

READOC: A Unified Benchmark for Realistic Document Structured Extraction

Zichao Li^{1,2,*}, Aizier Abulaiti^{1,2,*}, Yaojie Lu¹, Xuanang Chen^{1,†},
Jia Zheng^{1,†}, Hongyu Lin¹, Xianpei Han¹, Shanshan Jiang³, Bin Dong³, Le Sun¹
¹Chinese Information Processing Laboratory, Institute of Software, Chinese Academy of Sciences
²University of Chinese Academy of Sciences
³Ricoh Software Research Center Beijing Co., Ltd
{lizichao2022,aizier2022,luyaojie,chenxuanang,zhengjia}@iscas.ac.cn
{hongyu,xianpei,sunle}@iscas.ac.cn {shanshan.jiang,bin.dong}@cn.ricoh.com

Abstract

Document Structured Extraction (DSE) aims to extract structured content from raw documents. Despite the emergence of numerous DSE systems, their unified evaluation remains inadequate, significantly hindering the field’s advancement. This problem is largely attributed to existing benchmark paradigms, which exhibit fragmented and localized characteristics. To offer a thorough evaluation of DSE systems, we introduce a novel benchmark named READOC, which defines DSE as a realistic task of converting unstructured PDFs into semantically rich Markdown. The READOC dataset is derived from 3,576 diverse and real-world documents from arXiv, GitHub, and Zenodo. In addition, we develop a DSE Evaluation S³uite comprising Standardization, Segmentation and Scoring modules, to conduct a unified evaluation of state-of-the-art DSE approaches. By evaluating a range of pipeline tools, expert visual models, and general Vision-Language Models, we identify the gap between current work and the unified, realistic DSE objective for the first time. We aspire that READOC will catalyze future research in DSE, fostering more comprehensive and practical solutions.

1 Introduction

The wealth of knowledge preserved in documents is immeasurable. Document Structured Extraction (DSE), which involves converting raw documents into machine-readable structured text (Tkaczyk et al., 2015; Zhong et al., 2019; Shen et al., 2022; Lo et al., 2023), is increasingly crucial in real-world scenarios. It facilitates building extensive knowledge bases (Wang et al., 2020), constructing high-quality corpora (Jain et al., 2020), and plays a pivotal role in Retrieval-Augmented Generation

*Equal contribution.

†Corresponding author.

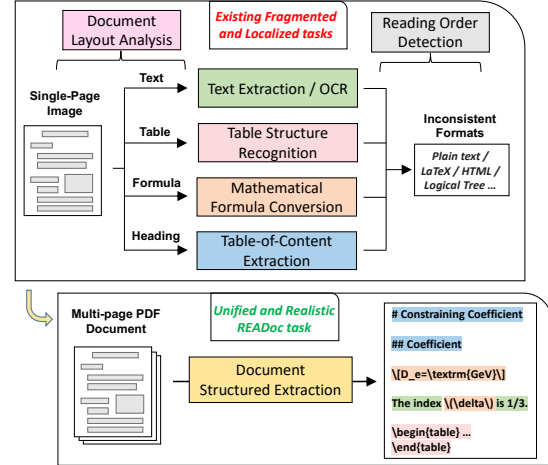


Figure 1: A comparison between fragmented and localized DSE task views and the READOC benchmark paradigm.

(RAG) (Gao et al., 2023) applications for Large Language Models (LLMs) (Achiam et al., 2023).

Recent progress in Document AI (Appalaraju et al., 2021; Huang et al., 2022; Ye et al., 2023; Hu et al., 2024) has led to the creation of numerous DSE systems (Breezedeus, 2022; Paruchuri and Lampa, 2023; Blecher et al., 2023). However, the absence of unified evaluation in real-world scenarios has left uncertainty about their performance levels and hindered their further development. This issue is largely due to the limitations of prevailing benchmark paradigms, which exhibit **fragmented** and **unrealistic** characteristics. Firstly, as depicted in Figure 1, **existing benchmarks typically fragment DSE into distinct subtasks**, including document layout analysis (Zhong et al., 2019), optical character recognition (Karatzas et al., 2015), table-of-contents extraction (Hu et al., 2022), reading order detection (Wang et al., 2021), table recognition (Smock et al., 2022) and formula conversion (Deng et al., 2017). Due to their narrow focus, diverse data sources, and inconsistent formats, existing benchmarks lack a unified framework to com-

Benchmark	Coverage						Task Paradigm	
	Layout Analysis	Character Recognition	Table Recognition	Formula Conversion	ToC. Extraction	Order Detection	INPUT	OUTPUT
PubLayNet	✓	✗	✗	✗	✗	✗	Page image	Layout blocks
DocBank	✓	✗	✗	✗	✗	✗	Page image & text	Layout blocks
Robust Reading	✗	✓	✗	✗	✗	✗	Text image	Plain text
PubTabNet	✗	✗	✓	✗	✗	✗	Table image	HTML table
TabLeX	✗	✗	✓	✗	✗	✗	Table image	TEX table
Im2Latex-100K	✗	✗	✗	✓	✗	✗	Formula image	TEX formula
ReadingBank	✗	✗	✗	✗	✗	✓	Unsorted tokens	Sorted tokens
HRDoc	✓	✗	✗	✗	✓	✗	Doc images & text	Logical tree
READOC	✓	✓	✓	✓	✓	✓	Realistic PDF doc	Markdown text

Table 1: A comparison between READOC and existing DSE benchmarks, including PubLayNet (Zhong et al., 2019), DocBank (Li et al., 2020), Robust Reading (Karatzas et al., 2015), PubTabNet (Zhong et al., 2020), TabLeX (Desai et al., 2021), Im2Latex-100K (Deng et al., 2017), ReadingBank (Wang et al., 2021), and HRDoc (Ma et al., 2023).

prehensively evaluate DSE systems. Additionally, **current research often unrealistically targets localized regions**, such as layout blocks or tables within a single page. This approach overlooks the complexity of real-world documents, which typically span multiple pages with hierarchical headings and require long-range dependencies to construct a global structure. Benchmarks focusing on isolated pages or blocks fail to provide realistic evaluations.

To address these issues, we introduce READOC, a unified benchmark designed to quantify the gap between existing work and the goal of REAListic Document Structured Extraction. READOC formulates DSE as an end-to-end task, converting multi-page PDFs into structured Markdown. We automatically construct 3,576 PDF-Markdown pairs from arXiv, GitHub, and Zenodo, covering diverse types, years, and topics to reflect real-world complexity. Additionally, we develop a DSE evaluation S³uite with three modules: Standardization, Segmentation, and Scoring, enabling unified evaluation of diverse DSE systems, including pipeline tools (Paruchuri and Lampa, 2023), expert models (Blecher et al., 2023), and Vision-Language Models (Achiam et al., 2023).

Our contributions are three-fold: 1) READOC is the first benchmark to frame DSE as a PDF-to-Markdown paradigm, which is realistic, end-to-end, and incorporates diverse data. 2) An evaluation S³uite is proposed to support the unified assessment of various DSE systems and to quantify multiple capabilities required for DSE. 3) We present the gap between current research and realistic DSE, emphasizing the importance of explor-

ing new modeling paradigms. The code and data are publicly available at <https://github.com/icip-cas/READOC>.

2 Related Work

2.1 Task Views and Benchmarks

DSE is a crucial task, yet existing benchmarks focus on discrete subtasks: document layout analysis (Zhong et al., 2019; Li et al., 2020) identifies layout blocks; optical character recognition (Karatzas et al., 2015) extracts text from images; table structure recognition (Zhong et al., 2020) transforms tables into structured formats; mathematical formula conversion (Deng et al., 2017) converts formulas into semantic formats; table-of-contents (ToC) extraction (Ma et al., 2023) constructs hierarchical heading trees; reading order detection (Wang et al., 2021) sorts page elements by reading order. We summarize relevant benchmarks in Table 1. However, their heterogeneity complicates unified DSE evaluation.

Recent research conceptualizes DSE as a single-page image-to-markup task (Blecher et al., 2023; Lee et al., 2023), and targeted benchmarks such as OmniDocBench (Ouyang et al., 2024) have emerged. Although OmniDocBench demonstrates commendable richness in evaluation and diversity of data, its single-page approach falls short in handling multi-page or lengthy documents often encountered in real-world scenarios.

2.2 Methods for Document Structured Extraction

Due to the intricacies of textual, graphical, and layout information within documents (Xu et al., 2020),

a universally accepted method for DSE has yet to emerge. A common simplistic strategy involves leveraging external parsing engines (PyMuPDF, 2024) to extract text and metadata from digital-born PDFs. With the rise of deep learning techniques, Numerous systems (Li et al., 2022; Paruchuri and Lampa, 2023; Contributors, 2024) have integrated a series of submodels into a pipeline, with each submodel dedicated to a specific subtask of DSE. Recent advancements have shifted towards end-to-end DSE methodologies, with some works leveraging Transformer (Vaswani et al., 2017) expert models to convert document page images directly into structured formats such as HTML (Lee et al., 2023) or Markdown (Blecher et al., 2023).

Recently, large Vision-Language Models have garnered widespread attention (Achiam et al., 2023; Liu et al., 2024c), with document understanding (Feng et al., 2023; Hu et al., 2024) being a key focus of their capabilities. Efforts to enhance VLMs’ document understanding capability have focused on methods like tailored training tasks (Ye et al., 2023) and resolution adaptation (Li et al., 2024), leading to notable improvements. VLMs obtain impressive results on various DSE subtasks, such as OCR (Liu et al., 2024d), table recognition (Zhao et al., 2024), and formula conversion (Xia et al., 2024). Additionally, some research has explored converting page images into structured text using VLMs (Lv et al., 2023; Wei et al., 2024a; Liu et al., 2024a). However, the lack of a unified benchmark leaves uncertainty about the gap between VLMs’ current capabilities and realistic DSE needs.

3 Task Definition

We establish a realistic task paradigm for end-to-end DSE, using raw PDF documents as input due to their prevalence and unstructured nature, which poses challenges with dispersed content and multimodal information. On the other hand, we employ Markdown as the output format, leveraging its lightweight markup to represent structural elements like headings and lists. We adopt a variant of Markdown (Blecher et al., 2023) that supports \LaTeX syntax for tables and formulas. Markdown, as the target format, can be chunked, indexed as flat text, or directly ingested by LLMs.

In summary, READOC uniformly defines DSE as a task that takes a complete PDF file as input and generates structured text in Markdown format, which is well-defined, practical, and challenging

for DSE systems. Examples of task inputs and outputs are provided in Appendix A.

4 Benchmark Construction

READOC is a unified DSE benchmark derived from real-world documents. We select heterogeneous documents from arXiv preprints¹, GitHub READMEs² and Zenodo’s Open Research Repository³, which are then automatically processed to construct PDF-Markdown pairs. READOC consists of 3,576 documents: 1,009 in the READOC-arXiv subset, 1,224 in the READOC-GitHub subset, and 1343 in the READOC-Zenodo subset. Each subset offers unique challenges: READOC-arXiv features complex academic structures such as formulas and tables, with diverse multi-column layout templates. However, its heading styles are simple, often following easily recognizable patterns like "1.1 Introduction." In contrast, READOC-GitHub includes only basic elements like paragraphs and headings, and presents a uniform single-column layout style. However, building its ToC structure is challenging due to varied and often unmarked heading styles. READOC-Zenodo contains longer documents (many exceeding 30 pages) and diverse types, such as posters, reports, theses, and books, challenging layout analysis. Additionally, its 27 languages further increase text processing and semantic understanding complexity. Each subset exhibits significant diversity in types, topics, eras, and so on, establishing READOC as a robust benchmark. We describe more construction details in Appendix B.

4.1 Document Collection and Processing

READOC-arXiv. For the arXiv preprints, the collection process involves using type keywords such as “Conference” and “Journal” to select papers and filtering for English language. Preprints without \LaTeX files or with unclear main \LaTeX files are excluded. The selected documents are first converted from \LaTeX to HTML using LaTeXXML⁴, followed by a conversion from HTML to Markdown using a modified version of the Nougat (Blecher et al., 2023) process. Only documents that complete this process without any errors and maintain correct table syntax after conversion are included.

¹<https://arxiv.org/>

²<https://github.com/>

³<https://zenodo.org/>

⁴<https://github.com/bruceMiller/LaTeXML>

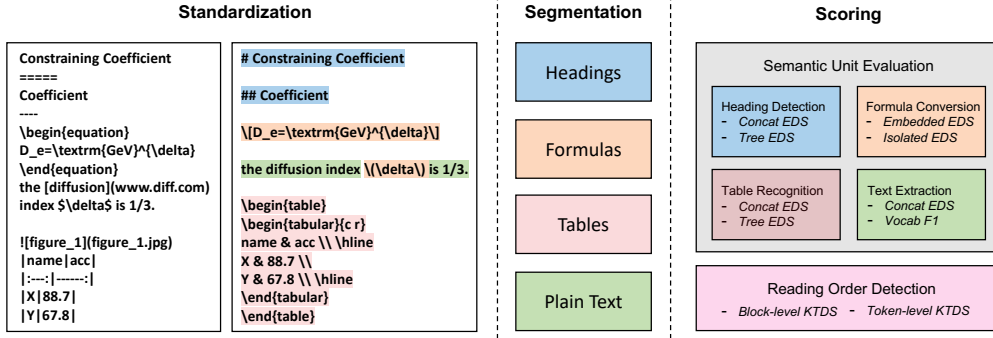


Figure 3: The three modules of the READOC Evaluation S³uite: Standardization, Segmentation and Scoring.

of texts generated by different DSE systems. The standardization includes aligning the boundaries of formulas, such as $\$, \backslashbegin{equation}$, and \backslash ; unifying different Markdown heading styles; aligning Markdown tables with \backslashbegin{table} table formats for evaluation consistency; removing images and eliminating the link syntax, etc. This module ensures that the focus remains on the core DSE capabilities, rather than being clouded by formatting disunity.

5.2 Segmentation

The second module of the S³uite divides the standardized Markdown text into distinct semantic units. To facilitate READOC in highlighting various specialized DSE capabilities of the models within a single document, we divide both the output and the ground truth Markdown text into four units: Headings of different levels, Formulas in both embedded and isolated forms, Tables, and residual Plain Text encompassing basic text and simple formatting such as bold, italic, and lists.

5.3 Scoring

The scoring module comprises two submodules:

The **Semantic Unit Evaluation** submodule implicitly assesses the models' layout analysis ability of identifying semantic units and explicitly measures four specialized capabilities: 1) **Text Extraction** refers to extracting plain text through PDF bytecode parsing or visual methods (e.g., OCR tools, VLMs). We measure edit distance similarity (EDS) after concatenating plain text and the F1 score of the plain text vocabulary. 2) **Heading Detection** involves detecting headings and constructing a hierarchical ToC tree. We measure EDS after concatenating all headings and construct ToC trees to calculate tree edit distance similarity (TEDS). 3) **Formula Conversion** involves transforming mathematical formulas into \backslash TeX format. We measure

EDS by concatenating all embedded formulas, as well as all isolated formulas. 4) **Table Recognition** involves identifying tables' structures. We evaluate EDS after concatenating all tables. Then, we convert tables into structural trees, use a maximum bipartite matching algorithm to find the optimal mapping between tables and calculate TEDS for the matched tables.

The **Reading Order Detection** submodule determines whether the model extracts document elements in the correct order. We first segment each document into blocks based on semantic unit boundaries and line breaks, create two ordered lists from these blocks, and calculate Kendall's Tau Distance Similarity (KTDS) between the lists. Additionally, we divide each document into sequential tokens, construct lists based on the positions of co-occurring tokens at their first appearance, and compute KTDS between the token-level lists.

6 Experiments

6.1 Compared Methods

Baselines. We employ three simple and widely applicable methods as baselines. The first is PyMuPDF4LLM (PyMuPDF, 2024), a PDF bytecode parsing engine that converts digital-born PDFs into Markdown using embedded metadata. The second is Tesseract (Smith, 2007), an OCR tool for text extraction and basic page segmentation. The third is MarkItDown (microsoft, 2024), a versatile tool designed for converting various file types into Markdown format.

Pipeline Tools. We evaluate three tools that support PDF-to-Markdown functionality, which integrate complex engineering with PDF parsing engines and advanced deep learning submodels. The pipeline tools we evaluate include the following. Marker (Paruchuri and Lampa, 2023), MinerU

Methods	Semantic Unit Evaluation								Reading Order		Average
	Text		Heading		Formula		Table		Block	Token	
	Concat	Vocab	Concat	Tree	Embed	Isolate	Concat	Tree			
Baselines											
PyMuPDF4LLM	66.66	74.27	27.86	20.77	0.07	0.02	23.27	15.83	87.70	89.09	40.55
Tesseract OCR	78.85	76.51	1.26	0.30	0.12	0.00	0.00	0.00	96.70	97.59	35.13
MarkItDown	73.88	79.64	2.90	0.93	0.10	0.00	0.34	0.08	97.13	97.08	35.21
Pipeline Tools											
MinerU	88.32	91.22	67.06	41.97	62.77	70.76	59.34	52.85	98.52	97.90	73.07
Pix2Text	85.85	83.72	63.23	34.53	43.18	37.45	54.08	47.35	97.68	96.78	64.39
Marker	79.11	82.71	63.60	39.39	3.47	48.74	64.61	72.36	98.04	97.74	64.98
Docling	79.73	85.39	68.74	38.33	0.23	0.0	54.09	66.56	98.05	97.18	58.83
Expert Visual Models											
Nougat-small	87.35	92.00	86.40	87.88	76.52	79.39	55.63	52.35	97.97	98.36	81.38
Nougat-base	88.03	92.29	86.60	88.50	76.19	79.47	54.40	52.30	97.98	98.41	81.42
GOT-OCR 2.0	84.47	86.24	66.69	57.68	53.48	56.23	50.40	34.50	97.73	97.50	68.49
Vision-Language Models											
DeepSeek-VL-7B-Chat	31.89	39.96	23.66	12.53	17.01	16.94	22.96	16.47	88.76	66.75	33.69
MiniCPM-Llama3-V2.5	58.91	70.87	26.33	7.68	16.70	17.90	27.89	24.91	95.26	93.02	43.95
LLaVa-1.6-Vicuna-13B	27.51	37.09	8.92	6.27	17.80	11.68	23.78	16.23	76.63	51.68	27.76
InternVL-Chat-V1.5	53.06	68.44	25.03	13.57	33.13	24.37	40.44	34.35	94.61	91.31	47.83
GPT-4o-mini	79.44	84.37	31.77	18.65	42.23	41.67	47.81	39.85	97.69	96.35	57.98

Table 3: Evaluation of various Document Structured Extraction systems on READOC-arXiv.

Methods	Semantic Unit Evaluation						Reading Order		Average
	Text		Heading		Table		Block	Token	
	Concat	Vocab	Concat	Tree	Concat	Tree			
MinerU	57.28	59.95	30.73	22.83	38.75	26.82	65.46	66.87	46.08
Marker	59.34	61.68	30.28	18.29	40.68	29.10	65.77	66.40	46.44
Nougat-base	57.54	66.87	35.99	26.98	13.99	11.55	93.56	93.01	49.94
GPT-4o-mini	64.16	71.76	25.07	15.4	45.75	31.88	95.23	95.2	55.56

Table 5: Evaluation of three representative Document Structured Extraction systems on READOC-Zenodo.

(Contributors, 2024), Pix2Text (Breezedeus, 2022) and Docling (Auer et al., 2024).

Expert Visual Models. We evaluate Nougat-small and Nougat-base (Blecher et al., 2023), specialized transformer models trained on arXiv academic documents under the single-page image-to-Markdown paradigm, with parameter sizes of 250M and 350M, respectively. Additionally, we evaluate the GOT-OCR 2.0 model (Wei et al., 2024b), a unified, end-to-end OCR system with 580M parameters.

Vision-Language Models. We evaluate VLMs using the same single-page image-to-Markdown paradigm as expert models. For efficiency, we select open-source models with fewer than 30 billion parameters and lightweight proprietary models. Only VLMs with basic instruction-following and preliminary Markdown understanding are retained. The retained VLMs include: the open-source

models DeepSeek-VL-7B-Chat (Lu et al., 2024), MiniCPM-Llama3-V2.5 (Hu et al., 2023), LLaVa-1.6-Vicuna-13B (Liu et al., 2024b), InternVL-Chat-V1.5 (Chen et al., 2024), and the proprietary model GPT-4o-mini (Achiam et al., 2023). More implementation details are included in Appendix D.

6.2 Experimental Results

Results for READOC-arXiv, READOC-GitHub and READOC-Zenodo are presented in Table 3, 4 and 5, respectively. We draw insights from both DSE system categories and specialized capabilities.

From the point of DSE system categories, we observe that: 1) **Pipeline tools are often plagued by complex engineering challenges.** Docling falls short in recognizing embedded and isolated formulas, while Marker struggles with detecting embedded formulas, issues that are overlooked within the complex pipeline designs. In contrast, MinerU performs significantly better than other pipeline

Methods	Semantic Unit Evaluation				Reading Order		Avg.
	Text		Heading		Block	Token	
	Concat	Vocab	Concat	Tree			
Baselines							
PyMuPDF4LLM	85.21	77.27	12.13	11.05	98.61	98.43	63.78
Tesseract OCR	82.06	82.32	6.65	4.25	98.48	99.01	62.13
MarkItDown	85.23	85.69	2.07	0.64	99.4	99.42	62.07
Pipeline Tools							
MinerU	84.46	84.78	67.24	47.15	99.51	99.18	80.39
Pix2Text	78.99	78.51	60.53	39.42	97.94	97.38	75.46
Marker	89.50	88.11	72.81	37.51	99.03	99.13	81.02
Docling	44.53	50.73	40.17	22.8	70.77	70.19	49.87
Expert Visual Models							
Nougat-small	76.04	77.11	62.81	38.73	98.33	96.12	74.86
Nougat-base	75.26	76.79	62.39	37.01	97.65	95.62	74.12
GOT-OCR 2.0	78.18	84.53	62.08	50.06	98.40	98.60	78.64
Vision-Language Models							
DeepSeek-VL-7B-Chat	34.85	41.97	38.69	20.15	96.47	84.55	52.78
MiniCPM-Llama3-V2.5	71.61	78.88	39.66	22.99	97.91	97.97	68.17
LLaVa-1.6-Vicuna-13B	37.03	52.88	27.49	16.13	95.65	90.45	53.27
InternVL-Chat-V1.5	72.59	78.27	56.85	32.70	98.39	97.69	72.75
GPT-4o-mini	85.06	89.41	62.65	42.04	98.81	99.05	79.50

Table 4: Evaluation of various DSE systems on READOC-GitHub.

Methods	Modeling Paradigm	Semantic Unit Evaluation								Reading Order		Average
		Text		Heading		Formula		Table		Block	Token	
		Concat	Vocab	Concat	Tree	Embed	Isolate	Concat	Tree			
GPT-4o-mini	Single Page	77.95	83.10	40.63	23.82	39.03	40.05	54.55	44.70	97.79	96.61	59.82
	Multiple Pages	70.28	78.66	62.71	51.87	33.22	33.62	43.50	34.09	98.87	96.31	60.31

Table 6: Comparison of page-level modeling paradigms, for documents within 5 pages of READOC-arXiv.

models in both embedded and isolated formula recognition. Besides, Pix2Text encounters program crashes and fails to process certain PDFs (3 files on READOC-arXiv and 4 on READOC-GitHub), posing a significant usability issue. 2) **Expert models struggle with generalization and scalability issues.** Nougat performs well on READOC-arXiv but declines markedly on READOC-GitHub (i.e., from 81.42 to 74.12) and READOC-Zenodo (e.g., from 81.42 to 49.94), which features simpler layouts and fewer semantic units, indicating poor transfer learning ability. Moreover, scaling up from Nougat-small to Nougat-base does not boost performance (+0.04 on READOC-arXiv, -0.74 on READOC-GitHub). Additionally, while GOT-OCR 2.0 scores lower than Nougat-small and Nougat-base on READOC-arXiv, it outperforms them on the READOC-GitHub subset. 3) **VLMs generally underperform in complex academic documents.** The best-performing open-source model, InternVL-Chat-V1.5, scores 47.83 on READOC-arXiv, while the proprietary model

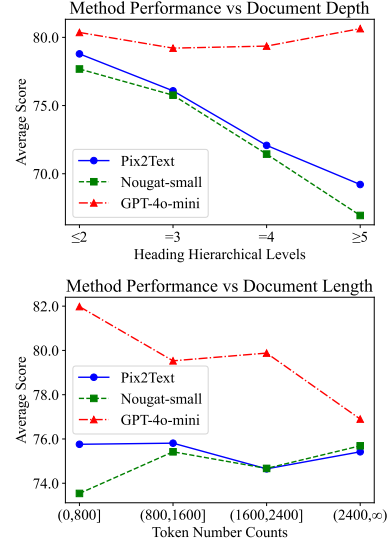


Figure 4: Relationship between DSE systems' performance and the depth or length of documents in READOC-GitHub.

GPT-4o-mini scores 57.98, both of which are lower than the pipeline tools. On READOC-Zenodo, GPT-4o-mini achieves an average score of 55.56, suggesting its potential in handling multilingual and multi-format documents.

From the perspective of specialized capabilities, we observe that: 1) **Building hierarchical ToC trees from a global perspective remains a significant challenge**, as existing systems predominantly focus on single-page images. Pipeline tools lack modules to assess heading depth, leading to substantial drops in Tree EDS compared to Concat EDS. Expert models can exhibit strong ToC construction for specific documents, which is more a superficial imitation rather than a semantic understanding of the logical structure. On READOC-arXiv, Nougat-base scores 88.50 in TEDS, while on READOC-GitHub, it drops to 37.01. 2) **Understanding localized structured data such as tables and formulas is relatively difficult.** VLMs perform poorly on these tasks. Even the expert model Nougat-base, trained on arXiv documents,

Methods	Semantic Unit Evaluation (Avg.)		
	Single-col.	Multi-col.	Drop ↓
MinerU	52.75	48.34	4.41
Pix2Text	56.93	55.34	1.59
Marker	56.70	53.39	3.31
Nougat-small	80.51	74.09	6.42
Nougat-base	80.48	74.16	6.32
InternVL-Chat-V1.5	39.96	32.83	7.13
GPT-4o-mini	49.23	47.08	2.15

Table 7: Relationship between DSE systems’ semantic unit evaluation and the layout of documents in READOC-arXiv.

Methods	Modeling Paradigm	Heading	
		Concat	Tree
Pix2Text	Single Page	63.38	38.89
Nougat-small	Single Page	64.87	39.53
GPT-4o-mini	Single Page	69.78	45.12
	Multiple Pages	83.71	68.11

Table 8: Comparison of page-level modeling paradigms, for documents within 5 pages of READOC-GitHub.

has shown only modest performance, with average metrics of 65.56 on these two tasks. 3) **Reading Order Detection is a relatively easy capability to acquire.** The baseline tool Tesseract, which uses heuristics for page segmentation, scores 96.70 and 98.48 in block-level KTDS on READOC-arXiv and READOC-GitHub, respectively.

6.3 Fine-grained Analysis

Impact of Document Length and Depth. Figure 4 displays the results of three representative DSE systems on READOC-GitHub. Pipeline tools and expert models exhibit similar performance trends, remaining stable with variations in document length but declining sharply as document depth increases. In contrast, VLMs demonstrate stability with changes in document depth, while their performance decreases as document length increases. This indicates that different DSE systems exhibit distinct shortcomings in realistic scenarios, which have not been previously revealed.

Impact of Document Layout. Using pdfplumber (Singer-Vine and The pdfplumber contributors, 2024) and heuristic rules, we classify documents in READOC-arXiv into single- and multi-column categories, representing different layout complexities. Table 7 illustrates the average semantic unit evaluation scores across the two document types. All

Methods	Time Cost (s) per Document	NVIDIA GPU Devices
Marker	23.86	1× Titan RTX (24GB)
MinerU	30.96	
Nougat-small	51.34	
Nougat-base	101.07	
Pix2Text	188.10	
MiniCPM-Llama3-V2.5	392.37	1× A100 (80GB)
InternVL-Chat-V1.5	1,182.02	2× A100 (80GB)

Table 9: Comparison of Efficiency of DSE systems.

systems show performance degradation on complex multi-column documents, highlighting that our semantic unit evaluation implicitly measures the layout analysis capability. Among the systems, GPT-4o-mini exhibits the best layout analysis capability, while InternVL-Chat-V1.5 shows the most significant performance decline, reflecting substantial differences in performance levels among VLMs.

Exploration of the Multi-Page Paradigm. Previous researches focus on processing single pages and perform poorly in constructing global ToC trees. To explore how DSE systems might utilize global information, we investigate a paradigm where multiple pages are processed simultaneously. Specifically, we employ GPT-4o-mini to receive all page images of the document at once and convert them into Markdown text. We conduct experiments on documents with up to 5 pages, as shown in Tables 6 and 8. While this method significantly enhances global ToC construction compared to the single-page paradigm, processing multiple pages simultaneously reduces local fine-grained capabilities, such as table and formula conversion, indicating that DSE modeling for multi-page documents still needs further development.

Considerations of Efficiency. DSE focuses on practical efficiency, which may involve real-time RAG calls or large-scale corpus construction. We sample 50 documents from READOC-arXiv and measure the throughput of DSE systems, as shown in Table 9. Despite significant advancements in GPU memory and computational power, VLMs remains considerably lower efficiency compared to pipeline tools and expert Transformer models, indicating the need for future improvements in not only performance but also efficiency.

6.4 Case Study

We compare results from four representative systems in Figure 5. Key observations include: 1)

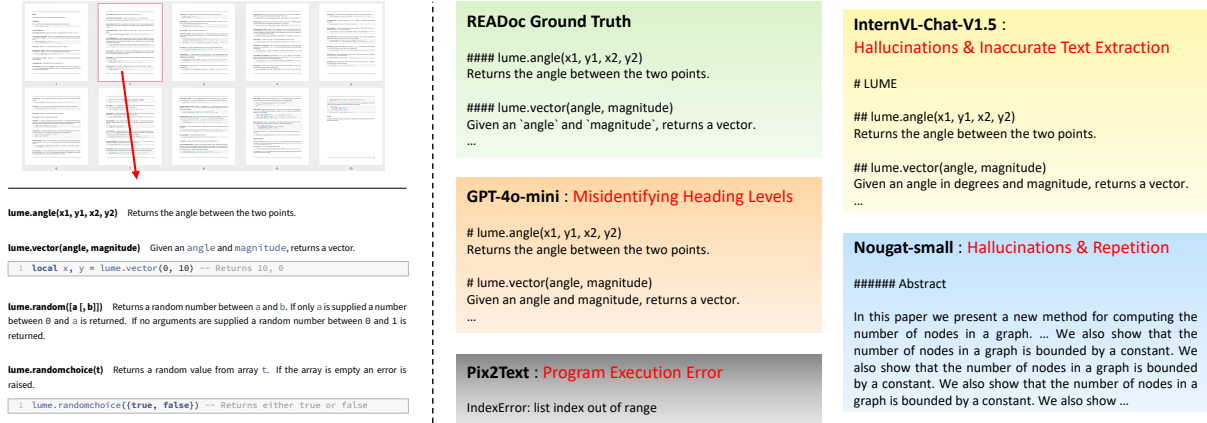


Figure 5: A case study from READoc-GitHub. More cases are in Appendix D.

GPT-4o-mini misclassifies heading levels, revealing the limitations of single-page paradigms that fail to globally perceive the document’s logical structure. 2) InternVL-Chat-V1.5 exhibits hallucinations and inaccurate text extraction, illustrating the differences in DSE capabilities between open-source and proprietary VLMs. 3) Nougat-small fabricates content completely unrelated to the original document, reflecting the poor generalization ability of expert models. 4) Pix2Text triggers an execution error, demonstrating the complexity in developing pipeline tools.

7 Conclusion

This paper proposes READOC, a novel benchmark that frames document structured extraction as a realistic, end-to-end task, i.e., transforming unstructured PDFs into semantically rich Markdown text. Based on an evaluation S³uite, We conduct a unified evaluation of state-of-the-art approaches, including pipeline tools, expert models and general VLMs. Our experimental results reveal critical gaps in current researches when applied to realistic scenarios and underscore the importance of exploring new research paradigms.

Limitations

Our work has two main limitations: 1) First, there exists some noise in the PDF-Markdown pairs generated through the automated framework. Although we have implemented various filtering and post-processing methods to minimize this impact, it remains difficult to eliminate completely. In fact, this is a common challenge in current document processing benchmark fields, and we will continue to explore more efficient and accurate processing

paradigms in the future. 2) Second, although we introduce a Standardization module in our DSE evaluation S³uite to unify the outputs from different models to the ground truth format, there are still some scenarios we cannot fully account for. We plan to introduce more comprehensive format unification modules in future work.

Ethics Statement

All the data, tools and model weights we use come from publicly available sources. We only use them for evaluation purposes in document structured extraction. When using these resources for this study, we strictly adhere to their licensing agreements.

Acknowledgment

We sincerely thank the reviewers for their insightful comments and valuable suggestions. This work was supported by Beijing Municipal Science and Technology Project (Nos. Z231100010323002), Beijing Natural Science Foundation (L243006), the Natural Science Foundation of China (No. 62306303, 62476265) and the Basic Research Program of ISCAS (Grant No. ISCAS-JCZD-202401).

References

- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altschmidt, Sam Altman, Shyamal Anadkat, et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Srikar Appalaraju, Bhavan Jasani, Bhargava Urala Kota, Yusheng Xie, and R Manmatha. 2021. Docformer: End-to-end transformer for document understanding. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 993–1003.

- Christoph Auer, Maksym Lysak, Ahmed Nassar, Michele Dolfi, Nikolaos Livathinos, Panos Vagenas, Cesar Berrospi Ramis, Matteo Omenetti, Fabian Lindlbauer, Kasper Dinkla, Lokesh Mishra, Yusik Kim, Shubham Gupta, Rafael Teixeira de Lima, Valery Weber, Lucas Morin, Ingmar Meijer, Viktor Kuropiatnyk, and Peter W. J. Staar. 2024. *Docling technical report*. Preprint, arXiv:2408.09869.
- Lukas Blecher, Guillem Cucurull, Thomas Scialom, and Robert Stojnic. 2023. Nougat: Neural optical understanding for academic documents. *arXiv preprint arXiv:2308.13418*.
- Breezedeus. 2022. Pix2text (p2t). <https://github.com/breezedeus/pix2text>. Accessed: 2024-07-18.
- Zhe Chen, Weiyun Wang, Hao Tian, Shenglong Ye, Zhangwei Gao, Erfei Cui, Wenwen Tong, Kongzhi Hu, Jiapeng Luo, Zheng Ma, et al. 2024. How far are we to gpt-4v? closing the gap to commercial multimodal models with open-source suites. *arXiv preprint arXiv:2404.16821*.
- MinerU Contributors. 2024. MinerU: A one-stop, open-source, high-quality data extraction tool. <https://github.com/opendatalab/MinerU>. Accessed: 2025-04-14.
- Yuntian Deng, Anssi Kanervisto, Jeffrey Ling, and Alexander M Rush. 2017. Image-to-markup generation with coarse-to-fine attention. In *International Conference on Machine Learning*, pages 980–989. PMLR.
- Harsh Desai, Pratik Kayal, and Mayank Singh. 2021. Tablex: a benchmark dataset for structure and content information extraction from scientific tables. In *Document Analysis and Recognition-ICDAR 2021: 16th International Conference, Lausanne, Switzerland, September 5–10, 2021, Proceedings, Part II 16*, pages 554–569. Springer.
- Hao Feng, Qi Liu, Hao Liu, Wengang Zhou, Houqiang Li, and Can Huang. 2023. Docpedia: Unleashing the power of large multimodal model in the frequency domain for versatile document understanding. *arXiv preprint arXiv:2311.11810*.
- Yunfan Gao, Yun Xiong, Xinyu Gao, Kangxiang Jia, Jinliu Pan, Yuxi Bi, Yi Dai, Jiawei Sun, and Haofen Wang. 2023. Retrieval-augmented generation for large language models: A survey. *arXiv preprint arXiv:2312.10997*.
- Anwen Hu, Haiyang Xu, Jiabo Ye, Ming Yan, Liang Zhang, Bo Zhang, Chen Li, Ji Zhang, Qin Jin, Fei Huang, et al. 2024. mplug-docowl 1.5: Unified structure learning for ocr-free document understanding. *arXiv preprint arXiv:2403.12895*.
- Jinyi Hu, Yuan Yao, Chongyi Wang, Shan Wang, Yinxu Pan, Qianyu Chen, Tianyu Yu, Hanghao Wu, Yue Zhao, Haoye Zhang, Xu Han, Yankai Lin, Jiao Xue, Dahai Li, Zhiyuan Liu, and Maosong Sun. 2023. Large multilingual models pivot zero-shot multimodal learning across languages. *arXiv preprint arXiv:2308.12038*.
- Pengfei Hu, Zhenrong Zhang, Jianshu Zhang, Jun Du, and Jiajia Wu. 2022. Multimodal tree decoder for table of contents extraction in document images. In *2022 26th international conference on pattern recognition (ICPR)*, pages 1756–1762. IEEE.
- Yupan Huang, Tengchao Lv, Lei Cui, Yutong Lu, and Furu Wei. 2022. Layoutlmv3: Pre-training for document ai with unified text and image masking. In *Proceedings of the 30th ACM International Conference on Multimedia*, pages 4083–4091.
- Sarthak Jain, Madeleine van Zuylen, Hannaneh Hajishirzi, and Iz Beltagy. 2020. Scirex: A challenge dataset for document-level information extraction. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7506–7516.
- Dimosthenis Karatzas, Lluís Gomez-Bigorda, Angelos Nicolaou, Suman Ghosh, Andrew Bagdanov, Masakazu Iwamura, Jiri Matas, Lukas Neumann, Vijay Ramaseshan Chandrasekhar, Shijian Lu, et al. 2015. Icdar 2015 competition on robust reading. In *2015 13th international conference on document analysis and recognition (ICDAR)*, pages 1156–1160. IEEE.
- Kenton Lee, Mandar Joshi, Iulia Raluca Turc, Hexiang Hu, Fangyu Liu, Julian Martin Eisenschlos, Urvasi Khandelwal, Peter Shaw, Ming-Wei Chang, and Kristina Toutanova. 2023. Pix2struct: Screen-shot parsing as pretraining for visual language understanding. In *International Conference on Machine Learning*, pages 18893–18912. PMLR.
- Chenxia Li, Ruoyu Guo, Jun Zhou, Mengtao An, Yuning Du, Lingfeng Zhu, Yi Liu, Xiaoguang Hu, and Dianhai Yu. 2022. Pp-structurev2: A stronger document analysis system. *arXiv preprint arXiv:2210.05391*.
- Minghao Li, Yiheng Xu, Lei Cui, Shaohan Huang, Furu Wei, Zhoujun Li, and Ming Zhou. 2020. Docbank: A benchmark dataset for document layout analysis. In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 949–960.
- Zhang Li, Biao Yang, Qiang Liu, Zhiyin Ma, Shuo Zhang, Jingxu Yang, Yabo Sun, Yuliang Liu, and Xiang Bai. 2024. Monkey: Image resolution and text label are important things for large multi-modal models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 26763–26773.
- Chenglong Liu, Haoran Wei, Jinyue Chen, Lingyu Kong, Zheng Ge, Zining Zhu, Liang Zhao, Jianjian Sun, Chunrui Han, and Xiangyu Zhang. 2024a. Focus anywhere for fine-grained multi-page document understanding. *arXiv preprint arXiv:2405.14295*.

- Haotian Liu, Chunyuan Li, Yuheng Li, Bo Li, Yuanhan Zhang, Sheng Shen, and Yong Jae Lee. 2024b. Llava-next: Improved reasoning, ocr, and world knowledge.
- Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. 2024c. Visual instruction tuning. *Advances in neural information processing systems*, 36.
- Yuliang Liu, Biao Yang, Qiang Liu, Zhang Li, Zhiyin Ma, Shuo Zhang, and Xiang Bai. 2024d. Textmonkey: An ocr-free large multimodal model for understanding document. *arXiv preprint arXiv:2403.04473*.
- Kyle Lo, Zejiang Shen, Benjamin Newman, Joseph Z Chang, Russell Authur, Erin Bransom, Stefan Candra, Yoganand Chandrasekhar, Regan Huff, Bailey Kuehl, et al. 2023. Papermage: A unified toolkit for processing, representing, and manipulating visually-rich scientific documents. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 495–507.
- Haoyu Lu, Wen Liu, Bo Zhang, Bingxuan Wang, Kai Dong, Bo Liu, Jingxiang Sun, Tongzheng Ren, Zhuoshu Li, Yaofeng Sun, et al. 2024. Deepseek-vl: towards real-world vision-language understanding. *arXiv preprint arXiv:2403.05525*.
- Tengchao Lv, Yupan Huang, Jingye Chen, Lei Cui, Shuming Ma, Yaoyao Chang, Shaohan Huang, Wenhui Wang, Li Dong, Weiyao Luo, et al. 2023. Kosmos-2.5: A multimodal literate model. *arXiv preprint arXiv:2309.11419*.
- Jiefeng Ma, Jun Du, Pengfei Hu, Zhenrong Zhang, Jian-shu Zhang, Huihui Zhu, and Cong Liu. 2023. Hrdoc: Dataset and baseline method toward hierarchical reconstruction of document structures. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pages 1870–1877.
- microsoft. 2024. Markitdown. <https://github.com/microsoft/markitdown>. Accessed: 2025-01-20.
- Linke Ouyang, Yuan Qu, Hongbin Zhou, Jiawei Zhu, Rui Zhang, Qunshu Lin, Bin Wang, Zhiyuan Zhao, Man Jiang, Xiaomeng Zhao, et al. 2024. Omnidocbench: Benchmarking diverse pdf document parsing with comprehensive annotations. *arXiv preprint arXiv:2412.07626*.
- Vik Paruchuri and Samuel Lampa. 2023. Marker: Convert pdf to markdown quickly with high accuracy. <https://github.com/VikParuchuri/marker>. Accessed: 2025-04-14.
- PyMuPDF. 2024. Pymupdf4llm. <https://github.com/pymupdf/RAG>. Accessed: 2024-07-26.
- Zejiang Shen, Kyle Lo, Lucy Lu Wang, Bailey Kuehl, Daniel S Weld, and Doug Downey. 2022. Vila: Improving structured content extraction from scientific pdfs using visual layout groups. *Transactions of the Association for Computational Linguistics*, 10:376–392.
- Jeremy Singer-Vine and The pdfplumber contributors. 2024. *pdfplumber*. Accessed: 2024-08-05.
- Ray Smith. 2007. *An overview of the tesseract ocr engine*. In *ICDAR '07: Proceedings of the Ninth International Conference on Document Analysis and Recognition*, pages 629–633, Washington, DC, USA. IEEE Computer Society.
- Brandon Smock, Rohith Pesala, and Robin Abraham. 2022. Pubtables-1m: Towards comprehensive table extraction from unstructured documents. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4634–4642.
- Dominika Tkaczyk, Paweł Szostek, Mateusz Fedorczyk, Piotr Jan Dendek, and Łukasz Bolikowski. 2015. Cermine: automatic extraction of structured metadata from scientific literature. *International Journal on Document Analysis and Recognition (IJ-DAR)*, 18:317–335.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.
- Lucy Lu Wang, Kyle Lo, Yoganand Chandrasekhar, Russell Reas, Jiangjiang Yang, Doug Burdick, Darrin Eide, Kathryn Funk, Yannis Katsis, Rodney Kinney, et al. 2020. Cord-19: The covid-19 open research dataset. In *Annual Meeting of the Association for Computational Linguistics*.
- Zilong Wang, Yiheng Xu, Lei Cui, Jingbo Shang, and Furu Wei. 2021. Layoutreader: Pre-training of text and layout for reading order detection. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 4735–4744.
- Haoran Wei, Lingyu Kong, Jinyue Chen, Liang Zhao, Zheng Ge, En Yu, Jianjian Sun, Chunrui Han, and Xiangyu Zhang. 2024a. Small language model meets with reinforced vision vocabulary. *arXiv preprint arXiv:2401.12503*.
- Haoran Wei, Chenglong Liu, Jinyue Chen, Jia Wang, Lingyu Kong, Yanming Xu, Zheng Ge, Liang Zhao, Jianjian Sun, Yuang Peng, Chunrui Han, and Xiangyu Zhang. 2024b. *General ocr theory: Towards ocr-2.0 via a unified end-to-end model*. *Preprint*, arXiv:2409.01704.
- Renqiu Xia, Song Mao, Xiangchao Yan, Hongbin Zhou, Bo Zhang, Haoyang Peng, Jiahao Pi, Daocheng Fu, Wenjie Wu, Hancheng Ye, et al. 2024. Docgenome: An open large-scale scientific document benchmark for training and testing multi-modal large language models. *arXiv preprint arXiv:2406.11633*.
- Yiheng Xu, Minghao Li, Lei Cui, Shaohan Huang, Furu Wei, and Ming Zhou. 2020. Layoutlm: Pre-training of text and layout for document image understanding. In *Proceedings of the 26th ACM SIGKDD international conference on knowledge discovery & data mining*, pages 1192–1200.

Jiabo Ye, Anwen Hu, Haiyang Xu, Qinghao Ye, Ming Yan, Guohai Xu, Chenliang Li, Junfeng Tian, Qi Qian, Ji Zhang, et al. 2023. Ureader: Universal ocr-free visually-situated language understanding with multimodal large language model. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 2841–2858.

Weichao Zhao, Hao Feng, Qi Liu, Jingqun Tang, Shu Wei, Binghong Wu, Lei Liao, Yongjie Ye, Hao Liu, Houqiang Li, et al. 2024. Tabpedia: Towards comprehensive visual table understanding with concept synergy. *arXiv preprint arXiv:2406.01326*.

Xu Zhong, Elaheh ShafieiBavani, and Antonio Jimeno Yepes. 2020. Image-based table recognition: data, model, and evaluation. In *European conference on computer vision*, pages 564–580. Springer.

Xu Zhong, Jianbin Tang, and Antonio Jimeno Yepes. 2019. Publaynet: largest dataset ever for document layout analysis. In *2019 International conference on document analysis and recognition (ICDAR)*, pages 1015–1022. IEEE.

A Task Example

This section provides an example of input and output for the Document Structured Extraction (DSE) task as defined by READOC. As illustrated in Figure 6, DSE systems are required to process a multi-page, real-world PDF document as input and produce a structured Markdown text as output.

B Details of Dataset

B.1 Construction Details

Type Keywords of READOC-arXiv. As we aim to improve data diversity in READOC across types, disciplines, topics, and so on, we use custom keywords to capture the document category types in READOC-arXiv, as illustrated in Table 10. We treat the comments of each arXiv preprint as retrieval targets, determining whether they contain the keywords to map them to specific document types. The document type distribution of READOC-arXiv is illustrated in Figure 10.

Modification of Nougat process. As described in the main body of the paper, we modify Nougat’s source code (Blecher et al., 2023) to convert arXiv documents in HTML format into Markdown text. Specifically, our modifications include: 1) improved support for converting the *Tableofcontents* command; 2) recognition and conversion of nested lists; 3) support for identifying and converting subtables and subfigures; and 4) removal of line breaks within headings in the Markdown format.

Document Type	Keywords
Workshop	<i>Workshop</i>
Conference	<i>Conference</i>
Journal	<i>Journal</i>
Dissertation	<i>Thesis, Dissertation</i>
Guide	<i>Handbook, Manual, Guide, Tutorial, Technical Note</i>
Others	-

Table 10: Document type keywords of READOC-arXiv.

B.2 Additional Data Statistics

In addition to the diversity in disciplines, topics and languages, we provide supplementary statistical information here. This includes the distribution of the READOC dataset by year, as shown in Figure 7, and the distribution by page count, depth, and length, as illustrated in Figure 8 and Figure 9.

C Details of Evaluation Suite

C.1 Standardization Details

We standardized the Markdown text output from various tools and models according to a set of specific rules:

- **Alignment of different formula boundaries.** For isolated formulas, we standardize the starting boundary by converting $\$, \begin{equation}, \begin{gather}, \begin{multline}, \begin{equation}$, to \backslash , with a similar format for the ending boundary. For embedded formulas, we convert the starting boundary $\$$ to \backslash , with a similar format for the ending boundary.
- **Heading formatting.** We standardize document headings by converting them to text lines beginning with consecutive “#” symbols, where the number of “#” indicates the heading level.
- **Removal of certain elements.** External URLs are removed from the link format and image references are also removed.
- **Table formatting.** We standardize tables by converting Markdown-compliant tables to LaTeX format. The reason for using LaTeX as

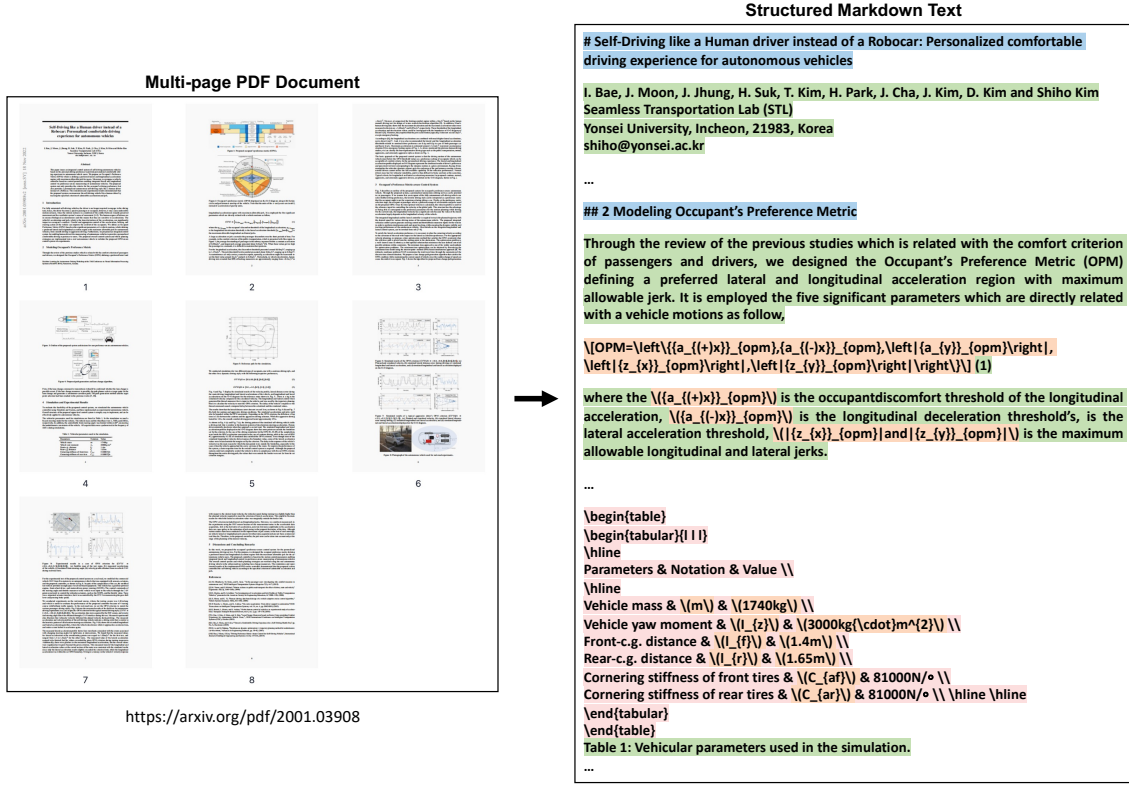


Figure 6: Input-Output example of READOC task.

the standardized table format is that, compared to native Markdown, it can represent more complex features such as multi-row and multi-column layouts.

C.2 Segmentation Details

The segmentation module primarily based on a series of regular expressions written in Python, as illustrated in Figure 11.

C.3 Scoring Details

This section mainly discusses the metric calculations involved in the scoring module. Specifically, we focus on the following three similarity measures:

1) **Edit Distance Similarity (EDS)**: The edit distance $ED(A, B)$ is defined as the minimum number of single-character edits (insertions, deletions, or substitutions) required to change string A into string B . The edit distance similarity is calculated using the formula:

$$EDS(A, B) = 1 - \frac{ED(A, B)}{\max(|A|, |B|)}$$

where $|A|$ and $|B|$ are the lengths of strings A and B , respectively.

2) **Tree Edit Distance Similarity (TEDS)**: The tree edit distance $TED(T_1, T_2)$ is the minimum number of operations needed to transform one tree T_1 into another tree T_2 . The operations typically include insertion, deletion, and relabeling of nodes. The tree edit distance similarity is computed as follows:


$$TEDS(T_1, T_2) = 1 - \frac{TED(T_1, T_2)}{\max(|T_1|, |T_2|)}$$

where $|T_1|$ and $|T_2|$ represent the sizes of trees T_1 and T_2 , respectively.


3) **Kendall's Tau Distance Similarity (KTDS)**: The Kendall's Tau distance measures the ordinal association between two rankings. The number of discordant pairs K_d is defined as the number of pairs where x_i and x_j are in different orders. The Kendall's Tau distance similarity is given by the formula:

$$KTDS(X, Y) = 1 - \frac{2 \cdot K_d}{n(n-1)}$$

where n is the total number of items.



(a)



(b)

Figure 7: Visualization of year distribution in READOC. (a) Document publication years of READOC-arXiv. (b) Document publication years of READOC-GitHub.

Parameter	Value
top_p	0.8
top_k	100
temperature	0.7
do_sample	True
repetition_penalty	1.05

Table 11: Generation parameters of Vision-Language Models.

D Details of Experiments

D.1 Implementation of VLMs

This section primarily supplements the evaluation process of Vision-Language Models. Specifically, we adopt a one-page image-to-Markdown approach, concatenating the Markdown text generated for each page. The prompts we use are illustrated in Figure 12, and we uniformly apply the generation parameters for all VLMs as shown in Table 11.

D.2 Additional Case Study

We provide two additional case analyses, as shown in Figures 13 and 14. It is evident that none of the DSE systems demonstrated satisfactory performance, as they exhibited various types of errors, including Missing Content, Misidentifying Headings, Inaccurate Text Extraction, Confused Reading Order, Hallucinations, and Repetitions.

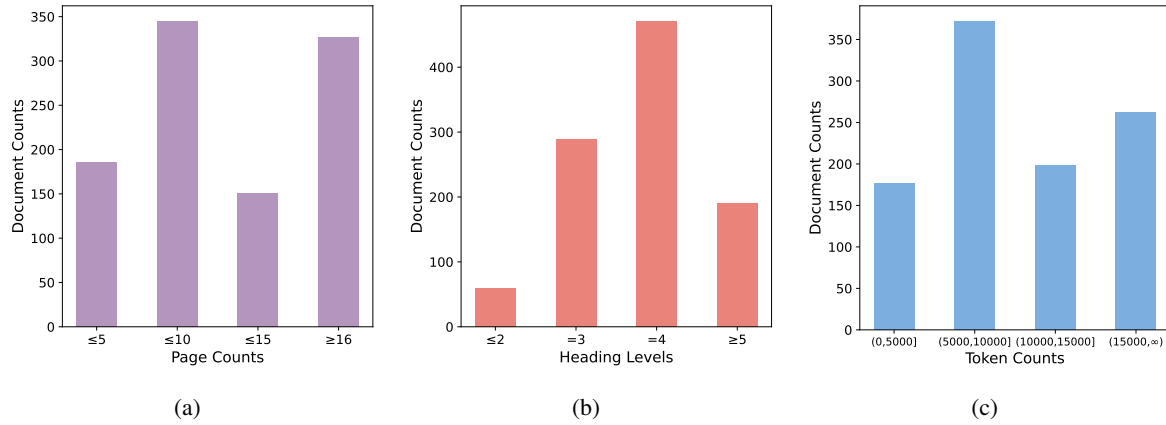


Figure 8: Visualization of year distribution in READOC-arXiv. (a) Document page counts of READOC-arXiv. (b) Document depth (heading levels) of READOC-arXiv. (c) Document Length (token counts) of READOC-arXiv.

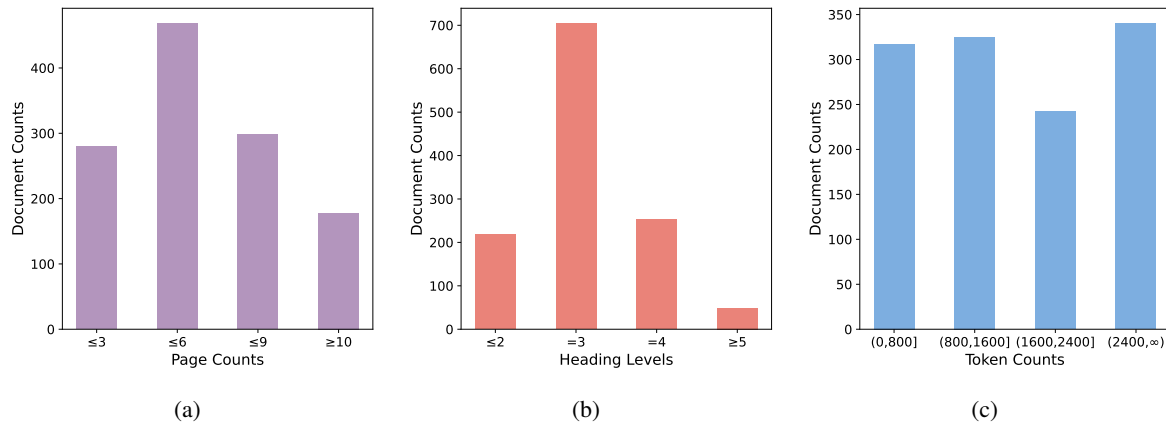


Figure 9: Visualization of year distribution in READOC-GitHub. (a) Document page counts of READOC-GitHub. (b) Document depth (heading levels) of READOC-GitHub. (c) Document Length (token counts) of READOC-GitHub.




Figure 10: Document types of READOC-arXiv.

Headings
<code>re.compile(r'^((#(1,6)) +(.+))', re.MULTILINE)</code>
Embedded Formulas
<code>re.compile(r'"\\ (((.*?)(?<!\ \\))\\ ")')</code>
Isolated Formulas
<code>re.compile(r'"\\ (((.*?)(?<!\ \\))\\ ")', re.S)</code>
Tables
<code>re.compile(r'\n\\begin\\{table\\}((?:!\\begin\\{table\\})\\.)*?\\end\\{table\\}', re.S)</code>
Plain Text
(The remaining part is treated as plain text)

Figure 11: Segmentation module's regular expressions.

Prompt for READoc-arXiv

This image displays a document page, convert the page content into Markdown format. Use continuous # to denote headings at each level. Display tables following markdown or latex format. Use \$ or $\backslash()$ to surround inline math, and use \$\$ or $\backslash[]$ to surround isolated math block. Don't explain, directly output the Markdown-format content.

Prompt for READoc-GitHub

This image displays a document page, convert the page content into Markdown format. Use continuous # to denote headings at each level. If it's a blank page, don't output anything. Otherwise, directly output the Markdown-format content without explanation.

Figure 12: Prompts of Vision-Language Models.

