

DocR1: Evidence Page-Guided GRPO for Multi-Page Document Understanding

Junyu Xiong^{1*}, Yonghui Wang^{1*}, Weichao Zhao¹, Chenyu Liu²,
Bing Yin^{2†}, Wengang Zhou^{1†}, Houqiang Li¹

¹ University of Science and Technology of China

² iFLYTEK Research

{xiongjyu, wyh1998, saruka}@mail.ustc.edu.cn, {cylu7, bingyin}@iflytek.com, {zhwg, lihq}@ustc.edu.cn

Abstract

Understanding multi-page documents poses a significant challenge for multimodal large language models (MLLMs), as it requires fine-grained visual comprehension and multi-hop reasoning across pages. While prior work has explored reinforcement learning (RL) for enhancing advanced reasoning in MLLMs, its application to multi-page document understanding remains underexplored. In this paper, we introduce **DocR1**, an MLLM trained with a novel RL framework, **Evidence Page-Guided GRPO (EviGRPO)**. EviGRPO incorporates an evidence-aware reward mechanism that promotes a coarse-to-fine reasoning strategy, guiding the model to first retrieve relevant pages before generating answers. This training paradigm enables us to build high-quality models with limited supervision. To support this, we design a two-stage annotation pipeline and a curriculum learning strategy, based on which we construct two datasets: EviBench, a high-quality training set with 4.8k examples, and ArxivFullQA, an evaluation benchmark with 8.6k QA pairs based on scientific papers. Extensive experiments across a wide range of benchmarks demonstrate that DocR1 achieves state-of-the-art performance on multi-page tasks, while consistently maintaining strong results on single-page benchmarks.

Introduction

Documents are ubiquitous in everyday life, encompassing scanned forms, tables, charts, and PDFs. Consequently, document understanding is a critical capability for multimodal large language models (MLLMs). While recent MLLMs have achieved strong performance on single-page document tasks (Hu et al. 2024a; Zhou et al. 2024; Liu et al. 2024; Wang et al. 2023; Zhao et al. 2024; Feng et al. 2024), these tasks often cover a subset of real-world applications. In practice, many tasks involve multi-page documents—such as scientific papers, contracts, and reports—which require not only fine-grained visual understanding but also the ability to retrieve relevant pages and reason across dispersed content.

Recent advances, including OpenAI’s o1 (Jaech et al. 2024) and DeepSeek-R1 (Guo et al. 2025), highlight the growing importance of advanced reasoning in LLMs. Parallel work has begun to explore reinforcement learning (RL)

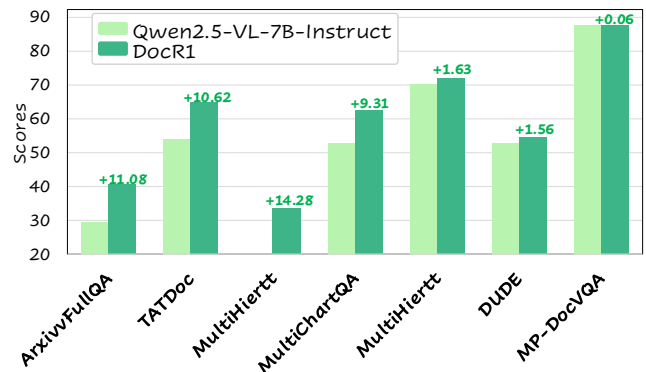


Figure 1: Our DocR1 has significant improvements on various benchmarks of multi-page document understanding.

in MLLMs. For example, VisualRFT (Liu et al. 2025) introduced verifiable reward functions for perception tasks and applied GRPO (Shao et al. 2024) to enhance object detection and visual classification. Similarly, Vision-R1 (Huang et al. 2025) integrated GRPO with a PTST training strategy to improve mathematical reasoning capabilities. However, these methods are restricted to single-image inputs and focus on either perception or symbolic reasoning. For multi-image tasks, methods like Video-R1 (Feng et al. 2025) and VideoChat-R1 (Li et al. 2025) have applied GRPO in the video domain, incorporating customized reward functions to capture spatiotemporal dependencies. Nevertheless, applying RL to multi-page document understanding—where effective evidence retrieval and multi-hop reasoning are both critical—remains largely unexplored.

To bridge this gap, we propose **DocR1**, a novel method equipped with the **Evidence Page-Guided GRPO (EviGRPO)** RL framework, specifically designed to enhance multi-page reasoning in MLLMs. EviGRPO extends the GRPO paradigm by introducing an evidence-aware reward mechanism that promotes a human-like, coarse-to-fine reasoning process: first forming a global understanding of the document, then retrieving the relevant pages, and finally reasoning over them to derive the answer. Specifically, we define three verifiable rewards to guide model optimization: format consistency, evidence page accuracy, and answer ac-

*These authors contributed equally.

†Corresponding author.

curacy. To support this process, we design a rigorous two-stage annotation pipeline comprising a generation step and a verification step. In the generation step, an MLLM is prompted to produce initial annotations. In the verification step, the same MLLM is prompted to validate the generated annotations, ensuring higher data quality through self-checking. Using this pipeline, we construct two datasets: EviBench, a minimally supervised yet high-quality training set containing 1.3k single-page and 3.5k multi-page samples; and ArxivFullQA, a new benchmark with 8.6k QA samples designed to evaluate document-level reasoning over full scientific papers. In addition, we introduce a two-stage curriculum training strategy to enable more effective learning under limited supervision. The model is first trained on single-page data to internalize the desired output format and reasoning style, and then trained on multi-page data to develop its multi-page reasoning capabilities.

Our proposed DocR1 is an MLLM tailored for complex document-level reasoning. Rather than producing only final answers, DocR1 explicitly outputs intermediate reasoning traces and evidence localization, thereby improving both interpretability and reliability. As shown in Figure 2, the model follows a structured decision-making path aligned with human reading behavior. Comprehensive evaluations across both multi-page and single-page benchmarks validate the effectiveness of our method. Specifically, DocR1 achieves state-of-the-art performance on multi-page benchmarks, attaining an average score of 59.36, outperforming the baseline by 6.93 points.

Our contributions can be summarized as follows:

- We present DocR1, an MLLM designed to generate structured outputs comprising the models’s thought process, selected evidence pages, and final answers.
- We propose EviGRPO, an RL framework specifically designed for multi-page document understanding in MLLMs. EviGRPO adopts a coarse-to-fine reasoning strategy, where the model first identifies the relevant evidence pages before generating an answer.
- We develop a two-stage annotation pipeline and a curriculum training strategy that together enable high-quality data labeling and efficient model training under limited supervision.

Related Work

Document Understanding with MLLMs

MLLMs have recently advanced document understanding by enabling unified modeling of textual and visual information, without relying on traditional OCR engines. Several recent works have explored MLLM-based document understanding. UReader (Ye et al. 2023b) pioneered OCR-free, end-to-end instruction tuning across multiple tasks. The mPLUG-DocOwl series (Ye et al. 2023a; Hu et al. 2024a,b) improved structure awareness and computational efficiency through structural modeling and visual token compression. TextMonkey (Liu et al. 2024) introduced shifted window attention to preserve semantic continuity across image patches. DOGE (Zhou et al. 2024) proposed a document-oriented grounding framework with high-quality training

data, while Doc-CoB (Mo et al. 2025) enabled dynamic region selection for focused, step-by-step reasoning. However, most methods are still limited to single-page documents. Our work targets more complex multi-page document tasks in real scenarios.

Reinforcement Learning for MLLMs

RL has shown promise in enhancing the reasoning capabilities of LLMs. While early approaches relied on pretrained reward models, recent work such as DeepSeek-R1 (Guo et al. 2025) demonstrates that simple, rule-based rewards can provide scalable and verifiable supervision. Building on this insight, researchers have extended RL to multimodal large language models (MLLMs). VisualRFT (Liu et al. 2025) introduced visually grounded, verifiable reward functions for fine-grained classification and few-shot detection using GRPO. Vision-R1 (Huang et al. 2025) further leveraged GRPO with a PTST strategy to improve mathematical reasoning on a curated multimodal dataset. In the video domain, Video-R1 (Feng et al. 2025) proposed T-GRPO to explicitly incorporate temporal cues during training, marking the first systematic application of GRPO to video-based MLLMs. VideoChat-R1 (Li et al. 2025) also applied GRPO to enhance spatiotemporal reasoning, achieving strong performance on temporal grounding and tracking tasks. While these studies focus on perception and video reasoning, RL has not yet been explored for multi-page document understanding. Our work aims to fill this gap.

Method

EviGRPO

While the GRPO algorithm (Shao et al. 2024) has demonstrated promise in enhancing reasoning capabilities, its direct application to multi-page document understanding remains suboptimal. To address this gap, we introduce **Evidence Page-Guided GRPO (EviGRPO)**, a variant of GRPO tailored specifically for multi-page document tasks, as illustrated in Figure 2.

When engaging in multi-page reading comprehension, humans typically begin by identifying the pages likely to contain the answer, and then focus on locating the specific regions that correspond to the question and answer within those pages. Inspired by this “coarse-to-fine” reading strategy, EviGRPO mimics the human approach by first selecting a small set of potentially relevant pages at a coarse level, followed by fine-grained reasoning over the selected content. This hierarchical reading paradigm facilitates a more efficient and accurate understanding of multi-page documents.

EviGRPO builds upon the original GRPO structure, which includes format and accuracy rewards, by introducing an additional evidence-aware reward that encourages grounding answers on relevant document pages. For each training sample, consisting of a question q and N document images, EviGRPO first generate G candidate responses $\mathcal{O} = \{o_1, o_2, \dots, o_G\}$ using the current policy. For each candidate response o_i , the verifiable total reward r_i is defined as the sum of three components:

$$r_i = r_i^{\text{format}} + r_i^{\text{acc}} + r_i^{\text{evi}}. \quad (1)$$

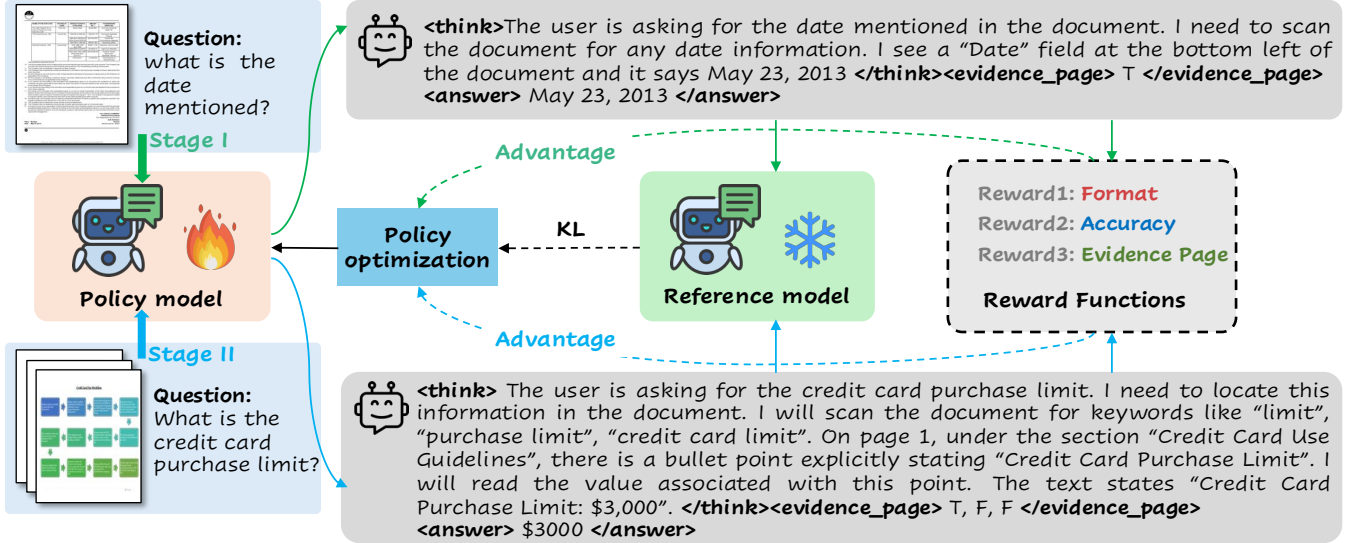


Figure 2: Our proposed EviGRPO training framework adopts a two-stage strategy to progressively enhance the model’s multi-page reasoning capabilities.

System: You will be given one or more images along with a question. Your task is to understand the visual content and answer the question. First, think carefully about the question and present your reasoning in `<think>` and `</think>`. Next, identify how many pages (images) are provided, and for each page, determine whether it contains relevant evidence to answer the question. List your judgments in `<evidence_page>` and `</evidence_page>` using a comma-separated sequence of T (True) or F (False), one for each page, in order (e.g., T, F, T, F). Finally, provide your answer in `<answer>` and `</answer>`. The answer should be one or more words or phrases. **User:** prompt. **Assistant:**

Figure 3: The page selection format of EviGRPO.

where $r_i^{\text{format}} = 1$ if the model’s output strictly adheres to the formatting rules specified in Figure 3, and $r_i^{\text{format}} = 0$ otherwise. The accuracy reward r_i^{acc} is defined as the ANLS score between the model-generated answer and the ground-truth. The evidence-page reward r_i^{evi} measures whether the model correctly identifies the supporting pages. Let N'_i denote the number of predicted evidence pages, and let P_i and G_i represent the sets of predicted and ground-truth evidence pages, respectively. We define r_i^{evi} as the F1-style overlap between the predicted and ground-truth evidence page sets. Specifically, it is computed as:

$$r_i^{\text{evi}} = \begin{cases} \frac{2 \cdot |P_i \cap G_i|}{|P_i| + |G_i|}, & \text{if } |P_i| + |G_i| > 0 \text{ and } N = N'_i \\ 0, & \text{otherwise} \end{cases} \quad (2)$$

We adopt the F1 score rather than accuracy to mitigate reward hacking. In multi-page documents, where evidence

pages are typically sparse, predicting all pages as irrelevant can lead to artificially high accuracy. The F1-based reward addresses this issue by balancing precision and recall. To further encourage fine-grained reasoning, the model is required to assess each image individually, labeling each as either relevant (T) or irrelevant (F). To enforce this behavior, we set the reward to zero if the number of predicted judgments N'_i does not match the total number of input images N . Next, following the GRPO framework, EviGRPO computes the mean and standard deviation of the total rewards across the response set \mathcal{O} , and normalizes each individual reward to obtain its corresponding advantage value A_i :

$$A_i = \frac{r_i - \text{mean}(\{r_i\}_{i=1}^G)}{\text{std}(\{r_i\}_{i=1}^G)}. \quad (3)$$

The final policy optimization objective is to maximize the expected return. To regularize the update, a KL divergence term $D_{\text{KL}}(\cdot \| \cdot)$ is introduced to constrain the optimized policy π_θ from diverging excessively from the reference policy π_{ref} . Moreover, a clipping term is applied to avoid overly large gradient steps, ensuring stable training. The resulting objective is formulated as:

$$\max_{\pi_\theta} \mathbb{E}_{\mathcal{O} \sim \pi_{\theta_{\text{old}}}(p)} \left[\min \left(\text{clip} \left(\frac{\pi_\theta(o_i)}{\pi_{\theta_{\text{old}}}(o_i)}, 1 - \epsilon, 1 + \epsilon \right) A_i, \frac{\pi_\theta(o_i)}{\pi_{\theta_{\text{old}}}(o_i)} A_i \right) - \beta D_{\text{KL}}(\pi_\theta \| \pi_{\text{ref}}) \right] \quad (4)$$

Data Construction

High-quality training data is critical for improving the multi-page document understanding capabilities of MLLMs. However, most existing open-source datasets are not directly compatible with our EviGRPO framework due to mismatches in input-output formats and reward structures. To

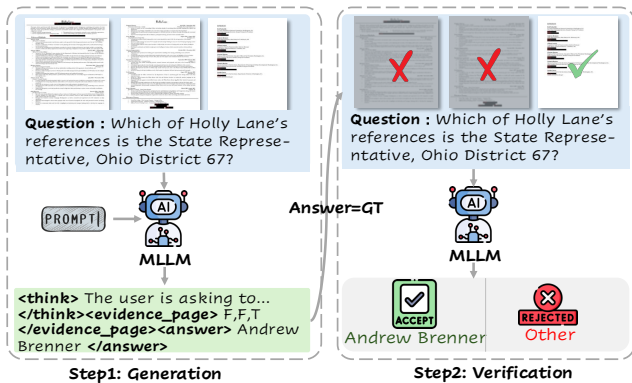


Figure 4: The two-stage data annotation pipeline consists of a data generation process and an annotation verification process, designed to ensure the quality of the annotations.

address this, we design a new data construction pipeline and introduce EviBench, the training dataset used in our study. To further ensure comprehensive evaluation across diverse multi-page document scenarios, we construct a new evaluation dataset, ArxivFullQA, using the same annotation pipeline.

Data Annotation Pipeline As illustrated in Figure 4, we develop a rigorous two-stage data annotation pipeline. In the first stage, given an input and a task-specific prompt, we employ the Gemini 2.5 Flash model (Comanici et al. 2025) to generate a target output. The sample advances to the next stage only if its predicted answer is consistent with the ground truth. In the second stage, the same MLLM is prompted again with the annotated content in order to verify its accuracy. The annotation is retained only if the model’s output once again aligns with the ground truth under this controlled setting.

EviBench for Training As shown in Table 1, based on the proposed annotation pipeline, we presents the statistics of our annotated dataset. The dataset includes both single-page and multi-page document samples. Specifically, for single-page documents, we curate 13 widely-used datasets, including DocVQA (Mathew, Karatzas, and Jawahar 2021), InfographicsVQA (Mathew et al. 2022), ChartQA (Masry et al. 2022), DeepForm (Svetlichnaya 2020), DVQA (Kafle et al. 2018), FigureQA (Kahou et al. 2017), KleisterCharity (Stanislawek et al. 2021), OCRVQA (Mishra et al. 2019), TabFact (Chen et al. 2019), TextCaps (Sidorov et al. 2020), TextVQA (Singh et al. 2019), VisualMRC (Tanaka, Nishida, and Yoshida 2021), and WikiTableQuestions (Pasupat and Liang 2015). These datasets cover a diverse range of document types. For instance, DocVQA and DVQA contain scanned pages and visual charts, while DeepForm, TabFact, and WikiTableQuestions focus on structured tabular formats. This diversity enables the model to learn from various layouts and visual structures commonly encountered in real-world documents. For the multi-page document domain, we integrate three widely used datasets: DUDE (Van Landeghem et al. 2023), MP-DocVQA (Tito, Karatzas, and Val-

Category	Dataset	Images	#Samples
Single	DocVQA	1	100
	InfographicVQA	1	100
	ChartQA	1	100
	DeepForm	1	100
	DVQA	1	100
	FigureQA	1	100
	KleisterCharity	1	100
	OCRVQA	1	100
	TabFact	1	100
	TextCaps	1	100
	TextVQA	1	100
	VisualMRC	1	100
	WikiTableQuestions	1	100
Total			1300
Multi	DUDE	1–21	1000
	MP-DocVQA	1–36	500
	TATDoc	1–3	500
	SlideVQ	20	500
	MultihierTT	3–7	500
	ArxivFullQA _{train}	1–29	500
Total			3500

Table 1: Statistics of each dataset used for training.

veny 2023), and TATDoc (Zhu et al. 2022), which are specifically designed for document-level reasoning across multi-page documents. To further enhance visual diversity, we also include SlideVQA (Tanaka et al. 2023), a dataset based on multi-slide presentations, and MultihierTT (Zhao et al. 2022), which focuses on complex hierarchical visual structures in multi-page charts. In addition, academic paper reading is a crucial task in the multi-page document domain. To improve the model’s capability in this area, we annotate a training subset, denoted as ArxivFullQA_{train}, from the DocMatrix dataset (Laurençon et al. 2024), following the same annotation details described in the next section.

Specifically, we adopt the system prompt shown in Figure 3. For each of the 13 single-page datasets, we annotate 100 samples. For the multi-page datasets, we annotated 500 samples each, except for DUDE, which received 1,000 annotations due to its higher complexity. In total, this process yielded a reasoning-annotated dataset EviBench consisting of 1.3k single-page and 3.5k multi-page document samples.

ArxivFullQA for Testing A key application of multi-page document understanding is the comprehension of scientific papers. However, large-scale multi-page benchmarks are still limited. To fill this gap, we curate a subset from the large-scale DocMatrix dataset (Laurençon et al. 2024), which contains scientific papers sourced from ArXiv. Unlike EviBench, the annotation details for ArxivFullQA differ in two key aspects. First, the prompt used for annotation is specifically tailored to scientific papers. Second, in step 1, the input consists solely of LaTeX-formatted text extracted from DocMatrix, rather than image-based document representations, in order to improve the accuracy of the generated annotations. This enriched textual input guides the model in generating QA pairs across seven categories: factual, reasoning, comparison, summary, procedural, motivation, and

Model	Parameter	DocVQA	InfoVQA	WTQ	TabFact	TextVQA	VisualMRC
UReader	7B	65.4	42.2	29.4	67.6	57.6	221.7
TextMonkey	9B	73.0	28.6	31.9	-	65.9	-
DocOwl-1.5-Chat	8B	82.2	50.7	40.6	80.2	68.6	246.4
DocOwl-2	8B	80.7	46.4	36.5	78.2	66.7	217.4
Qwen2.5-VL-Instruct	7B	95.1	82.1	62.1	78.0	84.9	277.1
DocR1	7B	95.1	82.6	63.1	79.6	81.0	251.6

Table 2: Performance comparison of 6 common single-page document benchmarks. InfoVQA is the abbreviation of the InformationVQA dataset, and WTQ is the abbreviation of the WikiTableQuestions dataset. The best results are highlighted in **bold**.

Models	Parameter	MP-DocVQA	DUDE	SlideVQA	MultiChartQA	MultiHiertt	TATDoc	ArxivFullQA	Avg.
LLaVA-NeXT-Inter	7B	39.38	25.35	29.19	28.31	9.31	10.63	5.48	21.09
LEOPARD-Idefics2	8B	66.06	40.74	34.93	18.03	10.09	2.85	14.88	26.80
mPlug-DocOwl2	8B	67.98	31.25	29.55	4.85	8.08	22.16	7.12	24.43
LLaVA-oneVison	7B	49.38	31.43	45.94	32.73	10.53	15.42	10.38	27.97
InternVL3-Instruct	38B	75.72	45.88	65.01	61.13	15.95	40.05	19.70	46.21
Qwen2.5-VL-Instruct	32B	84.79	51.14	68.67	30.10	21.25	50.38	31.28	46.80
Qwen2.5-VL-Instruct	7B	87.39	52.83	70.33	53.10	19.60	54.18	29.57	52.43
DocR1	7B	87.45	54.39	71.96	62.41	33.88	64.80	40.65	59.36

Table 3: Performance comparison on 7 text-rich multi-page image datasets. All models are evaluated using ANLS metric. The best results are highlighted in **bold**.

result questions. In Step 2, the input switches from LaTeX text to the visual format of the full paper, and the question is taken from the QA pair annotated in Step 1. The data is retained only if the model’s answer matches the previously annotated answer. Following this process, we curate a new evaluation dataset, ArxivFullQA, comprising 8.6k high-quality multi-page QA samples.

Training Recipe

We initialize our training with the Qwen2.5-VL-Instruct model (Bai et al. 2025) for two main reasons. First, large-scale chain-of-thought training data for multi-page document understanding is extremely scarce, and collecting such annotations is highly resource-intensive. Second, this instruction-tuned model already exhibits a moderate level of reasoning ability, making it a practical alternative to the conventional “cold start” phase in GRPO. To further adapt the model to our task, we adopt a two-stage curriculum training strategy within the EviGRPO framework, as illustrated in Figure 2. In the first stage, we train the model for one epoch using only single-page data. This step activates its latent reasoning capability while aligning its outputs with the expected answer format. In the second stage, we continue training on multi-page data for another epoch to enhance the model’s ability to reason over longer contexts and across multiple pages.

Experiments

Implementation Details We conduct our experiments using 8 NVIDIA A100 GPUs. During training, we set the batch size to 16, generate $G = 8$ candidate completions per sample, and use the KL penalty coefficient $\beta = 0.04$. To improve computational efficiency, we constrain the input

image resolution to a maximum of $1024 \times 28 \times 28$ during both training and evaluation.

Benchmarks We evaluate our approach using both single-page and multi-page benchmarks. For single-page evaluation, we consider six widely used datasets: DocVQA (Mathew, Karatzas, and Jawahar 2021), InformationVQA (Mathew et al. 2022), WikiTableQuestions (Pasupat and Liang 2015), TabFact (Chen et al. 2019), TextVQA (Singh et al. 2019), and VisualMRC (Tanaka, Nishida, and Yoshida 2021). For multi-page document understanding, we assess our method on six publicly available benchmarks: DUDE (Van Landeghem et al. 2023), MP-DocVQA (Tito, Karatzas, and Valveny 2023), TATDoc (Zhu et al. 2022), Multi-ChartQA (Zhu et al. 2024), MultiHiertt (Zhao et al. 2022), and SlideVQA (Tanaka et al. 2023), as well as our newly annotated benchmark, ArxivFullQA, to comprehensively evaluate its effectiveness in processing complex, multi-page inputs.

Main Result

Quantitative Results As shown in Table 2, DocR1 demonstrates strong performance across six widely used single-page document benchmarks. It performs comparably to Qwen2.5-VL-Instruct, with particularly notable improvements on WikiTableQuestions and TabFact, where it outperforms the baseline by 1.0 and 1.6 points, respectively. In Table 3, our method achieves state-of-the-art performance across all multi-page document understanding benchmarks. Remarkably, despite using significantly fewer parameters, DocR1 consistently surpasses much larger models such as Qwen2.5-VL-Instruct (32B) and InternVL-3-Instruct (38B). It attains an average score of 59.36, representing an absolute improvement of 6.93 over the baseline. Notably, our method

Images: (2)

Question: What is the Long-term debt, less current maturities average?

Ground Truth: 1597.5

Qwen2.5-VL-Instruct-7B: \$ 4,570

DocR1: <think>The table shows the Long-term debt, less current maturities for June 1, 2019 as \$641 and for June 2, 2018 as \$2,554. To find the average, we add these two values (\$641 + \$2,554 = \$3,195) and divide by 2: $\$3,195 / 2 = \$1,597.5$.</think>

<evidence_page>T,F</evidence_page><answer>1597.5</answer>

(a)

Images: (3)

Question: How many of the same countries are there in the first chart and the second chart?

Ground Truth: 4

Qwen2.5-VL-Instruct-7B: 5

DocR1: <think>The first chart shows data for the United Kingdom, United States, China, India, France, and Germany. The second chart also shows data for the United States, United Kingdom, China, and India. Therefore, the number of the same countries in both charts is 4.</think>

<evidence_page>T, T, F</evidence_page><answer>4</answer>

(c)

Images: (20)

Question: Which region of the world accounts for the highest percentage of revenues in the year 12% GROWTH is achieved?

Ground Truth: N. America

Qwen2.5-VL-Instruct-7B: Rest of world

DocR1: <think>The image shows a breakdown of Alcatel-Lucent's global revenue by region for the year 2010. The chart indicates that North America accounted for approximately 36% of the company's total revenue, while Europe contributed around 32%, Asia Pacific made up about 18%, and the rest of the world represented roughly 14%.</think>

<evidence_page>T, F, F, F, F, T, F, F, F, F, F, F, F, F, F, F, F, F, F, F</evidence_page><answer>N. America</answer>

(b)

Images: (4)

Question: Does TRICARE subtotal keeps increasing each year between 2003 and 2004 ?

Ground Truth: yes

Qwen2.5-VL-Instruct-7B: no

DocR1: <think>TRICARE subtotal increased from \$327,810 in 2003 to \$427,977 in 2004.</think>

<evidence_page>T, T, F, F</evidence_page><answer>yes</answer>

(d)

Figure 5: Examples on multi-page document QA task. It can be seen that DocR1 can not only answer the questions correctly, but also provide relatively accurate evidence pages.

yields improvements of over 10 points on four particularly challenging benchmarks—MultiChartQA, MultiHiertt, TATDoc, and ArxivFullQA. These results highlight the effectiveness of the proposed EviGRPO framework and training strategy. With only 4.8k annotated samples, DocR1 significantly improves multi-page reasoning capabilities while preserving, and in some cases slightly enhancing performance on single-page benchmarks. This demonstrates not only high data efficiency but also strong generalization from coarse-grained page-level retrieval to fine-grained semantic understanding.

Qualitative Results As illustrated in Figure 5, we provide visual comparisons between DocR1 and the baseline model across different benchmarks. In Figure 5(a), given the question “What is the Long-term debt, less current maturities average?”, the baseline model fails to locate the relevant content and outputs an incorrect answer. In contrast, DocR1 accurately identifies the first page as the sole evidence source, extracts the correct financial figure, and generates a coherent reasoning trace that leads to the correct answer. Similar trends are observed in the remaining three examples. Our evidence-first strategy significantly enhances the model’s coarse-to-fine reasoning ability, allowing it to concentrate on a small set of relevant pages. Even in challenging cases such as Figure 5(b), which involves 20 pages, DocR1 effectively identifies page 7 as the critical source of evidence and constructs an accurate reasoning path grounded in its content.

These results demonstrate the model’s strong capability in localizing sparse information and performing structured reasoning in complex multi-page scenarios.

Ablation Studies

Training Paradigm: SFT vs. GRPO vs. EviGRPO To assess the impact of different training paradigms, we compare Supervised Fine-Tuning (SFT), standard GRPO, and our proposed EviGRPO using the same set of mixed single-page and multi-page training data (denoted as mixdata). As shown in Table 4, SFT achieves only modest improvements on some benchmarks, primarily due to its limited data efficiency. Notably, ArxivFullQA is the only benchmark where SFT shows a significant gain, likely due to the model’s exposure to previously unseen data domains. In contrast, both GRPO and EviGRPO demonstrate substantial performance gains on nearly all benchmarks, underscoring the effectiveness of RL in enhancing reasoning capabilities by leveraging limited supervision more effectively. Furthermore, EviGRPO generally outperforms standard GRPO, demonstrating that its coarse-to-fine reasoning mechanism enables the model to identify relevant evidence pages prior to answer generation, thereby leading to more accurate predictions.

Data Composition: Single vs. Multi-page To assess the effectiveness of different training subsets, we compare EviGRPO trained with only single-page data or multi-page data. As shown in Table 4, training solely on single-page data

Method	MP-DocVQA	DUDE	SlideVQA	MultiChartQA	MultiHiertt	TATDoc	ArxivFullQA
Qwen2.5-VL-Instruct							
Baseline	87.39	52.83	70.33	53.10	19.60	54.18	29.57
SFT(w/ mixdata)	87.25(-0.14)	52.72(-0.11)	71.30(+0.97)	51.34(-1.76)	20.76(+1.16)	55.61(+1.43)	37.93(+8.36)
GRPO(w/ mixdata)	87.49(+0.10)	53.14(+0.31)	70.92(+0.59)	54.26(+1.16)	29.02(+9.42)	60.03(+5.85)	38.11(+8.54)
EviGRPO							
w/ single	76.95(-10.44)	46.45(-6.38)	49.11(-21.22)	49.09(-4.01)	25.26(+5.66)	56.11(+1.93)	27.04(-2.53)
w/ multi	86.82(-0.57)	52.84(+0.01)	70.71(+0.38)	60.82(+7.72)	28.33(+8.73)	64.15(+9.97)	40.03(+10.46)
w/ mixdata	86.85(-0.54)	52.93(+0.10)	71.24(+0.91)	61.54(+8.44)	29.49(+9.89)	64.51(+10.33)	39.93(+10.36)
Page Selection							
PSF-1	87.06(-0.33)	52.57(-0.26)	69.84(-0.49)	56.27(+3.17)	32.94(+13.34)	60.93(+6.75)	40.36(+10.69)
PSF-2	86.51(-0.88)	52.45(-0.38)	71.22(+0.89)	61.64(+8.54)	32.58(+12.98)	61.84(+7.66)	41.02(+11.45)
Ours	87.45(+0.06)	54.39(+1.56)	71.96(+1.63)	62.41(+9.31)	33.88(+14.28)	64.80(+10.62)	40.65(+11.08)

Table 4: Ablation experiments on training paradigms, data composition, training strategies, and page selection formats (PSF). To ensure consistency, all EviGRPO experiments adopt the same page selection format as Ours.

severely compromises the model’s multi-page reasoning capabilities, as such data lacks the need for evidence page identification—leading to significant performance drops on most multi-page tasks. In contrast, training with only multi-page data yields performance gains across many benchmarks, but the improvements are consistently smaller than those achieved when both data types are combined. These results underscore the importance of leveraging both single-page and multi-page data to fully unlock the model’s multi-page reasoning potential.

Training Strategy: Mixed vs. Curriculum To explore different strategies for combining single-page and multi-page data, we further compare our two-stage curriculum training approach with a baseline that trains on mixed data simultaneously. As shown in Table 4, although the mixed strategy also yields improvements across most benchmarks, our curriculum-based method consistently delivers greater and more generalized gains. This can be attributed to the “cold-start” effect of single-page training, which helps the model become familiar with the required output format and reasoning style before tackling more complex multi-page tasks.

Page Selection Format Variants In addition, we compare three page selection formats designed to guide evidence identification in multi-page document reasoning. PSF-1 prompts the model to directly output the indices of relevant pages (e.g., “1, 3”). PSF-2 requires the model to assign a binary label “T” (True) or “F” (False), with the total number of images explicitly provided. Ours differs from PSF-2 by omitting the image count. As shown in Table 4, PSF-1 encourages sparse selection, often leading to the omission of relevant pages due to the lack of enforced per-image decisions. PSF-2 alleviates this limitation by requiring exhaustive labeling; however, the known number of input images enables the model to heuristically align its outputs without performing genuine reasoning. In contrast, our method strengthens the reasoning requirement by omitting the image count, compelling the model to first infer the number of input pages before assigning labels accordingly.

Dataset	Qwen2.5-VL-Instruct	DocR1
MP-DocVQA	46.13	91.69(+45.56)
DUDE	56.34	85.33(+28.99)
MultiHiertt	55.62	97.38(+41.76)
Avg	52.70	91.47(+38.77)

Table 5: Evidence page recall on three multi-page document benchmarks. Recall is used as the evaluation metric to measure whether the ground truth evidence pages are successfully retrieved by the model.

Evidence Page Selection Accuracy To quantitatively evaluate the accuracy of the evidence page selection mechanism, we compute the recall of predicted evidence pages across three multi-page document understanding benchmarks. As shown in Table 5, DocR1 achieves an average recall improvement of 38.77 points over the baseline, demonstrating its superior ability to identify ground-truth evidence pages. These results underscore the reliability of the coarse-to-fine reasoning strategy employed in our framework.

Conclusion

In this work, we propose DocR1, an MLLM specifically designed for multi-page document understanding, trained under our newly introduced RL framework, EviGRPO. By incorporating evidence-aware rewards, EviGRPO encourages a human-like coarse-to-fine reasoning process, enabling the model to effectively retrieve and reason over relevant content. Supported by a two-stage annotation pipeline and a curriculum training strategy, our approach enables efficient learning under limited supervision. Extensive experiments across multiple benchmarks demonstrate that DocR1 achieves state-of-the-art performance on multi-page tasks while maintaining strong capabilities on single-page inputs. This work highlights the potential of RL to advance document-level reasoning in MLLMs.

Acknowledgments

This work was supported by the Youth Innovation Promotion Association CAS. It was also supported by the GPU cluster built by MCC Lab of Information Science and Technology Institution, USTC and the Supercomputing Center of USTC.

References

- Bai, S.; Chen, K.; Liu, X.; Wang, J.; Ge, W.; Song, S.; Dang, K.; Wang, P.; Wang, S.; Tang, J.; et al. 2025. Qwen2. 5-vl technical report. *arXiv preprint arXiv:2502.13923*.
- Chen, W.; Wang, H.; Chen, J.; Zhang, Y.; Wang, H.; Li, S.; Zhou, X.; and Wang, W. Y. 2019. Tabfact: A large-scale dataset for table-based fact verification. *arXiv preprint arXiv:1909.02164*.
- Comanici, G.; Bieber, E.; Schaekermann, M.; Pasupat, I.; Sachdeva, N.; Dhillon, I.; Blistein, M.; Ram, O.; Zhang, D.; Rosen, E.; 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*.
- Feng, H.; Liu, Q.; Liu, H.; Tang, J.; Zhou, W.; Li, H.; and Huang, C. 2024. Docpedia: Unleashing the power of large multimodal model in the frequency domain for versatile document understanding. *Science China Information Sciences*, 67: 220106.
- Feng, K.; Gong, K.; Li, B.; Guo, Z.; Wang, Y.; Peng, T.; Wu, J.; Zhang, X.; Wang, B.; and Yue, X. 2025. Video-R1: Reinforcing video reasoning in mlms. *arXiv preprint arXiv:2503.21776*.
- Guo, D.; Yang, D.; Zhang, H.; Song, J.; Zhang, R.; Xu, R.; Zhu, Q.; Ma, S.; Wang, P.; Bi, X.; et al. 2025. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*.
- Hu, A.; Xu, H.; Ye, J.; Yan, M.; Zhang, L.; Zhang, B.; Li, C.; Zhang, J.; Jin, Q.; Huang, F.; et al. 2024a. mplug-docowl 1.5: Unified structure learning for ocr-free document understanding. *arXiv preprint arXiv:2403.12895*.
- Hu, A.; Xu, H.; Zhang, L.; Ye, J.; Yan, M.; Zhang, J.; Jin, Q.; Huang, F.; and Zhou, J. 2024b. mplug-docowl2: High-resolution compressing for ocr-free multi-page document understanding. *arXiv preprint arXiv:2409.03420*.
- Huang, W.; Jia, B.; Zhai, Z.; Cao, S.; Ye, Z.; Zhao, F.; Xu, Z.; Hu, Y.; and Lin, S. 2025. Vision-R1: Incentivizing reasoning capability in multimodal large language models. *arXiv preprint arXiv:2503.06749*.
- Jaech, A.; Kalai, A.; Lerer, A.; Richardson, A.; El-Kishky, A.; Low, A.; Helyar, A.; Madry, A.; Beutel, A.; Carney, A.; et al. 2024. Openai o1 system card. *arXiv preprint arXiv:2412.16720*.
- Kafle, K.; Price, B.; Cohen, S.; and Kanan, C. 2018. Dvqa: Understanding data visualizations via question answering. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 5648–5656.
- Kahou, S. E.; Michalski, V.; Atkinson, A.; Kádár, Á.; Trischler, A.; and Bengio, Y. 2017. FigureQA: An annotated figure dataset for visual reasoning. *arXiv preprint arXiv:1710.07300*.
- Laurençon, H.; Marafioti, A.; Sanh, V.; and Tronchon, L. 2024. Building and better understanding vision-language models: insights and future directions. *arXiv preprint arXiv:2408.12637*.
- Li, X.; Yan, Z.; Meng, D.; Dong, L.; Zeng, X.; He, Y.; Wang, Y.; Qiao, Y.; Wang, Y.; and Wang, L. 2025. Videochat-R1: Enhancing spatio-temporal perception via reinforcement fine-tuning. *arXiv preprint arXiv:2504.06958*.
- Liu, Y.; Yang, B.; Liu, Q.; Li, Z.; Ma, Z.; Zhang, S.; and Bai, X. 2024. TextMonkey: An ocr-free large multimodal model for understanding document. *arXiv preprint arXiv:2403.04473*.
- Liu, Z.; Sun, Z.; Zang, Y.; Dong, X.; Cao, Y.; Duan, H.; Lin, D.; and Wang, J. 2025. Visual-RFT: Visual reinforcement fine-tuning. *arXiv preprint arXiv:2503.01785*.
- Masry, A.; Long, D. X.; Tan, J. Q.; Joty, S.; and Hoque, E. 2022. Chartqa: A benchmark for question answering about charts with visual and logical reasoning. *arXiv preprint arXiv:2203.10244*.
- Mathew, M.; Bagal, V.; Tito, R.; Karatzas, D.; Valveny, E.; and Jawahar, C. 2022. Infographicvqa. In *Proceedings of the IEEE Winter Conference on Applications of Computer Vision*, 1697–1706.
- Mathew, M.; Karatzas, D.; and Jawahar, C. 2021. Docvqa: A dataset for vqa on document images. In *Proceedings of the IEEE Winter Conference on Applications of Computer Vision*, 2200–2209.
- Mishra, A.; Shekhar, S.; Singh, A. K.; and Chakraborty, A. 2019. Ocr-vqa: Visual question answering by reading text in images. In *Proceedings of the International Conference on Document Analysis and Recognition*, 947–952.
- Mo, Y.; Shao, Z.; Ye, K.; Mao, X.; Zhang, B.; Xing, H.; Ye, P.; Huang, G.; Chen, K.; Huan, Z.; et al. 2025. Doc-CoB: Enhancing Multi-Modal Document Understanding with Visual Chain-of-Boxes Reasoning. *arXiv preprint arXiv:2505.18603*.
- Pasupat, P.; and Liang, P. 2015. Compositional semantic parsing on semi-structured tables. *arXiv preprint arXiv:1508.00305*.
- Shao, Z.; Wang, P.; Zhu, Q.; Xu, R.; Song, J.; Bi, X.; Zhang, H.; Zhang, M.; Li, Y.; Wu, Y.; et al. 2024. Deepseekmath: Pushing the limits of mathematical reasoning in open language models. *arXiv preprint arXiv:2402.03300*.
- Sidorov, O.; Hu, R.; Rohrbach, M.; and Singh, A. 2020. Textcaps: a dataset for image captioning with reading comprehension. In *Proceedings of the European Conference on Computer Vision*, 742–758.
- Singh, A.; Natarajan, V.; Shah, M.; Jiang, Y.; Chen, X.; Batra, D.; Parikh, D.; and Rohrbach, M. 2019. Towards vqa models that can read. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 8317–8326.

Stanisławek, T.; Graliński, F.; Wróblewska, A.; Lipiński, D.; Kaliska, A.; Rosalska, P.; Topolski, B.; and Biecek, P. 2021. Kleister: key information extraction datasets involving long documents with complex layouts. In *Proceedings of the International Conference on Document Analysis and Recognition*, 564–579.

Svetlichnaya, S. 2020. DeepForm: Understand structured documents at scale. *arXiv preprint*.

Tanaka, R.; Nishida, K.; Nishida, K.; Hasegawa, T.; Saito, I.; and Saito, K. 2023. Slidevqa: A dataset for document visual question answering on multiple images. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 13636–13645.

Tanaka, R.; Nishida, K.; and Yoshida, S. 2021. VisualMRC: Machine reading comprehension on document images. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 13878–13888.

Tito, R.; Karatzas, D.; and Valveny, E. 2023. Hierarchical multimodal transformers for multipage docvqa. *Pattern Recognition*, 144: 109834.

Van Landeghem, J.; Tito, R.; Borchmann, Ł.; Pietruszka, M.; Joziak, P.; Powalski, R.; Jurkiewicz, D.; Coustaty, M.; Anckaert, B.; Valveny, E.; et al. 2023. Document understanding dataset and evaluation (dude). In *Proceedings of the IEEE International Conference on Computer Vision*, 19528–19540.

Wang, Y.; Zhou, W.; Feng, H.; Zhou, K.; and Li, H. 2023. Towards improving document understanding: An exploration on text-grounding via mllms. *arXiv preprint arXiv:2311.13194*.

Ye, J.; Hu, A.; Xu, H.; Ye, Q.; Yan, M.; Dan, Y.; Zhao, C.; Xu, G.; Li, C.; Tian, J.; et al. 2023a. mplug-docowl: Modularized multimodal large language model for document understanding. *arXiv preprint arXiv:2307.02499*.

Ye, J.; Hu, A.; Xu, H.; Ye, Q.; Yan, M.; Xu, G.; Li, C.; Tian, J.; Qian, Q.; Zhang, J.; et al. 2023b. UReader: Universal ocr-free visually-situated language understanding with multimodal large language model. *arXiv preprint arXiv:2310.05126*.

Zhao, W.; Feng, H.; Liu, Q.; Tang, J.; Wu, B.; Liao, L.; Wei, S.; Ye, Y.; Liu, H.; Zhou, W.; et al. 2024. Tabpedia: Towards comprehensive visual table understanding with concept synergy. In *Proceedings of the Advances in Neural Information Processing Systems*, 7185–7212.

Zhao, Y.; Li, Y.; Li, C.; and Zhang, R. 2022. MultiHiertt: Numerical reasoning over multi hierarchical tabular and textual data. *arXiv preprint arXiv:2206.01347*.

Zhou, Y.; Chen, Y.; Lin, H.; Yang, S.; Zhu, L.; Qi, Z.; Ma, C.; and Shan, Y. 2024. Doge: Towards versatile visual document grounding and referring. *arXiv preprint arXiv:2411.17125*.

Zhu, F.; Lei, W.; Feng, F.; Wang, C.; Zhang, H.; and Chua, T.-S. 2022. Towards complex document understanding by discrete reasoning. In *Proceedings of the ACM International Conference on Multimedia*, 4857–4866.

Zhu, Z.; Jia, M.; Zhang, Z.; Li, L.; and Jiang, M. 2024. MultiChartQA: Benchmarking Vision-Language Models on Multi-Chart Problems. *arXiv preprint arXiv:2410.14179*.