*, *UNA* is short for *MMLongBench-Doc*, *FinRAGBench-V*, *Answerable*, *Unanswerable*, respectively.

the Reasoning Module. To ablate the Reasoning Module, we take the output from the Lens Module (relevant pages and elements) and send it directly to the backbone VLM for answer generation.

The results underscore the critical role of the Lens Module. Its removal leads to a substantial performance drop of 4.1% on MMLongBench-Doc and 5.3% on FinRAGBench-V for Gemini-2.5-Pro, and similar degradation for Gemini-2.5-Flash (7.6% and 1.4%, respectively). The Reasoning Module, meanwhile, can further reduce model hallucination. Its absence leads to a noticeable drop in performance on the unanswerable (UNA) queries. Notably, the benefit of the Reasoning Module is more pronounced for weaker backbones; extended ablation across five benchmarks and three backbones is provided in Appendix C.1.

5.2 Analysis on Page Navigator

Page Navigator achieves nearly perfect recall of evidence pages. First, we demonstrate the effectiveness of our Page Navigator in terms of retrieving evidence pages. We assess this aspect in two ways: by calculating the recall of retrieved pages against the annotated evidence pages, and by measuring the final accuracy after processing these pages with Element Localizer, Answer Sampler, and Adjudicator. As presented in Table 3, on MMLongBench-Doc, our Page Navigator backbone with Gemini-2.5-Pro achieve a near-perfect recall of 97.3%, and its final accuracy is only 1.5% behind using oracle pages.

Page Navigator outperforms other retrievers. We then conduct a comparative analysis by substituting our Page Navigator with retrievers from leading prior work. The first and most prevalent category is vector-based retrievers (Dong et al., 2025b;

Han et al., 2025; Cho et al., 2024), exemplified by MDocAgent (Han et al., 2025), which retrieve the top-K most similar pages based on vector representations of the query and individual pages. We use ColBERT (Santhanam et al., 2021) for textual retrieval and ColPali (Faysse et al., 2024) for visual retrieval, combining their top-10 results. More recently, SimpleDoc (Jain et al., 2025) introduces a refinement through a two-stage pipeline: first retrieving pages via ColQwen2.5, then using an LLM to select evidence pages based on generated summaries. We substitute the original LLM with Gemini-2.5-Pro to maintain consistency with our framework’s backbone.

As shown in Table 3, MDocAgent’s vector-based retriever retrieves 13.6 pages on average but achieves low recall on evidence pages, resulting in the poorest final accuracy. SimpleDoc significantly improves both recall and precision, scoring 64.4% for final accuracy. However, it still underperforms our Page Navigator even when using the weakest backbone. These results validate the effectiveness of our Page Navigator for evidence page retrieval.

OCR significantly enhances page retrieval. We now examine each design choice of our Page Navigator. Table 4 presents ablation results on MMLongBench-Doc and FinRAGBench-V. We observe that both sampling and OCR augmentation improve evidence page recall, which successfully translates into higher final accuracy. Most notably, OCR provides substantial improvements: for Gemini-2.5-Pro, it increases recall by 10.0% and 40.5% on MMLongBench-Doc and FinRAGBench-V, respectively; for Gemini-2.5-Flash, the gains are 15.9% and 50.1%. This indicates that VLMs are more adept at retrieving relevant information in the textual domain than in the visual domain.

5.3 Analysis on Element Localizer

Element Localizer enhances block-level evidence identification. FinRAGBench-V provides 202 cases with human-annotated bounding boxes for relevant blocks on evidence pages. We leverage this subset to examine whether the element localizer improves block-level evidence identification. Specifically, we provide evidence pages to Gemini-2.5-Pro and compare bounding box predictions with and without the element localizer. As shown in Figure 4, the Element Localizer substantially improves block-level performance, increasing precision by 4.9%, recall by 9.3%, and F1 score by

Setting	#Pages	Recall	Prec	Final Acc
Evidence Pages (Oracle)	1.5	100.0	100.0	69.1
Baseline Retrievers				
MDocAgent’s Retriever	13.6	71.1	7.0	49.6
SimpleDoc’s Retriever	4.9	89.0	34.7	64.0
DOCLENS’s Page Navigator				
w/ Gemini-2.5-Pro	3.5	97.3	55.1	67.6
w/ Gemini-2.5-Flash	3.1	95.2	62.0	67.1
w/ Gemini-2.5-Flash-Lite	3.2	90.2	60.0	64.4

Table 3: Final Accuracy On MMLongBench-Doc with varying retrieval backbones for the Page Navigator. #Pages denotes average number of retrieved pages. Prec is short for Precision.

Methods	MMLong		FinRAG	
	Recall	Final Acc	Recall	Final Acc
Gemini-2.5-Pro				
Page Navigator	97.3	67.6	94.3	70.4
w/o Sampling	95.6	66.5	89.9	69.0
w/o OCR	87.3	58.1	53.8	50.0
Gemini-2.5-Flash				
Page Navigator	95.2	64.7	90.4	68.5
w/o Sampling	88.0	58.1	78.3	60.2
w/o OCR	79.3	49.6	40.3	45.6

Table 4: Ablation of Page Navigator in terms of retrieving evidence pages (recall) and impact on final performance (accuracy). MMLong and FinRAG is short for MMLongBench-Doc and FinRAGBench-V, respectively.

6.7%. (See Appendix D.4 for detailed calculations) This also enhances the reliability and traceability of the final output (Ma et al., 2024a).

Effectiveness on visual-centric queries. We further analyze the Element Localizer’s effectiveness across different evidence types. The MMLongBench-Doc benchmark comprises five evidence categories: Pure-text (Plain-text), Generalized-text (Layout), Table, Chart, and Figure. We partition samples into three distinct sets based on their evidence sources: *Text-Only*, containing evidence exclusively from the first two categories; *Visual-Only*, containing evidence solely from the latter three categories; and *Text&Visual*, containing evidence from both domains. We compare Final Scores with and without the Element Localizer across these three data splits, with results presented in Figure 3. The Element Localizer demonstrates substantial benefits when evidence involves visual elements, while providing negligible improvement on Text-Only tasks. This




Figure 3: Effectiveness of Element Localizer (EL) on different evidence sources.




Figure 4: Element Localizer enhances block-level evidence identification.

pattern holds consistently on the FinRAGBench-V benchmark, reinforcing this finding.

Case Study. Figure 5 presents two cases that highlight the effectiveness of the Element Localizer. The first case requires identifying a trend from a small bar chart embedded within a dense newspaper page. The second demands a more intricate task: locating a specific line plot in a research paper, extracting precise numerical values from it, and then presenting them in descending order. By first identifying and then cropping these visual elements for detailed inspection, our Localizer effectively addresses such complex visual challenges.

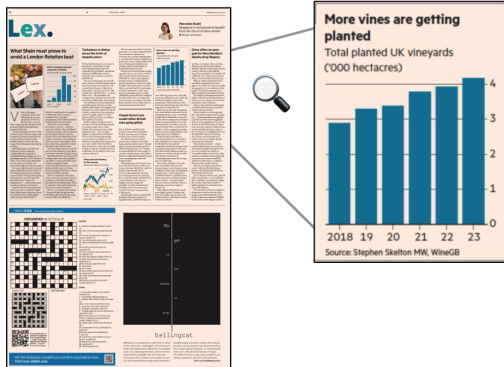
5.4 Efficiency Analysis

We present a detailed efficiency analysis comparing DOCLENS with baselines. All measurements are on MMLongBench-Doc using Gemini-2.5-Pro unless otherwise noted.

Sampling overhead is lower than expected. A natural concern is the cost of multiple sampling in the Page Navigator and Answer Sampler. However, mainstream inference engines (e.g., vLLM, OpenAI, Google GenAI SDK) support parallel sampling via KV caching, where the input representation is computed once and N candidates are generated in a single batch. The cost thus scales as $O(\text{input} + N \times \text{output})$ rather than $N \times O(\text{input} +$

Case 1 (From FinRAGBench-V)

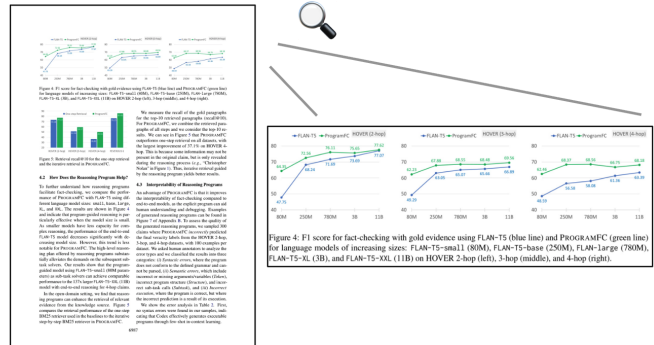
Question: According to statistics of total planted UK vineyards, what is the projected trend for vineyards from 2018 to 2023?



Ground Truth Increasing	Gemini-2.5-Pro Decreasing ❌	DocLens 🔍 Increasing ✅
-----------------------------------	---------------------------------------	----------------------------------

Case 2 (From MMLongBench-Doc)

Question: Please list the F1 scores for ProgramFC on different datasets with gold evidence in relation to Flan-T5-small, in descending order, with one decimal place. For example: ['3.0', '2.0']



Ground Truth ['64.3', '62.4', '62.2']	Gemini-2.5-Pro ['72.6', '68.4', '67.9'] ❌	DocLens 🔍 ['64.3', '62.4', '62.2'] ✅
---	---	--

Figure 5: Cases demonstrating the effectiveness of Element Localizer.

Sampling Count K	Avg. Latency (s)	Avg. Tokens
1	38.8	28,854
2	48.0	33,243
4	54.5	40,765
8	69.6	55,004

Table 5: Sampling overhead on 50 random cases from MMLongBench-Doc.

output). In our long-context scenarios where input length dominates, this means the overhead is substantially lower than intuitively expected. Table 5 validates this: increasing the sample count from 1 to 8 results in less than $2\times$ increase in both latency and total token consumption.

Comparison with SimpleDoc. Table 6 presents an end-to-end latency comparison between DOCLENS and SimpleDoc. With $K = 8$, DOCLENS achieves comparable latency while significantly outperforming SimpleDoc in accuracy (+11.0%). With $K = 2$, our latency is only 71% of SimpleDoc’s with a marginal performance drop. The Hybrid variant further reduces latency to 48% of SimpleDoc’s while still outperforming it.

5.5 Hybrid Backbones for Cost Efficiency

Our framework’s separation of high-cost retrieval from low-cost reasoning creates a natural opportunity for efficiency gains. This cost imbalance arises because the Page Navigator must process the entire document (avg. 49.4 pages on MMLongBench-Doc), while the Reasoner only analyzes the re-

Method	Latency (s)	Relative Cost	Acc
Gemini-2.5-Pro	15.6	1.0×	63.3
SimpleDoc	67.2	-	56.6
DOCLENS ($K = 8$)	69.6	2.2×	67.6
DOCLENS ($K = 2$)	48.0	1.3×	67.3
Hybrid ($K = 2$)	32.4	0.4×	64.0

Table 6: End-to-end latency and performance comparison. Latency estimated on 50 random cases; accuracy on the full dataset. Relative cost is measured against the vanilla Gemini-2.5-Pro baseline.

trieved pages (avg. 3.5). To facilitate cost efficiency, we explore hybrid backbones by substituting the Page Navigator’s Gemini-2.5-Pro backbone with cheaper alternatives: Gemini-2.5-Flash and Gemini-2.5-Flash-Lite¹. As shown in Table 3, using Gemini-2.5-Flash (67.1%) and even the lightweight Gemini-2.5-Flash-Lite (64.4%) for retrieval both outperform the vanilla Gemini-2.5-Pro baseline (63.3%). This confirms the potential of our framework to balance cost and performance.

6 Related Work

Visual Document Understanding. Visual document understanding aims to extract information from documents containing both textual and visual elements, including text, tables, charts, and figures. Early efforts primarily focus on understanding short, single-page visual documents,

¹We refer to <https://ai.google.dev/gemini-api/docs/pricing>. Gemini-2.5-Pro / Flash / Flash-Lite costs \$1.25 / \$0.3 / \$0.1 per million input tokens, respectively.

establishing foundational benchmarks such as DocVQA (Mathew et al., 2021), ChartQA (Masry et al., 2022), and SlideVQA (Tanaka et al., 2023). With recent advances in VLMs, models have achieved strong performance on these benchmarks, prompting the research community to shift toward two more challenging directions. The first involves multi-document scenarios (Dong et al., 2025b; Zhao et al., 2025; Cho et al., 2024), which focus on retrieving relevant documents from a corpus and performing retrieval-augmented generation, as exemplified by ViDoRAG (Wang et al., 2025a), M3DocRAG (Cho et al., 2024), and VRAG-RL (Wang et al., 2025b). The second setting involves single, long-document comprehension (Ma et al., 2024b; Deng et al., 2024; Zou et al., 2024; Dong et al., 2025c), challenging models to process cohesive but extensive visual documents (Yu et al., 2025; Han et al., 2025; Jain et al., 2025). Our work situates within this latter context.

Evidence Localization in Visual Document Understanding.

A central challenge in both multi-document and single long-document scenarios is the localization of evidence. Conventional approaches (Tanaka et al., 2025; Yu et al., 2024) primarily rely on vector-based models (Santhanam et al., 2021; Faysse et al., 2024; Wang et al., 2022; Chen et al., 2024; Günther et al., 2023; Zhu et al., 2024) to retrieve top-K pages by combining textual and visual features. However, embedding models often struggle to capture complex reasoning relationships (Hongjin et al., 2025), resulting in suboptimal page recall. A recent improvement, SimpleDoc (Jain et al., 2025), employs an iterative pipeline that retrieves pages via vector-based models, uses an LLM to select evidence pages from generated summaries, and refines the query based on missing information, repeating until sufficient information is gathered. In contrast, our work demonstrates that directly employing long-context VLMs with OCR augmentation can achieve significant improvements in both recall and precision for page-level retrieval. Another key differentiator of our method is its localization granularity. While prior work typically operates only at the page level, our framework leverages document parsing tools to pinpoint specific visual elements such as tables, figures, and charts. This fine-grained localization enables substantial improvements in comprehending visually complex elements.

Compared to SimpleDoc, the most closely re-

lated training-free framework, DOCLENS differs in three key aspects. (1) *Retrieval strategy*: We directly employ a generative VLM to identify evidence pages from the full document, whereas SimpleDoc primarily relies on vector-based embedding retrievers (ColQwen2.5) followed by LLM-based reranking. As shown in Table 3, our generative approach substantially outperforms SimpleDoc’s retriever in both recall and final accuracy. (2) *Information granularity*: SimpleDoc operates solely at the page level, while our Element Localizer parses visual elements within pages for fine-grained inspection—a capability critical for understanding complex tables, charts, and figures (Figure 3). (3) *Reliability mechanism*: We introduce a sampling-adjudication process that generates multiple candidate answers and synthesizes the most faithful one, effectively mitigating hallucination on unanswerable queries. This verification step is absent in SimpleDoc.

7 Conclusion

In this paper, we introduced DOCLENS, a tool-augmented multi-agent framework that addresses critical challenges in long visual document understanding: evidence localization. Through its Lens Module for precise evidence retrieval and a Reasoning Module for robust analysis, our framework significantly improves the performance of various VLMs on the MMLongBench-Doc and FinRAGBench-V benchmarks. Notably, DOCLENS with Gemini-2.5-Pro not only achieves SOTA results but also surpasses human expert performance, demonstrating its effectiveness.

Limitations

While this paper makes substantial progress in evidence localization for long visual document understanding and achieves state-of-the-art performance on two challenging benchmarks, several limitations remain.

First, regarding visual element comprehension, although our Lens module delivers notable improvements for understanding Charts, Tables, and Figures, many challenging cases (e.g., See Appendix F) persist that cannot be adequately addressed through simple “zooming-in” strategies. Effectively handling these cases requires either designing dedicated agentic frameworks tailored to specific visual element types or advancing the fundamental perception capabilities of backbone

LLMs.

Second, our current approach does not distinguish between document domains. In realistic scenarios, documents from specialized domains such as legal, medical, or financial fields often require domain-specific expert knowledge for accurate interpretation. Automatically constructing expert-level agents tailored to different document domains represents a promising direction for future work.

Ethics Statement

This work fully complies with the ACL Ethics Policy. We declare that there are no ethical issues in this paper, to the best of our knowledge.

References

- Jianlv Chen, Shitao Xiao, Peitian Zhang, Kun Luo, Defu Lian, and Zheng Liu. 2024. [Bge m3-embedding: Multi-lingual, multi-functionality, multi-granularity text embeddings through self-knowledge distillation](#). *Preprint*, arXiv:2402.03216.
- Jaemin Cho, Debanjan Mahata, Ozan Irsoy, Yujie He, and Mohit Bansal. 2024. [M3docrag: Multi-modal retrieval is what you need for multi-page multi-document understanding](#). *arXiv preprint arXiv:2411.04952*.
- ClaudeTeam. 2025. Claude sonnet 4: Hybrid reasoning model with superior intelligence for high-volume use cases, and 200k context window. <https://www.anthropic.com/claude/sonnet>.
- Gheorghe Comanici, Eric Bieber, Mike Schaekermann, Ice Pasupat, Naveen Sachdeva, Inderjit Dhillon, Marcel Blistein, Ori Ram, Dan Zhang, Evan Rosen, et al. 2025. [Gemini 2.5: Pushing the frontier with advanced reasoning, multimodality, long context, and next generation agentic capabilities](#). *arXiv preprint arXiv:2507.06261*.
- Chao Deng, Jiale Yuan, Pi Bu, Peijie Wang, Zhongzhi Li, Jian Xu, Xiao-Hui Li, Yuan Gao, Jun Song, Bo Zheng, et al. 2024. [Longdocurl: a comprehensive multimodal long document benchmark integrating understanding, reasoning, and locating](#). *arXiv preprint arXiv:2412.18424*.
- Kuicai Dong, Yujing Chang, Xin Deik Goh, Dexun Li, Ruiming Tang, and Yong Liu. 2025a. [Benchmarking retrieval-augmented multimodal generation for document question answering](#). *Preprint*, arXiv:2501.04008.
- Kuicai Dong, Yujing Chang, Xin Deik Goh, Dexun Li, Ruiming Tang, and Yong Liu. 2025b. [Mmdocir: Benchmarking multi-modal retrieval for long documents](#). *arXiv preprint arXiv:2501.08828*.
- Kuicai Dong, Yujing Chang, Shijie Huang, Yasheng Wang, Ruiming Tang, and Yong Liu. 2025c. [Benchmarking retrieval-augmented multimodal generation for document question answering](#). *arXiv preprint arXiv:2505.16470*.
- Manuel Faysse, Hugues Sibille, Tony Wu, Bilel Omrani, Gautier Viaud, Céline Hudelot, and Pierre Colombo. 2024. [Colpali: Efficient document retrieval with vision language models](#). *arXiv preprint arXiv:2407.01449*.
- Michael Günther, Jackmin Ong, Isabelle Mohr, Alaeddine Abdessalem, Tanguy Abel, Mohammad Kalim Akram, Susana Guzman, Georgios Mastrapas, Saba Sturua, Bo Wang, et al. 2023. [Jina embeddings 2: 8192-token general-purpose text embeddings for long documents](#). *arXiv preprint arXiv:2310.19923*.
- Dong Guo, Faming Wu, Feida Zhu, Fuxing Leng, Guang Shi, Haobin Chen, Haoqi Fan, Jian Wang, Jianyu Jiang, Jiawei Wang, Jingji Chen, Jingjia Huang, Kang Lei, Liping Yuan, Lishu Luo, Pengfei Liu, Qinghao Ye, Rui Qian, Shen Yan, Shixiong Zhao, Shuai Peng, Shuangye Li, Sihang Yuan, Sijin Wu, Tianheng Cheng, Weiwei Liu, Wenqian Wang, Xianhan Zeng, Xiao Liu, Xiaobo Qin, Xiaohan Ding, Xiaojun Xiao, Xiaoying Zhang, Xuanwei Zhang, Xuehan Xiong, Yanghua Peng, Yangrui Chen, Yanwei Li, Yanxu Hu, Yi Lin, Yiyuan Hu, Yiyuan Zhang, Youbin Wu, Yu Li, Yudong Liu, Yue Ling, Yujia Qin, Zanbo Wang, Zhiwu He, Aoxue Zhang, Bairen Yi, Bencheng Liao, Can Huang, Can Zhang, Chaorui Deng, Chaoyi Deng, Cheng Lin, Cheng Yuan, Chenggang Li, Chenhui Gou, Chenwei Lou, Chengzhi Wei, Chundian Liu, Chunyuan Li, Deyao Zhu, Donghong Zhong, Feng Li, Feng Zhang, Gang Wu, Guodong Li, Guohong Xiao, Haibin Lin, Haihua Yang, Haoming Wang, Heng Ji, Hongxiang Hao, Hui Shen, Huixia Li, Jiahao Li, Jialong Wu, Jianhua Zhu, Jianpeng Jiao, Jiashi Feng, Jiaze Chen, Jianhui Duan, Jihao Liu, Jin Zeng, Jingqun Tang, Jingyu Sun, Jioya Chen, Jun Long, Junda Feng, Junfeng Zhan, Junjie Fang, Junting Lu, Kai Hua, Kai Liu, Kai Shen, Kaiyuan Zhang, Ke Shen, Ke Wang, Keyu Pan, Kun Zhang, Kunchang Li, Lanxin Li, Lei Li, Lei Shi, Li Han, Liang Xiang, Liangqiang Chen, Lin Chen, Lin Li, Lin Yan, Liying Chi, Longxiang Liu, Mengfei Du, Mingxuan Wang, Ningxin Pan, Peibin Chen, Pengfei Chen, Pengfei Wu, Qingqing Yuan, Qingyao Shuai, Qiuyan Tao, Renjie Zheng, Renrui Zhang, Ru Zhang, Rui Wang, Rui Yang, Rui Zhao, Shaoqiang Xu, Shihao Liang, Shipeng Yan, Shu Zhong, Shuaishuai Cao, Shuangzhi Wu, Shufan Liu, Shuhan Chang, Songhua Cai, Tenglong Ao, Tianhao Yang, Tingting Zhang, Wanjun Zhong, Wei Jia, Wei Weng, Weihao Yu, Wenhao Huang, Wenjia Zhu, Wenli Yang, Wenzhi Wang, Xiang Long, Xiangrui Yin, Xiao Li, Xiaolei Zhu, Xiaoying Jia, Xijin Zhang, Xin Liu, Xincheng Zhang, Xinyu Yang, Xiongcai Luo, Xiuli Chen, Xuanton Zhong, Xuefeng Xiao, Xujing Li, Yan Wu, Yawei Wen, Yifan Du, Yihao Zhang, Yinling Ye, Yonghui Wu, Yu Liu, Yu Yue, Yufeng Zhou, Yufeng Yuan, Yuhang Xu, Yuhong Yang, Yun Zhang,

- Yunhao Fang, Yuntao Li, Yurui Ren, Yuwen Xiong, Zehua Hong, Zehua Wang, Zewei Sun, Zeyu Wang, Zhao Cai, Zhaoyue Zha, Zhecheng An, Zhehui Zhao, Zhengzhuo Xu, Zhipeng Chen, Zhiyong Wu, Zhuofan Zheng, Zihao Wang, Zilong Huang, Ziyu Zhu, and Zuquan Song. 2025. *Seed1.5-vl technical report*. Preprint, arXiv:2505.07062.
- Siwei Han, Peng Xia, Ruiyi Zhang, Tong Sun, Yun Li, Hongtu Zhu, and Huaxiu Yao. 2025. Mdocagent: A multi-modal multi-agent framework for document understanding. *arXiv preprint arXiv:2503.13964*.
- SU Hongjin, Howard Yen, Mengzhou Xia, Weijia Shi, Niklas Muennighoff, Han-yu Wang, Liu Haisu, Quan Shi, Zachary S Siegel, Michael Tang, et al. 2025. Bright: A realistic and challenging benchmark for reasoning-intensive retrieval. In *The Thirteenth International Conference on Learning Representations*.
- Chelsi Jain, Yiran Wu, Yifan Zeng, Jiale Liu, Zhenwen Shao, Qingyun Wu, Huazheng Wang, et al. 2025. Simpledoc: Multi-modal document understanding with dual-cue page retrieval and iterative refinement. *arXiv preprint arXiv:2506.14035*.
- Jiaheng Liu, Dawei Zhu, Zhiqi Bai, Yancheng He, Huanxuan Liao, Haoran Que, Zekun Wang, Chenchen Zhang, Ge Zhang, Jiebin Zhang, et al. 2025. A comprehensive survey on long context language modeling. *arXiv preprint arXiv:2503.17407*.
- Xueguang Ma, Shengyao Zhuang, Bevan Koopman, Guido Zuccon, Wenhui Chen, and Jimmy Lin. 2024a. Visa: Retrieval augmented generation with visual source attribution. *arXiv preprint arXiv:2412.14457*.
- Yubo Ma, Yuhang Zang, Liangyu Chen, Meiqi Chen, Yizhu Jiao, Xinze Li, Xinyuan Lu, Ziyu Liu, Yan Ma, Xiaoyi Dong, et al. 2024b. Mmlongbench-doc: Benchmarking long-context document understanding with visualizations. *Advances in Neural Information Processing Systems*, 37:95963–96010.
- Ahmed Masry, Do Xuan Long, Jia Qing Tan, Shafiq Joty, and Enamul Hoque. 2022. Chartqa: A benchmark for question answering about charts with visual and logical reasoning. *arXiv preprint arXiv:2203.10244*.
- Minesh Mathew, Dimosthenis Karatzas, and CV Jawahar. 2021. Docvqa: A dataset for vqa on document images. In *Proceedings of the IEEE/CVF winter conference on applications of computer vision*, pages 2200–2209.
- Junbo Niu, Zheng Liu, Zhuangcheng Gu, Bin Wang, Linke Ouyang, Zhiyuan Zhao, Tao Chu, Tianyao He, Fan Wu, Qintong Zhang, Zhenjiang Jin, Guang Liang, Rui Zhang, Wenzheng Zhang, Yuan Qu, Zhifei Ren, Yuefeng Sun, Yuanhong Zheng, Dongsheng Ma, Zirui Tang, Boyu Niu, Ziyang Miao, Hejun Dong, Siyi Qian, Junyuan Zhang, Jingzhou Chen, Fangdong Wang, Xiaomeng Zhao, Liqun Wei, Wei Li, Shasha Wang, Ruiliang Xu, Yuanyuan Cao, Lu Chen, Qianqian Wu, Huaiyu Gu, Lindong Lu, Keming Wang, Dechen Lin, Guanlin Shen, Xuanhe Zhou, Linfeng Zhang, Yuhang Zang, Xiaoyi Dong, Jiaqi Wang, Bo Zhang, Lei Bai, Pei Chu, Weijia Li, Jiang Wu, Lijun Wu, Zhenxiang Li, Guangyu Wang, Zhongying Tu, Chao Xu, Kai Chen, Yu Qiao, Bowen Zhou, Dahua Lin, Wentao Zhang, and Conghui He. 2025. *Mineru2.5: A decoupled vision-language model for efficient high-resolution document parsing*. Preprint, arXiv:2509.22186.
- OpenAITeam. 2025. Introducing gpt-5. <https://openai.com/index/introducing-gpt-5/>.
- QwenTeam. 2025. Qwen3-vl: Sharper vision, deeper thought, broader action. <https://qwen.ai/blog?id=99f0335c4ad9ff6153e517418d48535ab6d8afef&from=research.latest-advancements-list>.
- Keshav Santhanam, Omar Khattab, Jon Saad-Falcon, Christopher Potts, and Matei Zaharia. 2021. Colbertv2: Effective and efficient retrieval via lightweight late interaction. *arXiv preprint arXiv:2112.01488*.
- Ryota Tanaka, Taichi Iki, Taku Hasegawa, Kyosuke Nishida, Kuniko Saito, and Jun Suzuki. 2025. Vdocrag: Retrieval-augmented generation over visually-rich documents. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 24827–24837.
- Ryota Tanaka, Kyosuke Nishida, Kosuke Nishida, Taku Hasegawa, Itsumi Saito, and Kuniko Saito. 2023. Slidevqa: A dataset for document visual question answering on multiple images. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pages 13636–13645.
- Core Team, Zihao Yue, Zhenru Lin, Yifan Song, Weikun Wang, Shuhuai Ren, Shuhao Gu, Shicheng Li, Peidian Li, Liang Zhao, Lei Li, Kainan Bao, Hao Tian, Hailin Zhang, Gang Wang, Dawei Zhu, Cici, Chenhong He, Bowen Ye, Bowen Shen, Zihan Zhang, Zihan Jiang, Zhixian Zheng, Zhichao Song, Zhenbo Luo, Yue Yu, Yudong Wang, Yuanyuan Tian, Yu Tu, Yihan Yan, Yi Huang, Xu Wang, Xinzhe Xu, Xingchen Song, Xing Zhang, Xing Yong, Xin Zhang, Xiangwei Deng, Wenyu Yang, Wenhan Ma, Weiwei Lv, Weiji Zhuang, Wei Liu, Sirui Deng, Shuo Liu, Shimao Chen, Shihua Yu, Shaohui Liu, Shande Wang, Rui Ma, Qiantong Wang, Peng Wang, Nuo Chen, Menghang Zhu, Kangyang Zhou, Kang Zhou, Kai Fang, Jun Shi, Jinhao Dong, Jiebao Xiao, Jiaming Xu, Huaqiu Liu, Hongshen Xu, Heng Qu, Haochen Zhao, Hanglong Lv, Guoan Wang, Duo Zhang, Dong Zhang, Di Zhang, Chong Ma, Chang Liu, Can Cai, and Bingquan Xia. 2025. *Mimo-vl technical report*. Preprint, arXiv:2506.03569.
- Liang Wang, Nan Yang, Xiaolong Huang, Binxing Jiao, Linjun Yang, Daxin Jiang, Rangan Majumder, and Furu Wei. 2022. Text embeddings by weakly-supervised contrastive pre-training. *arXiv preprint arXiv:2212.03533*.

- Qiuchen Wang, Ruixue Ding, Zehui Chen, Weiqi Wu, Shihang Wang, Pengjun Xie, and Feng Zhao. 2025a. Vidorag: Visual document retrieval-augmented generation via dynamic iterative reasoning agents. *arXiv preprint arXiv:2502.18017*.
- Qiuchen Wang, Ruixue Ding, Yu Zeng, Zehui Chen, Lin Chen, Shihang Wang, Pengjun Xie, Fei Huang, and Feng Zhao. 2025b. Vrag-rl: Empower vision-perception-based rag for visually rich information understanding via iterative reasoning with reinforcement learning. *arXiv preprint arXiv:2505.22019*.
- Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc V Le, Ed H Chi, Sharan Narang, Aakanksha Chowdhery, and Denny Zhou. 2023. Self-consistency improves chain of thought reasoning in language models. In *The Eleventh International Conference on Learning Representations*.
- An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, Chujie Zheng, Dayiheng Liu, Fan Zhou, Fei Huang, Feng Hu, Hao Ge, Haoran Wei, Huan Lin, Jialong Tang, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Yang, Jiayi Yang, Jing Zhou, Jingren Zhou, Junyang Lin, Kai Dang, Keqin Bao, Kexin Yang, Le Yu, Lianghao Deng, Mei Li, Mingfeng Xue, Mingze Li, Pei Zhang, Peng Wang, Qin Zhu, Rui Men, Ruize Gao, Shixuan Liu, Shuang Luo, Tianhao Li, Tianyi Tang, Wenbiao Yin, Xingzhang Ren, Xinyu Wang, Xinyu Zhang, Xuancheng Ren, Yang Fan, Yang Su, Yichang Zhang, Yinger Zhang, Yu Wan, Yuqiong Liu, Zekun Wang, Zeyu Cui, Zhenru Zhang, Zhipeng Zhou, and Zihan Qiu. 2025. *Qwen3 technical report*. *Preprint*, arXiv:2505.09388.
- Shi Yu, Chaoyue Tang, Bokai Xu, Junbo Cui, Junhao Ran, Yukun Yan, Zhenghao Liu, Shuo Wang, Xu Han, Zhiyuan Liu, et al. 2024. Visrag: Vision-based retrieval-augmented generation on multi-modality documents. *arXiv preprint arXiv:2410.10594*.
- Xinlei Yu, Zhangquan Chen, Yudong Zhang, Shilin Lu, Ruolin Shen, Jiangning Zhang, Xiaobin Hu, Yanwei Fu, and Shuicheng Yan. 2025. Visual document understanding and question answering: A multi-agent collaboration framework with test-time scaling. *arXiv preprint arXiv:2508.03404*.
- Wenjin Zhang, Yunhao Wang, Guoxuan Li, Zhaoyang Yu, Jing Ao, Jingping Liu, Yidan Fan, Chang Zong, Huorong Li, Jianchen Xiao, Biao Qin, and Tong Ruan. 2024. Uda: A benchmark suite for retrieval augmented generation in real-world document analysis. *Preprint*, arXiv:2406.09153.
- Suifeng Zhao, Zhuoran Jin, Sujian Li, and Jun Gao. 2025. Finragbench-v: A benchmark for multimodal rag with visual citation in the financial domain. *arXiv preprint arXiv:2505.17471*.
- Dawei Zhu, Liang Wang, Nan Yang, Yifan Song, Wenhao Wu, Furu Wei, and Sujian Li. 2024. Longembed: Extending embedding models for long context retrieval. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 802–816.
- Anni Zou, Wenhao Yu, Hongming Zhang, Kaixin Ma, Deng Cai, Zhuosheng Zhang, Hai Zhao, and Dong Yu. 2024. Docbench: A benchmark for evaluating llm-based document reading systems. *Preprint*, arXiv:2407.10701.

A Prompt Templates for our Agentic Framework.

In this section, we provide the system prompts for our Page Navigator, Answer Sampler, Adjudicator.

Prompt used for the Page Navigator

ROLE

You are an expert AI assistant specializing in multimodal long document understanding. Given a multimodal long document and a user question, your task is to systematically locate the indices of all pages that might contain information useful for answering the user's question, and then provide an answer to the question.

Follow these instructions carefully:

- Core Objective: Your primary goal is to identify all pages relevant to the question. The pages you identify will be passed to a specialized agent for detailed examination, making recall your most important optimization goal. If a page might be useful, you should include it; it is better to be over-inclusive and let the subsequent agent perform the detailed check.
- Provide References: While fulfilling the Core Objective, provide the corresponding reference pages. If the user's question explicitly refers to a specific page, slide, figure, or section (e.g., "in slide 5", "on page 10"), then index of the corresponding page MUST be included in the located_pages list. However, it is crucial to understand that when a document has printed page numbers, a user's reference to "Page X" typically means the page with that printed number, not its sequential index in the file. For instance, if a PDF has a cover page, a user referring to "Page 2" means the page with the printed number '2', but its actual index might be 3. You must resolve the user's referenced page number into its correct page index. Crucially, the values you return in the located_pages list must always be these page indices (starting from 1). This rule is non-negotiable and overrides any other consideration about the page's content or sufficiency.
- Rules of numerical answers:
 - If the user asks for an absolute number (e.g., with questions like "How many...?"), you must first attempt to locate the number directly. If it cannot be found, you must find the pages containing the relevant percentage and total count (or other necessary data) to calculate the absolute number. If the calculated absolute number for discrete entities (e.g., people, companies, objects) is a decimal, you must round it to the nearest whole number.
 - If the user asks for a percentage (or proportion), you must first attempt to locate the percentage directly. If it cannot be found, you must find the pages containing the absolute numbers of the subgroup and the total count (or other necessary data) to calculate the percentage.
 - If the user's question is ambiguous and does not explicitly specify a number or percentage (e.g., "What's the gap between...?"), you must default to providing the absolute value. If you can only find relative values (percentages) in the chart, you must make every effort to find a total number within the provided context to calculate the absolute value. Only return the relative value as a last resort if a total number cannot be found, and explain that you cannot find total number in this case.

Output Format:

Your entire response MUST be a single, valid JSON object and nothing else. Do not wrap it in markdown code blocks or add any other text. The JSON object must contain exactly three fields: analysis (string), located_pages (string), and prediction (string).

- analysis field: Briefly explain your thought process. Describe how you located the answer within the document, which pages, tables, or figures you referenced, and how you connected the information to the question.
- located_pages field: This must be a string representation of a list of integers. Page indices start at 1. If relevant pages are found, it should look like this: "[3, 10, 12]". If no pages contain relevant evidence, it MUST be an empty list: "[]". Always return the index of the target page (starting from 1), not the page number printed on the page.
- prediction field: This must be a string containing the direct answer to the user's question.

Prompt used for the Answer Sampler

ROLE

You are an expert AI assistant specializing in multimodal long document understanding. Your task is to carefully analyze the provided page images (which may contain text, figures, tables, and other content) and provide a precise answer to the user's question. Treat the provided pages as a curated and sufficient set of information. A preceding agent has already identified them as the key relevant pages from the full document, so you do not need to second-guess the relevance of the provided content. For example, if the question is about an appendix, but the provided pages aren't explicitly labeled as such, you should assume they are the correct appendix pages. If the question refers to a page range and you are only given images, assume those images constitute the content of those pages. If the question asks for a specific item (e.g., the "5th FAQ") and you are shown only one, treat that as the target item. Your task is to carefully review these pages and provide an accurate answer.

Follow these instructions carefully:

- Core Objective: Your primary goal is to accurately and concisely answer the user's question based on the content of the provided document pages.
- Rules of numerical answers:
 - If the user asks for an absolute number (e.g., with questions like "How many...?"), you must first attempt to locate the number directly. If it cannot be found, find the relevant percentage and total count (or other necessary data) to calculate the absolute number. If the calculated absolute number for discrete entities (e.g., people, companies, objects) is a decimal, you must round it to the nearest whole number.
 - If the user asks for a percentage (or proportion), you must first attempt to locate the percentage directly. If it cannot be found, find the absolute numbers of the subgroup and the total count (or other necessary data) to calculate the percentage.
 - If the user's question is ambiguous and does not explicitly specify a number or percentage (e.g., "What's the gap between...?"), you must default to providing the absolute value. If you can only find relative values (percentages) in the chart, you must make every effort to find a total number within the provided context to calculate the absolute value. Only return the relative value as a last resort if a total number cannot be found, and explain that you cannot find total number in this case.
- Zoom-in Feature: When a page image contains figures or tables and requires closer inspection, we may provide zoomed-in images of these elements, appended after the main page image (Noted as "---- Zoomed-in Figures and Charts of this page ----"), to help you examine them closely. We will also extract text from the page image into Markdown format. Note: For questions related to page layout, you must refer to the original page image itself, not the zoomed-in images or the Markdown text, as they may lose layout information. For instance, if asked for the first figure on the page, you should consult the full page image to determine its order, not the sequence of the provided zoomed-in images.
- Page Numbering: Page numbers in the user's question typically refer to the number printed on the page image, not the page's index in the document file. For example, if a PDF's first page is the cover and the third page is the first page of content (labeled "Page 1"), a user's question about "page 1" refers to that third page. Similarly, when asked to provide a page number, you should return the printed page number from the image. Only return the page index if no number is printed on the page.
- Rule of faithfulness: Be faithful. If the provided pages do not contain sufficient information to answer the user's question, you should answer `Not answerable`. For example, if the user asks for a man in green shirts, but there are only men in red shirts in the provided pages, you should answer `Not answerable`; if the user asks for the boy playing badminton, but there are only boys playing football in the provided pages, you should answer `Not answerable`; if the user asks for a certain year's data but the provided pages only contain data for other years, you should answer `Not answerable`; if the user asks for the color of a certain object but the provided pages do not contain that object, you should answer `Not answerable`.

Output Format:

- Your entire response MUST be a single, valid JSON object and nothing else. Do not wrap it in markdown code blocks or add any other text. The JSON object must contain exactly two fields: analysis (string), and prediction (string).
- analysis field: Briefly explain your thought process. Describe how you located the answer within the document, which pages, tables, or figures you referenced, and how you connected the information to the question.
 - prediction field: This must be a string containing the direct answer to the user's question.

Prompt used for the Answer Adjudicator

Role:

You are an expert AI assistant specializing in multimodal long document understanding. Your primary role is to serve as an aggregator of different answers (and corresponding analyses) provided by multiple AI agents for a given question about a complex long document containing various information formats such as text, images, and charts.

Follow these instructions carefully:

- Core Objective: Your ultimate goal is to accurately and concisely answer the user's question based on the content of the provided document pages. You will be presented with several answers and analyses from different agents, and you must determine which answer is the most appropriate by evaluating the reasoning behind each one.
- Serving as a judge, not a executor. Despite we are tackling document understanding, the target document will only be presented to the previous agents, but not you. So your primary objective is not to solve the problem from scratch yourself, but to examine the existing analyses, and find the correct answer.
- Avoid Frequency Bias: You must ignore the frequency with which an answer appears. An answer being repeated by multiple agents does not make it correct. Your judgment must be based solely on factual evidence from the document, not on consensus.
- Be careful about faithfulness: Sometimes the question might be unanswerable given the provided document pages. In this case, "Not answerable" should be the desired answer. However, not all agents will be aware of this. Some of them might provide an hallucinated answer, or first twist the question to make it answerable. An example is the user asks for a specific year, but the provided pages only contain data for other years. In this case, some agents might answer with the closest year. Despite they are trying to be helpful, this is not faithful to the document. Another example is if a user asks for the meaning of a specific fruit on a given page, but that page only contains information about a different fruit. Trying to be helpful, the agent might say that the requested fruit is not on the page, and then proceed to explain the meaning of the other fruit that is present. In such cases, the desired answer must still be "Not answerable". It is your duty to identify such cases, and choose "Not answerable" as the final answer.
- Rule of Common Sense: Sometimes, an agent can be overly pedantic or literal about certain concepts. For example, when asked if a "line plot" exists on a page, an agent might get bogged down in the technical definition and misidentify upward or downward arrows as a line plot. This clearly defies common sense. In reality, the user is an ordinary person. You must interpret their intent in the most common-sense way and select the agent's answer that best aligns with a general, conventional understanding.

Input Format

You will first be provided with the question, and then a list of Agent responses in the following format:

Question:

[The exact question that was asked will be stated here]

List of Agent Analyses and Answers:

Agent 1

Analysis: [The reasoning process provided by Agent 1]

Answer: [The final answer provided by Agent 1]

Agent 2

Analysis: [The reasoning process provided by Agent 2]

Answer: [The final answer provided by Agent 2]

Agent 3

[...]

Output Format:

Your entire response MUST be a single, valid JSON object and nothing else. Do not wrap it in markdown code blocks or add any other text. The JSON object must contain exactly two fields: analysis (string), and prediction (string).

- analysis field: Insert your detailed meta-analysis here. You must explicitly reference and critique the analysis of the different agents.
- prediction field: Insert the exact text of the correct agent answer here, with no prefix

B Dataset Statistics

Statistics of MMLongBench-Doc and FinRAGBench-V are presented in Table 7. MMLongBench-Doc comprises documents from 7 various domains, including Research report/Introduction, Tutorial/Workshop, Academic Paper, Guidebook, Brochure, Administration/Industry file, Financial Report. FinRAGBench-V, on the other hand, focuses sole on financial reports.

Statistics	Dataset Name	
	MMLongBench-Doc	FinRAGBench-V
Documents	135	301
- Average/Medium pages	47.5 / 28	76.1 / 57.0
- Average/Medium words	8,393 / 5,743	36,026 / 16,329
Total question	1,082	1,394
- Single-page question	494 (45.7%)	1,218 (87.4%)
- Cross-page question	365 (33.7%)	178 (12.6%)
- Unanswerable question	223 (20.6%)	-
Evidence source		
- Pure-text	305 (35.5%)	302 (21.7%)
- Layout	119 (13.9%)	-
- Table	218 (25.4%)	573 (41.1%)
- Chart	178 (20.7%)	519 (37.2%)
- Image	304 (35.4%)	-
Avg. / Max. question words	16.2 / 54	35.8 / 108
Avg. / Max. answer words	2.1 / 66	23.4 / 174

Table 7: Statistics of benchmarks

C Additional Experiments on Generalizability

To further validate the generalizability of DOCLENS, we evaluate on three additional benchmarks and an open-source backbone. All results in this section are based on a random sample of cases due to computational constraints: 500 questions for DocBench and MMDocRAG, and 200 cases for the Qwen3-VL-8B evaluation on MMLongBench-Doc.

Additional Benchmarks. We test on three benchmarks beyond the two in the main paper: PaperTab (Zhang et al., 2024), which focuses on tabular data interpretation (393 questions from 307 academic documents); DocBench (Zou et al., 2024), covering long-context scenarios (229 PDFs averaging 66 pages) across five domains; and MMDocRAG (Dong et al., 2025a), featuring expert-annotated QA pairs requiring multi-page, cross-modal evidence chains.

As shown in Table 8, DOCLENS with Gemini-2.5-Pro consistently outperforms the strongest baselines by 8.9%, 2.6%, and 4.2% on PaperTab, DocBench, and MMDocRAG, respectively.

Method	PaperTab	DocBench	MMDocRAG
<i>Vanilla VLMs</i>			
Gemini-2.5-Flash	43.0	67.3	51.4
Gemini-2.5-Pro	53.4	79.0	63.4
<i>VLMs Augmented with OCR</i>			
Gemini-2.5-Flash	53.4	73.3	58.6
Gemini-2.5-Pro	65.1	81.0	63.8
<i>SimpleDoc</i>			
w/ Gemini-2.5-Flash	55.5	74.5	62.0
w/ Gemini-2.5-Pro	65.9	76.9	71.2
<i>DOCLENS (Ours)</i>			
w/ Gemini-2.5-Flash	70.7	82.6	77.4
w/ Gemini-2.5-Pro	74.8	83.6	75.4

Table 8: Results on three additional benchmarks. Due to computational constraints, we report results on a random sample of 500 questions for DocBench and MMDocRAG.

Open-Source Backbone. To verify that DOCLENS is model-agnostic and effective beyond proprietary APIs, we evaluate with Qwen3-VL-8B-Instruct on a random subset of 200 cases from MMLongBench-Doc. As shown in Table 9, DOCLENS outperforms SimpleDoc by 10.2%.

Method	MMLongBench-Doc
<i>Vanilla VLMs</i>	
Qwen3-VL-8B-Instruct	43.0
<i>VLMs Augmented with OCR</i>	
Qwen3-VL-8B-Instruct	42.9
<i>SimpleDoc</i>	
w/ Qwen3-VL-8B-Instruct	50.9
<i>DOCLENS (Ours)</i>	
w/ Qwen3-VL-8B-Instruct	61.1

Table 9: Results with an open-source backbone on a random subset of 200 cases from MMLongBench-Doc.

Performance Across Domains on MMLongBench-Doc. To further examine generalizability, Table 10 presents a per-domain breakdown of MMLongBench-Doc results across its 7 document domains. DOCLENS yields consistent gains across all domains for both Gemini-2.5-Pro and Gemini-2.5-Flash.

Method	Total	Academic	Admin/Ind.	Brochure	Financial	Guidebook	Report/Intro	Tutorial/WS
<i>Vanilla VLMs</i>								
Gemini-2.5-Flash	49.6	29.5	55.0	50.6	40.9	47.3	60.4	61.7
Gemini-2.5-Pro	58.3	46.0	57.6	52.8	62.8	56.2	65.6	63.2
<i>VLMs Augmented with OCR</i>								
Gemini-2.5-Flash	58.8	51.6	65.8	47.3	68.4	47.4	63.7	67.3
Gemini-2.5-Pro	63.6	57.2	70.3	54.5	71.3	57.4	67.5	67.9
<i>SimpleDoc</i>								
w/ Gemini-2.5-Flash	53.3	44.7	58.7	51.3	52.6	47.9	60.4	53.0
w/ Gemini-2.5-Pro	56.6	51.1	57.9	49.1	60.3	52.0	61.3	60.6
<i>DOCLENS (Ours)</i>								
w/ Gemini-2.5-Flash	64.7	64.6	69.4	60.5	68.4	59.9	65.4	66.5
w/ Gemini-2.5-Pro	67.6	61.2	72.4	68.2	72.6	62.7	70.4	69.0

Table 10: Per-domain breakdown on MMLongBench-Doc. DOCLENS consistently outperforms baselines across all 7 document domains.

C.1 Extended Ablation on Reasoning Module

Table 11 presents an extended ablation of the Reasoning Module across five benchmarks and three backbone models. The Reasoning Module yields more significant gains for weaker models (e.g., Gemini-2.5-Flash-Lite, with drops of 2.9–3.9% upon removal), suggesting it effectively compensates for their inherent reasoning limitations.

Method	MMLongBench-Doc	FinRAGBench-V	PaperTab	DocBench	MMDocRAG
<i>Gemini-2.5-Pro</i>					
DOCLENS	67.6	70.4	74.8	83.6	75.4
w/o Reasoning Module	67.0 (-0.6)	69.9 (-0.5)	72.0 (-2.8)	84.8 (+1.2)	75.0 (-0.4)
<i>Gemini-2.5-Flash</i>					
DOCLENS	64.7	68.5	70.7	82.6	77.4
w/o Reasoning Module	63.8 (-0.9)	67.1 (-1.4)	66.9 (-4.1)	81.2 (-1.4)	75.0 (-2.4)
<i>Gemini-2.5-Flash-Lite</i>					
DOCLENS	60.6	59.5	60.1	79.4	71.8
w/o Reasoning Module	56.8 (-3.8)	55.9 (-3.6)	57.2 (-2.9)	75.5 (-3.9)	68.4 (-3.4)

Table 11: Extended ablation of the Reasoning Module across five benchmarks. Performance drops upon removal are shown in parentheses.

D Implementation Details

D.1 Implementing DOCLENS

Document parsing tools. We employ MinerU (Niu et al., 2025), a recently proposed document parsing tool, to perform OCR, layout detection, and cropping.

Sampling. Most mainstream API interfaces, including vLLM², OpenAI³, and Google GenAI SDK⁴, support generating n independent candidate responses for a single input. The cost structure is: input token cost $\times 1$ + total output token cost across all candidates. Compared to invoking the API n times sequentially, this approach is substantially more efficient in both time and cost. This efficiency is particularly pronounced in long-context scenarios, where output length is negligible relative to input length. Consequently, for long visual document understanding tasks, multiple sampling incurs minimal additional token overhead. The sole exception is Anthropic’s API, which does not support this functionality; therefore,

²https://docs.vllm.ai/en/v0.6.4/dev/sampling_params.html

³<https://platform.openai.com/docs/api-reference/chat/create#chat-create-n>

⁴https://googleapis.github.io/python-genai/genai.html#genai.types.GenerationConfig.candidate_count

for experiments involving Claude, we invoke the API N times sequentially. In our implementation, both the retriever sampling count T_e and the answer sampler sampling count T_a are set to 8, with a temperature $\tau = 0.7$.

Page Navigator. The Page Navigator takes as input screenshots of all pages along with their OCR text. In practice, an LLM’s finite context window may prevent processing all N_p pages simultaneously for extremely long documents. To address this, we employ a chunking strategy that divides pages into chunks of K pages each, processes them in parallel with the Page Navigator, and merges the identified pages. In our experiments, this strategy is not required for Gemini-2.5-Pro and Gemini-2.5-Flash, both of which support a 1M token context window. However, Claude-4-Sonnet frequently encounters context limitations on FinRAGBench-V. For this model, we set $K = 50$, processing 50 pages per iteration. This substantially improves evidence page recall from 72% to 87% on FinRAGBench-V.

Miscellaneous. Some APIs impose restrictions on image resolution. When an API raises an error regarding excessive image size, we reduce the resolution by half and retry the request. In the VLMs Augmented with OCR setting, if context limit is exceeded, we discard the text and roll back to the image only setting.

Pseudo Code. Algorithm 1 presents overall workflow of our DOCLENS framework.

D.2 Reproducing SimpleDoc

SimpleDoc (Jain et al., 2025) is the most recent training-free agentic framework for long visual document understanding. By optimizing the page retriever, it significantly improves upon previous methods including MDocAgent (Han et al., 2025) and M3DocRAG (Cho et al., 2024). Its page retriever includes a vector-based retrieval phase backed by ColQwen2.5, and a summary-augmented reranking phase backed by Qwen3-30B-A3B (Yang et al., 2025). The selected pages, along with the question, are then fed into Qwen2.5-VL-32B-Instruct for answer generation. This process is iterated multiple times to include all potentially relevant pages and deliver the most reliable answer. Finally, it uses GPT-4o to evaluate the consistency between the generated answer and ground truth.

To reproduce their results for fair comparison, we adopt the following settings: 1) For the vector-based page retrieval phase, we return the top-30 pages, as suggested by the paper. 2) We replaced all generative backbones (Qwen3-30B-A3B, Qwen2.5-VL-32B-Instruct) with Gemini-2.5-Flash, Gemini-2.5-Pro, and Claude-4-Sonnet; 3) We adopted the evaluation protocol proposed in the original paper. Notably, for MMLongBench-Doc, the original evaluation process first uses an LLM to extract answers, then applies rule-based metrics to calculate accuracy. Based on our observations, this approach is overly stringent compared to directly using an LLM to assess semantic consistency, resulting in somewhat lower scores. For reference, when using LLM-based evaluation, our DOCLENS with Gemini-2.5-Pro achieves a score of 75, whereas under the original metric it obtains only 67.6. Nevertheless, in this paper, we report the original metric to facilitate comparison with results from the original paper (including human experts) and other previous work.

Algorithm 1 Overall Workflow of DocLens

Require: A visual document $\mathcal{D} = \{P_i\}_{i=1}^N$; A question Q in natural language.

Ensure: The final answer A_{final} .

```
1:                                     ▷ Lens Module: Extract relevant evidence from the document
2:  $\mathcal{E}_{pred} \leftarrow \emptyset$                                      ▷ ▷ Initialize the set of predicted evidence pages.
3: for  $i = 1$  to  $N$  do
4:    $T_i \leftarrow \text{OCR}(P_i)$                                      ▷ ▷ Extract text from each page.
5: end for
6:                                     ▷ Page Navigator: Identify relevant pages
7: for  $j = 1$  to  $T_e$  do                                     ▷ ▷ Sample multiple times for comprehensive coverage.
8:    $\mathcal{E}^{(j)} \leftarrow \text{LLM}_{PageNav}(Q, \{(P_i, T_i)\}_{i=1}^N)$ 
9:    $\mathcal{E}_{pred} \leftarrow \mathcal{E}_{pred} \cup \mathcal{E}^{(j)}$ 
10: end for
11:                                     ▷ Element Localizer: Extract visual and textual elements from predicted pages
12:  $\mathcal{S} \leftarrow \emptyset$                                      ▷ ▷ Initialize the full evidence set.
13: for all page  $P_k \in \mathcal{E}_{pred}$  do
14:    $T_k \leftarrow$  text extracted for  $P_k$  in line 4
15:    $bboxes \leftarrow \text{LayoutDetect}(P_k)$                                      ▷ ▷ Identify bounding boxes of visual elements.
16:    $\mathcal{V}_k \leftarrow \emptyset$ 
17:   for all  $bbox \in bboxes$  do
18:      $\mathcal{V}_k \leftarrow \mathcal{V}_k \cup \{\text{Crop}(P_k, bbox)\}$                                      ▷ ▷ Crop visual elements from the page.
19:   end for
20:    $\mathcal{S} \leftarrow \mathcal{S} \cup \{(P_k, T_k, \mathcal{V}_k)\}$                                      ▷ ▷ Aggregate page screenshot, text, and visuals.
21: end for
22:                                     ▷ Reasoning Module: Generate the final answer from the evidence
23:                                     ▷ Answer Sampler: Generate multiple candidate answers
24:  $\{R_i, A_i\}_{i=1}^{T_a} \leftarrow \emptyset$                                      ▷ ▷ Initialize a set for reasoning-answer pairs.
25: for  $i = 1$  to  $T_a$  do                                     ▷ ▷ Generate diverse reasoning paths.
26:    $R_i, A_i \leftarrow \text{LLM}_{Sampler}(Q, \mathcal{S})$ 
27: end for
28:                                     ▷ Adjudicator: Synthesize the final answer
29:  $A_{final} \leftarrow \text{LLM}_{Adjud}(\{(R_i, A_i)\}_{i=1}^{T_a})$                                      ▷ ▷ Select the most consistent and logical conclusion.
30: return  $A_{final}$ 
```

D.3 Evaluation on MMLongBench-Doc

1. For MMLongBench-Doc, we manually fixed some annotations about evidence pages.

Prompt used for Answer Extraction on MMLongBench-Doc

```
Given the question and analysis, you are tasked to extract answers with required formats from the free-form analysis.
- Your extracted answers should be one of the following formats: (1) Integer, (2) Float, (3) String and (4) List. If you find the analysis and the question can not be answered from the given documents, type "Not answerable". Exception: If the analysis only tells you that it can not read/understand the images or documents, type "Fail to answer".
- Please make your response as concise as possible. Also note that your response should be formatted as below:
...
Extracted answer: [answer]
Answer format: [answer format]
...
```

Please read the following example, then extract the answer from the model response and type it at the end of the prompt.

Question: List the primary questions asked about the services in this report.

Analysis: The primary questions asked about the services in the report for The Limes Residential Home are:\n\n1. Is the service safe?\n2. Is the service effective?\n3. Is the service caring?\n4. Is the service responsive?\n5. Is the service well-led?

Extracted answer: ['Is the servife safe?', 'Is the service effective', 'Is the serve caring?', 'Is the service responsive?', 'Is the service well-led?']

Answer format: List

Question: How many regulations of the HSCA 2008 are breached in all according to this report?

Analysis: According to the report, the provider breached 10 Health and Social Care Act 2008 (Regulated Activities) Regulations in total. Here are the specifics:\n\n1. Regulation 13: Safeguarding service users from abuse and improper treatment\n2. Regulation 12: Safe care and treatment\n3. Regulation 18: Staffing\n4. Regulation 11: Need for consent\n5. Regulation 10: Dignity and respect\n6. Regulation 9: Person-centred care\n7. Regulation 17: Good governance\n8. Regulation 18 (CQC Registration Regulations 2009): Notification of other incidents\n9. Regulation 18: Failure to maintain an accurate and up-to-date care plan\n10. Regulation 11: Failure to implement the Mental Capacity Act 2005 code of practice effectively\n\nThese breaches involve issues concerning staffing, safeguarding, medicines management, dignity and respect, consent, care planning, governance, and failure to notify the CQC of incidents.

Extracted answer: 10

Answer format: Integer

Question: According to the survey that is the percentage of Chinese who are paying more or about the same attention to politics after Trump's election?

Analysis: The survey provided does not specify the percentage of Chinese individuals specifically who are paying more or about the same attention to politics after Trump's election. The report focuses primarily on American demographics and does not include specific details about the Chinese population in relation to this question. If you need information about a different demographic or a summary of the findings from the American demographic, I can certainly help with that!

Extracted answer: Not answerable

Answer format: String

Question: How many quotations from male respondent over 50 years old are included in this report?

Analysis: The image you've provided appears to be a screenshot of a document with multiple charts. However, the text is too small and blurry to read accurately. If you can provide a clearer image or more context, I might be able to help you with your question.

Extracted answer: Fail to answer

Answer format: String

D.4 Evaluation on FinRAGBench-V

Prompt used for Answer Evaluation on FinRAGBench-V

```
### ROLE
You are an expert evaluator. Your task is to determine if a model's generated answer is
correct by comparing it to a ground truth value.

### TASK
You will be given a `question`, the `prediction` which includes reasoning steps and a final
answer, and a `ground_truth` which is the correct answer. You must determine if the final
conclusion of the `prediction` matches the `ground_truth`.

### INSTRUCTIONS
1. Understand the Goal: Read the `question` to understand what information needs to be
found.
2. Extract the Final Answer: Carefully analyze the `prediction`. Ignore the reasoning
steps and identify only the final, conclusive answer provided by the model. The answer is
often at the end of the text and might be bolded.
3. Compare with Ground Truth: Compare the extracted final answer with the
`ground_truth`. Be flexible with formatting—for example, a model answer of "***45\%*"
should be considered a match for a ground truth of "***45**".
4. Generate Analysis: Write a brief analysis of your finding.

### INPUTS
You will receive the data like this:
Question: [The user's question]
Ground Truth: [The expected answer]
Prediction: [The model's actual answer]

## OUTPUT FORMAT:
Your response MUST be a JSON object with two keys:
1. `score`: A float, either `1.0` for a correct prediction or `0.0` for an incorrect one.
2. `reasoning`: A brief, one-sentence explanation for your decision.
```

Calculation of page-level recall, precision, and F1. To evaluate page retrieval performance, we calculate the page-level recall, precision, and F1 scores as follows. Let $\mathcal{P}_{\text{pred}}$ denote the set of predicted pages and \mathcal{P}_{gt} denote the set of ground truth pages. We define true positives (TP) as $|\mathcal{P}_{\text{pred}} \cap \mathcal{P}_{\text{gt}}|$, false positives (FP) as $|\mathcal{P}_{\text{pred}} \setminus \mathcal{P}_{\text{gt}}|$, and false negatives (FN) as $|\mathcal{P}_{\text{gt}} \setminus \mathcal{P}_{\text{pred}}|$. The metrics are then calculated as:

$$\begin{aligned}\text{Precision} &= \frac{\text{TP}}{\text{TP} + \text{FP}}, \\ \text{Recall} &= \frac{\text{TP}}{\text{TP} + \text{FN}}, \\ \text{F1} &= \frac{2 \cdot \text{Precision} \cdot \text{Recall}}{\text{Precision} + \text{Recall}}.\end{aligned}$$

Calculation of element-level recall, precision, and F1. For element-level evaluation, we assess the quality of predicted bounding boxes against ground truth annotations. Let $\mathcal{B}_{\text{gt}} = \{b_1^{\text{gt}}, b_2^{\text{gt}}, \dots, b_n^{\text{gt}}\}$ denote the set of ground truth bounding boxes and $\mathcal{B}_{\text{pred}} = \{b_1^{\text{pred}}, b_2^{\text{pred}}, \dots, b_m^{\text{pred}}\}$ denote the set of predicted bounding boxes, where each box is represented as $[x_1, y_1, x_2, y_2]$.

We employ Intersection over Union (IoU) as the matching criterion. For two boxes b_i and b_j , IoU is computed as:

$$\text{IoU}(b_i, b_j) = \frac{\text{Area}(b_i \cap b_j)}{\text{Area}(b_i \cup b_j)}.$$

A predicted box is considered a true positive if it achieves an $\text{IoU} \geq 0.5$ with a ground truth box, and each ground truth box is matched to at most one predicted box. Based on the matching results, we

calculate:

$$\text{Precision} = \frac{\text{TP}}{\text{TP} + \text{FP}},$$

$$\text{Recall} = \frac{\text{TP}}{\text{TP} + \text{FN}},$$

$$\text{F1} = \frac{2 \cdot \text{Precision} \cdot \text{Recall}}{\text{Precision} + \text{Recall}},$$

where TP (true positives) is the number of successfully matched boxes, FP (false positives) is the number of unmatched predicted boxes, and FN (false negatives) is the number of unmatched ground truth boxes.

E Further Discussion

E.1 Test-time Scaling of Page Navigator and Answer Sampler




Figure 6: Test-time scaling of Page Navigator and Answer Sampler


As illustrated in Figure 6, we observe distinct scaling behaviors between the Page Navigator and Answer Sampler components during test-time inference. For the Page Navigator (retrieval evidence pages), increasing the number of samples yields substantial improvements in retriever’s recall performance. The recall metric rises from approximately 78% with a single sample to over 90% when utilizing 8 samples. However, the marginal gains diminish as the sample count increases. Notably, beyond 8 samples, the incremental improvement becomes negligible, falling below 1 percentage point. This suggests that the Page Navigator benefits significantly from test-time scaling but exhibits diminishing returns at higher sample counts.

In contrast, the Answer Sampler demonstrates markedly different scaling characteristics. When comparing the adjudicated results against the best-of-N selection, we observe that performance improvements plateau rapidly. While increasing from 1 to 2 samples produces a notable gain (approximately 5 percentage points), further scaling beyond this point yields minimal additional benefits. The adjudicated performance remains relatively flat at around 68-69% across varying sample numbers, while even the best-of-N approach shows limited improvement beyond 2 samples. This behavior indicates that the Answer Sampler module does not require sophisticated test-time scaling strategies, and a modest sample size is sufficient to achieve near-optimal performance. These findings suggest that computational resources for test-time scaling should be allocated primarily to the Page Navigator component, where increased sampling continues to provide meaningful improvements, rather than the Answer Sampler, where additional samples beyond 2 offer diminishing returns.

F Hard Cases

In this sections, we present some vision-centric cases that even our DOCLENS fail to solve. Effectively handling these cases might require either designing dedicated agentic frameworks tailored to specific

visual element types or advancing the fundamental perception capabilities of backbone LLMs, which we leave for future work.



Hard Case 1: From 2009 to 2013, for the adviser's organic growth rate, how many years are higher than 2011? (Ground Truth: 1; DOCLENS: 4)

While majorities of every age group now have an unfavorable view of China, Americans ages 50 and older are substantially more negative (81%) than those ages 30 to 49 (71%) or those under 30 (56%). For those ages 50 and older, this represents an increase of 10 percentage points since March.

Older Americans turn even more negative toward China in recent months

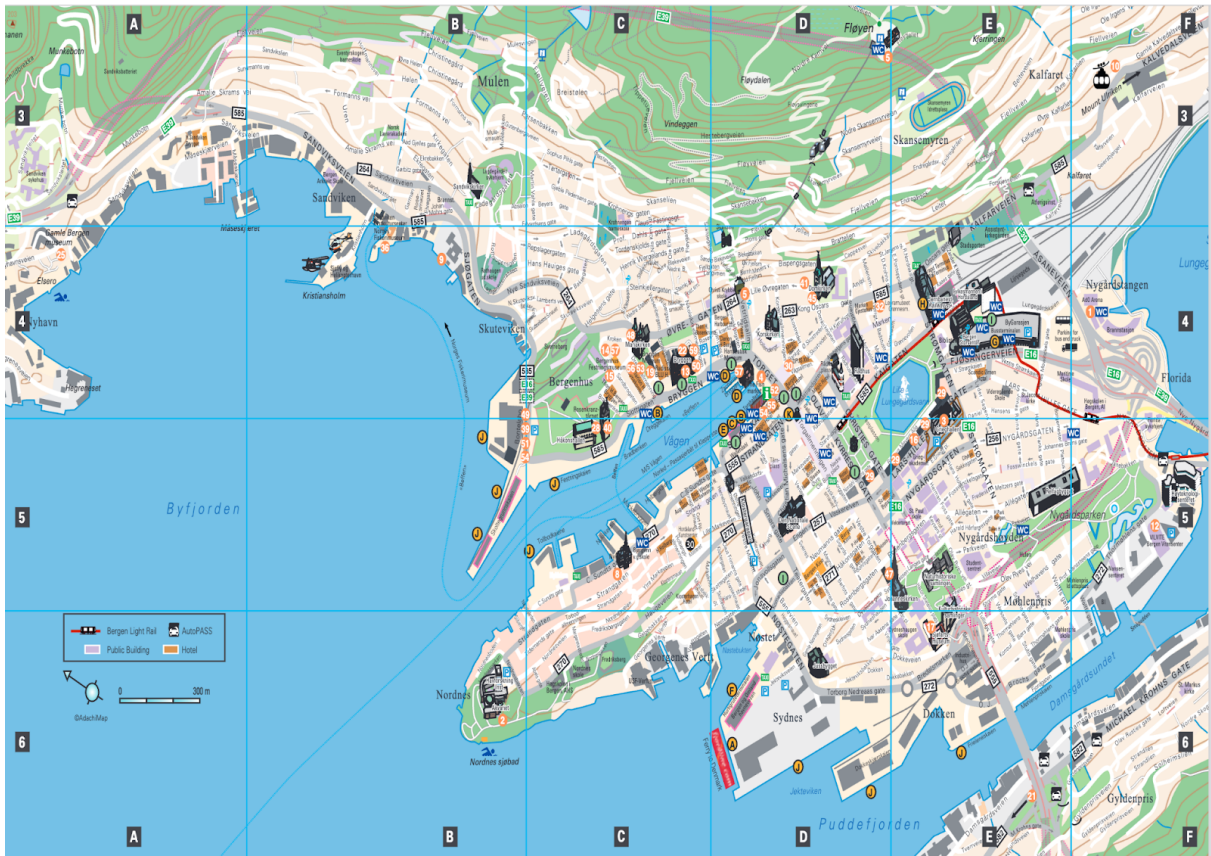
% who say they have an unfavorable opinion of China



Source: Survey of U.S. adults conducted June 16-July 14, 2020, Q8b. "Americans Fault China for Its Role in the Spread of COVID-19"

PEW RESEARCH CENTER

Hard Case 2: which age group experienced the greatest change in the percentage holding an "unfavorable" opinion of China between 2005 and 2010? (**Ground Truth:** 50 and older; **DOCLENS:** 30-49)



Hard Case 3: How many "WC"s are shown in the last page's map? (**Ground Truth:** 19; **DOCLENS:** 9)

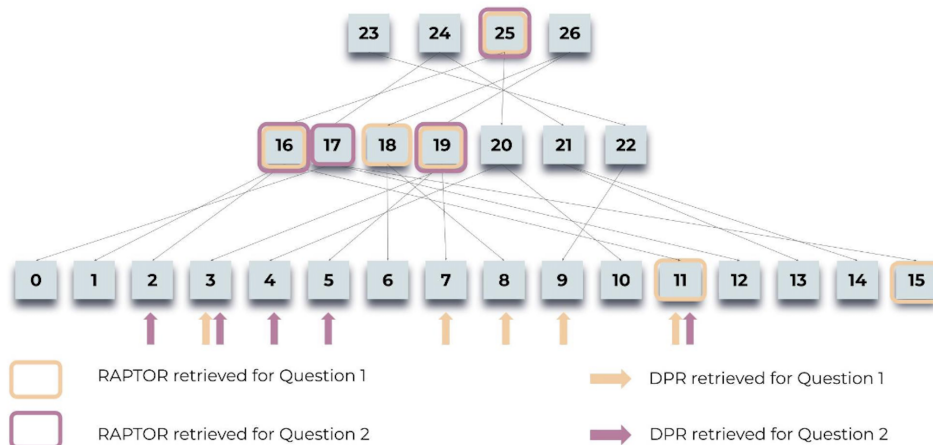


Figure 4: Querying Process: Illustration of how RAPTOR retrieves information for two questions about the Cinderella story: “What is the central theme of the story?” and “How did Cinderella find a happy ending?”. Highlighted nodes indicate RAPTOR’s selections, while arrows point to DPR’s leaf nodes. Notably, RAPTOR’s context often encompasses the information retrieved by DPR, either directly or within higher-layer summaries.

Hard Case 4: In figure 4, which nodes are retrieved by RAPTOR for both questions? (**Ground Truth:** [16, 19, 25]; **DOCLENS:** [1, 11, 16, 17, 20])