

Multimodal Large Language Models for Text-rich Image Understanding: A Comprehensive Review

Anonymous ACL submission

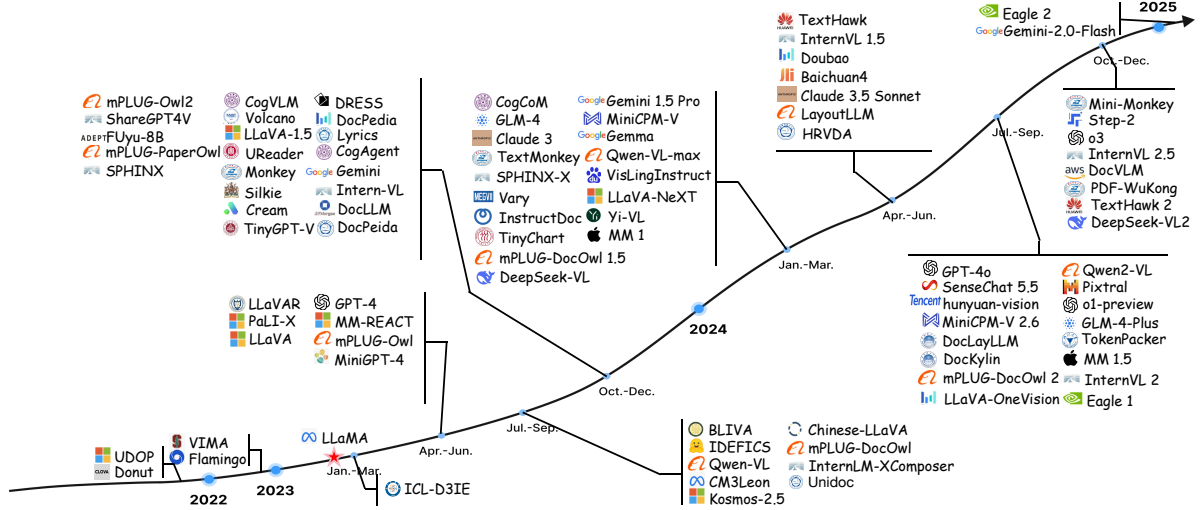


Figure 1: The development timeline of TIU MLLMs.

Abstract

The recent emergence of Multi-modal Large Language Models (MLLMs) has introduced a new dimension to the Text-rich Image Understanding (TIU) field, with models demonstrating impressive and inspiring performance. However, their rapid evolution and widespread adoption have made it increasingly challenging to keep up with the latest advancements. To address this, we present a systematic and comprehensive survey to facilitate further research on TIU MLLMs. Initially, we outline the timeline, architecture, and pipeline of nearly all TIU MLLMs. Then, we review the performance of selected models on mainstream benchmarks. Finally, we explore promising directions, challenges, and limitations within the field.

1 Introduction

Text-rich images play a pivotal role in real-world scenarios by efficiently conveying complex information and improving accessibility (Biten et al., 2019). Accurately interpreting these images is essential for automating information extraction, advancing AI systems, and optimizing user interactions. To formalize this research domain, we term

it **Text-rich Image Understanding (TIU)**, which encompasses two core capabilities: perception and understanding. The perception dimension focuses on visual recognition tasks, such as text detection (Liao et al., 2022), formula recognition (Truong et al., 2024), and document layout analysis (Yupan et al., 2022). The understanding dimension, conversely, requires semantic reasoning for applications like key information extraction and document-based visual question answering (*e.g.*, DocVQA (Mathew et al., 2021b), ChartQA (Masry et al., 2022), and TextVQA (Singh et al., 2019)).

Historically, perception and understanding tasks were handled separately through specialized models or multi-stage pipelines. Recent advances in vision-language models have unified these tasks within Visual Question Answering (VQA) paradigms, driving research towards the development of end-to-end universal models.

Figure 1 presents an evolutionary timeline delineating critical milestones in unified text-rich image understanding models. The trajectory reveals two distinct eras: (a) The pre-LLM period (2019-2022) characterized by specialized architectures

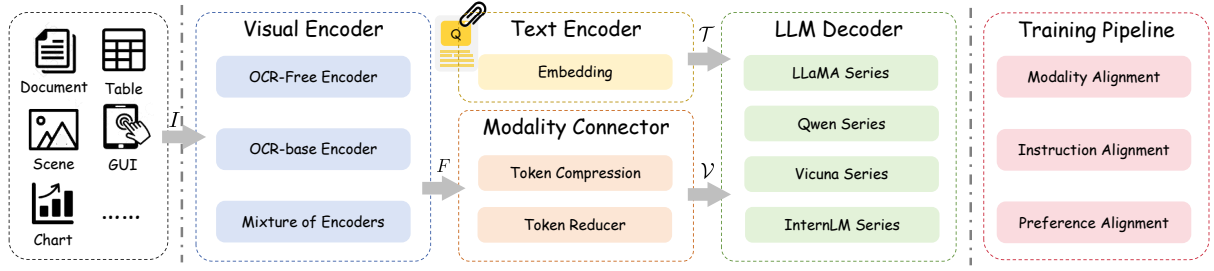


Figure 2: The general model architecture of MLLMs and the implementation choices for each component.

like LayoutLM (Xu et al., 2019) and Donut (Kim et al., 2021), which employed modality-specific pre-training objectives (masked language modeling, masked image modeling, *etc.*) coupled with OCR-derived supervision (text recognition, spatial order recovery, *etc.*). While effective in controlled settings, these models exhibited limited adaptability to open-domain scenarios due to their task-specific fine-tuning requirements and constrained cross-modal interaction mechanisms. (b) The post-LLM era (2023–present) is marked by the growing popularity of LLMs. Some studies propose Multi-modal Large Language Models (MLLMs), which integrate LLM with visual encoders to jointly process visual tokens and linguistic elements through unified attention mechanisms, achieving end-to-end sequence modeling.

This paradigm evolution addresses two critical limitations of earlier methods. First, the emergent MLLM framework eliminates modality-specific inductive biases through homogeneous token representation, enabling seamless multi-task integration. Second, the linguistic priors encoded in LLMs empower unprecedented zero-shot generalization and allow direct application to diverse tasks without task-specific tuning.

Although these MLLMs present impressive and inspiring results, their rapid evolution and broad adoption have made tracking cutting-edge advancements increasingly challenging. Therefore, a systematic review that is tailored for documents to summarize and analyse these methods is in demand. However, existing surveys on text-rich image understanding often exhibit narrow focus: they either analyze domain-specific scenarios (e.g., tables and figures (Huang et al., 2024a), charts (Huang et al., 2024b; Al-Shetairy et al., 2024), forms (Abdallah et al., 2024)) or emphasize unified deep learning frameworks (Subramani et al.; Ding et al., 2024).

Our systematic survey addresses the gap by providing the first comprehensive analysis of nearly all TIU MLLMs in four dimensions: Model Archi-

tectures (Section 2), Training Pipeline (Section 3), Datasets and Benchmarks (Section 4), Challenges and Trends (Section 5). This holds both academic and practical significance for advancing the field.

2 Model Architecture

TIU MLLM methods typically leverage pre-trained general visual foundation models to extract robust visual features or employ OCR engines to capture text and layout information from images. A modality connector is then used to align these visual features with the semantic space of the language features from the LLM. Finally, the combined visual-language features are fed into the LLM, which utilizes its powerful comprehension capabilities for semantic reasoning to generate the final answer. As illustrated in Figure 2, the framework of TIU MLLMs can be abstracted into three core components: Visual Encoder, Modality Connector, and LLM Decoder.

2.1 Visual Encoder

The Visual Encoder $\mathcal{F}(\cdot)$ is responsible for transforming input image I into feature representations V , expressed as $V = \mathcal{F}(\cdot)(I)$. As illustrated in Figure 3, these encoders can be categorized into OCR-free, OCR-based, or a hybrid approach.

OCR-free Encoder is widely used to extract high-level visual features, effectively capturing essential information about objects, scenes, and textures. The commonly used OCR-free encoders include (1) **CLIP** (Radford et al., 2021); (2) **ConvNeXt** (Woo et al., 2023); (3) **SAM** (Kirillov et al., 2023); (4) **DINOv2** (Oquab et al., 2023); (5) **Swin-T** (Liu et al., 2021); (6) **InternViT** (Chen et al., 2024d).

OCR-based Encoder processes textual content and layout information from OCR outputs through three primary paradigms: (1) **Direct Input** injects raw OCR texts into LLMs, though long sequences degrade inference efficiency (He et al., 2023b); (2) **Cross-Attention** dynamically selects salient content via attention mechanisms within LLMs

Model	Visual Encoder	Modality Connector	LLM Decoder	Training Pipeline	DocVQA	InfoVQA	ChartQA	TextVQA	Avg.
UReader (Ye et al., 2023b)	CLIP-ViT-L/14	Cross Attention	LLaMA-7B	MA+IA	65.4	42.2	59.3	57.6	56.13
DocLLM-1B (Wang et al., 2023)	-	-	Falcon-1B	MA+IA	61.4	-	-	-	-
DocLLM-7B (Wang et al., 2023)	-	-	LLaMA2-7B	MA+IA	69.5	-	-	-	-
Cream (Kim et al., 2023)	CLIP-ViT-L/14	Cross Attention	Vicuna-7B	MA+IA	79.5	43.5	63.0	-	-
LLaVA-13B (Liu et al., 2023c)	CLIP-ViT-L/14	MLP	Vicuna-13B	MA+IA	6.9	-	-	36.7	-
PaLI-X (Chen et al., 2023)	ViT-22B	MLP	UL2-32B	MA+IA	86.8	54.8	72.3	80.8	73.68
LLaVAR (Zhang et al., 2023)	CLIP-ViT-L/14	MLP	Vicuna-13B	MA+IA	11.6	-	-	48.5	-
Qwen-VL (Bai et al., 2023b)	ViT-bigG	Cross Attention	Qwen-7B	MA+IA	65.1	35.4	65.7	63.8	57.50
LLaVA-1.5-7B (Liu et al., 2023b)	CLIP-ViT-L	MLP	Vicuna1.5-7B	MA+IA	-	-	-	58.2	-
LLaVA-1.5-13B (Liu et al., 2023b)	CLIP-ViT-L	MLP	Vicuna1.5-13B	MA+IA	-	-	-	62.5	-
CogAgent (Hong et al., 2023)	EVA2-CLIP+CogVLM	MLP+Cross Attention	Vicuna-13B	MA+IA	81.6	44.5	68.4	76.1	67.65
Unidoc (Feng et al., 2023)	CLIP-ViT-L/14	MLP	Vicuna-13B	MA+IA	90.2	36.8	70.5	73.7	67.80
Monkey (Li et al., 2024e)	Vit-BigG	Cross Attention	Qwen-7B	MA+IA	66.5	36.1	65.1	67.6	58.83
Mini-Monkey (Huang et al., 2024c)	InternViT-300M	MLP	InternLM2-2B	IA	87.4	60.1	76.5	75.7	74.93
TextMonkey (Liu et al., 2024e)	Vit-BigG	Cross Attention	Qwen-7B	MA+IA	73.0	-	66.9	65.6	-
IDEFICS2 ((Laurençon et al., 2024))	SigLIP-SO400M	Cross Attention	Mistral-7B	MA+IA	74.0	-	-	73.0	-
LayoutLLM (Luo et al., 2024b)	LayoutLMv3-large	MLP	Vicuna1.5-7B	MA+IA	74.25	-	-	-	-
DocKylin (Zhang et al., 2024b)	Swin	MLP	Qwen-7B	MA+IA	77.3	46.6	66.8	-	-
DocLayLLM (Liao et al., 2024b)	LayoutLMV3	MLP	LLaMA3-8B	MA+IA	77.79	42.02	-	-	-
mPLUG-DocOwl (Hu et al., 2024a)	CLIP-ViT-L/14	Cross Attention	LLaMA-7B	MA+IA	62.2	38.2	57.4	52.6	52.60
mPLUG-DocOwl1.5 (Hu et al., 2024b)	CLIP-ViT-L/14	H-Reducer	LLaMA2-7B	MA+IA	82.2	50.7	70.2	68.6	67.93
mPLUG-DocOwl2 (Hu et al., 2024d)	CLIP-ViT-L/14	H-Reducer	LLaMA2-7B	MA+IA	80.7	46.4	70.0	66.7	65.95
Vary (Wei et al., 2024)	CLIP-ViT-L/14 + SAM	MLP	Qwen-7B	MA+IA	76.3	-	66.1	-	-
Eagle (Shi et al., 2024)	CLIP + ConvNeX + Pix2Struct + EVA2 + SAM	MLP	LLaMA3-8B	MA+IA	86.6	-	80.1	77.1	-
PDF-WuKong (Xie et al., 2024)	CLIP-ViT-L-14	Cross Attention	InternLM2-7B	MA+IA	85.1	61.3	80.0	-	-
InstructDoc (Tanaka et al., 2024b)	CLIP/Eva-CLIP-ViT	Cross Attention + MLP	Flan-T5/OPT	MA+IA	-	50.9	29.4	53.8	-
TextHawk (Yu et al., 2024a)	SigLIP	Cross Attention	InternLM-XComposer	MA+IA	76.4	50.6	66.6	-	-
TextHawk2 (Yu et al., 2024b)	SigLIP	Cross Attention	Qwen2-7B	MA+IA	89.6	67.8	81.4	75.1	78.48
MM1.5-1B (Zhang et al., 2024a)	CLIP-ViT-H	C-Abstractor	Private	MA+IA	81.0	50.5	67.2	72.5	67.80
MM1.5-3B (Zhang et al., 2024a)	CLIP-ViT-H	C-Abstractor	Private	MA+IA	87.7	58.5	74.2	76.5	74.23
MM1.5-7B (Zhang et al., 2024a)	CLIP-ViT-H	C-Abstractor	Private	MA+IA	88.1	59.5	78.6	76.5	75.68
MM1.5-30B (Zhang et al., 2024a)	CLIP-ViT-H	C-Abstractor	Private	MA+IA	91.4	67.3	83.6	79.2	80.38
HRVDA (Liu et al., 2024a)	Swin-L	MLP	LLaMA2-7B	MA+IA	72.1	43.5	67.6	73.3	64.13
InternVL1.5-26B (Chen et al., 2024c)	InternViT-6B	MLP	InternLM2-20B	MA+IA	90.9	72.5	83.8	80.6	81.95
InternVL2.5-1B (Chen et al., 2024b)	InternViT-300M	MLP	Qwen2.5-0.5B	MA+IA	84.8	56.0	75.9	72.0	72.18
InternVL2.5-2B (Chen et al., 2024b)	InternViT-300M	MLP	InternLM2.5-1.8B	MA+IA	88.7	60.9	79.2	74.3	75.78
InternVL2.5-4B (Chen et al., 2024b)	InternViT-300M	MLP	Qwen2.5-3B	MA+IA	91.6	72.1	84.0	76.8	81.13
InternVL2.5-8B (Chen et al., 2024b)	InternViT-300M	MLP	InternLM2.5-7B	MA+IA	93.0	77.6	84.8	79.1	83.63
InternVL2.5-26B (Chen et al., 2024b)	InternViT-6B	MLP	InternLM2.5-20B	MA+IA	94.0	79.8	87.2	82.4	85.85
InternVL2.5-38B (Chen et al., 2024b)	InternViT-6B	MLP	Qwen2.5-32B	MA+IA	95.3	83.6	88.2	82.7	87.45
InternVL2.5-78B (Chen et al., 2024b)	InternViT-6B	MLP	Qwen2.5-72B	MA+IA	95.1	84.1	88.3	83.4	87.73
InternVL2.5-8B-npo (Wang et al., 2024c)†	InternViT-300M	MLP	InternLM2.5-7B	PA	92.3	76.0	83.8	79.1	82.80
DocPeida (Feng et al., 2024)	Swin	MLP	Vicuna-7B	MA+IA	47.1	15.2	46.9	60.2	42.35
TinyChart (Zhang et al., 2024d)	SigLIP	MLP	Phi-2	IA	-	-	83.6	-	-
TokenPacker-7B (Li et al., 2024d)	CLIP-ViT-L/14	Cross Attention	Vicuna-7B	MA+IA	60.2	-	-	-	-
TokenPacker-13B (Li et al., 2024d)	CLIP-ViT-G/14	Cross Attention	Vicuna-13B	MA+IA	70.0	-	-	-	-
LLaVA-OneVision-0.5B (Li et al., 2024a)	SigLIP	MLP	qwen2-0.5B	MA+IA	70.0	41.8	61.4	-	-
LLaVA-OneVision-7B (Li et al., 2024a)	SigLIP	MLP	qwen2-7B	MA+IA	87.5	68.8	80.0	-	-
Qwen2-VL-2B (Wang et al., 2024b)	CLIP-ViT-G/14	Cross Attention	Qwen2-2B	MA+IA	90.1	65.5	73.5	79.7	77.20
Qwen2-VL-7B (Wang et al., 2024b)	CLIP-ViT-G/14	Cross Attention	Qwen2-7B	MA+IA	94.5	76.5	83.0	84.3	84.58
DocVLM (Nacson et al., 2024)	CLIP-ViT-G/14 + DocFormerV2	Cross Attention	Qwen2-7B	MA+IA	92.8	66.8	-	82.8	-
Qwen2-VL-72B (Wang et al., 2024b)	CLIP-ViT-G/14	Cross Attention	Qwen2-72B	MA+IA	96.5	84.5	88.3	85.5	88.70
DeepSeek-VL2-3B (Wu et al., 2024)	SigLIP-SO400M-384	Pixel-shuffle + MLP	DeepSeekMoE	MA+IA	88.9	66.1	81.0	80.7	79.18
DeepSeek-VL2-16B (Wu et al., 2024)	SigLIP-SO400M-384	Pixel-shuffle + MLP	DeepSeekMoE	MA+IA	92.3	75.8	84.5	83.4	84.00
DeepSeek-VL2-27B (Wu et al., 2024)	SigLIP-SO400M-384	Pixel-shuffle + MLP	DeepSeekMoE	MA+IA	93.3	78.1	86.0	84.2	85.40
Eagle2 (Li et al., 2025)	SigLIP + ConvNeXt	MLP	Qwen2.5-7B	MA+IA	92.6	77.2	86.4	83.0	84.80

Table 1: The summary of representative mainstream MLLMs, including the model architectures, training pipelines, and scores on the four most popular benchmarks of TIU. “Private” indicates that the MLLM utilizes a proprietary large model. “†” indicates the results are obtained by downloading official open-source model and testing it locally.

et al., 2024c) identifies important tokens and removes redundant ones. To evaluate the redundancy of image features, the similarity between image tokens is often utilized (Liu et al., 2024e). This method selects tokens that are highly unique and lack closely similar counterparts. Average pooling is the most special one.

2.3 LLM Decoder

The aligned features are fed into the LLM decoder together with the language embeddings for reasoning. We list the commonly used LLMs in MLLM: **LLaMA Series.** LLaMA (Touvron et al., 2023a,b; Dubey et al., 2024) is a series of open-source large language models developed by Meta, aimed at promoting openness and innovation in artificial intelligence technology, LLaMA series include models of varying parameter scales (e.g., 7B, 13B, 34B).

Qwen Series. Qwen (Bai et al., 2023a; Yang et al., 2024a), developed by Alibaba, is a multilingual LLM that supports both Chinese and English.

Vicuna Series. Vicuna (Zheng et al., 2023) is an open-source large language model built on LLaMA, developed by research teams from institutions including UC Berkeley, CMU, and Stanford.

InternLM series. InternLM (Cai et al., 2024) is an open-source large language model series developed by the Shanghai Artificial Intelligence Laboratory, with the latest version, InternLM 2.5, offering parameter sizes of 1.8B, 7B, and 20B.

3 Training Pipeline

The training pipeline of MLLM for TIU can be delineated into three main stages: 1) Modality Alignment (MA); 2) Instruction Alignment (IA); and 3) Preference Alignment (PA).

3.1 Modality Alignment

In this stage, previous works typically use OCR data to pre-train the MLLM, which aims to bridge the modality gap. The general alignment methods can be categorized into three types: recognition, localization, and parsing.

Read Full Text. UReader (Ye et al., 2023b) is the first to explore unified document-level understanding, which introduces the Read Full Text task in VQA for pre-training. Specifically, they include 1) reading all texts from top to bottom and left to right, and 2) reading the remaining texts based on given texts. Compared to reading the full text, some works (Lv et al., 2024; Hu et al., 2024a) propose a more structured reading approach by predicting the image markdown, not text transcriptions.

Reading Partial Text within Localization. Due to the length of document texts, instructions for reading the full text may risk truncation because of the limited token length in LLMs. To address these limitations, Park et al. (Park et al., 2024) introduced two novel tasks: Reading Partial Text (RPT) and Predicting Text Position (PTP). The former randomly selects and reads continuous portions of text in the reading order from top to bottom and left to right. For example, “Q: What is the text in the image between the first 30%, from 20% to 40%, or the last 16%?” For the PTP task, given a text segment, the MLLM aims to infer its relative position (percentage format) within the full text. For example, “Q: Where is the text query texts located within the image? A: 40% to 80%”. However, this approach can be somewhat obscure and challenging to express accurately.

Alternatively, some methods (Hu et al., 2024b; Yu et al., 2024a; Liu et al., 2024a) extract texts based on specific spatial positions, which are summarized into two types. 1) Text Recognition aims to extract the textual content from a given position in the image, ensuring that the model can accurately recognize and extract text within specific regions. 2) Text Grounding involves identifying the corresponding bounding box for specific text in the image, which assists the model in understanding the document layout.

Parsing. In document images, many elements (charts, formulas, and tables) may not be represented using plain text. An increasing number of researchers are now focusing on these element parsing. 1) Chart Parsing. Chart types include vertical bars, horizontal bars, lines, scatter plots, and pie

charts. Charts serve as visual representations of tables, and organizing text in reading order fails to capture their structure. To preserve their mathematical properties, researchers often convert charts into tables. This process involves breaking down the chart into x/y axes and their corresponding values, which can be represented in Markdown, CSV formats, or even converted into Python code. This approach enables models to better understand the chart’s specific meaning.

2) Table Parsing. Compared to charts, tables have a more standardized structure, where rows and columns form key-value pairs. Common formats for representing tables include LaTeX, Markdown, and HTML. Markdown is often used for simple tables due to its concise text format, while HTML can handle cells that span multiple rows and columns, despite its use of many paired tags like `<tr></tr>` and `<td></td>`. Some tables, with complex spanning, custom lines, spacing, or multi-page length, require LaTeX for representation. However, the diversity in LaTeX representations can make these tables challenging for models to fully understand.

3) Formula Parsing. Besides tables and charts, formulas are also commonly used. In the pre-training phase, models learn the LaTeX representation of formula images, enhancing their understanding of formulas. This provides a solid foundation for tasks involving formula computation and reasoning during the instruction alignment.

3.2 Instruction Alignment

Upon completing the modality alignment pre-training stage, the MLLM acquires basic visual recognition and dialogue capabilities. However, to achieve human-aligned intelligence, three critical capability gaps must be addressed: (1) Advanced multimodal perception and cross-modal reasoning abilities; (2) Prompt robustness across diverse formulations; (3) Zero-shot generalization for unseen task scenarios. To bridge these gaps, instruction alignment through supervised fine-tuning (SFT) has emerged as an effective paradigm. This phase typically unfreezes all model parameters and employs instructional data with structured templates.

To systematically address these challenges, we have categorized the current methods emerging in instruction alignment into three distinct levels:

1) Level 1: Visual-Semantic Anchoring. We categorize these instructions into two types: i) Answer within the image; and ii) Answer without the image. This type of instruction data where answers are lo-

cated directly within the image, assists MLLMs refine their accuracy in generating responses that are directly linked to specific visual content, reducing reliance on generic or contextually weak answers (Mathew et al., 2021b, 2022). Certain tasks require reasoning based on world knowledge and involve complex inference procedures, such as scientific question answering (Masry et al., 2022; Chen et al., 2021). Consequently, these instructions are designed with the common characteristic that the answer is not directly visible in the image. This encourages the model to utilize its linguistic comprehension and external knowledge, enhancing its advanced reasoning and inference capabilities. An example might be: “Q: How much higher is the red bar compared to the yellow bar in the chart, in terms of percentage? A: 12.1%.”

2) Level 2: Prompt Diversity Augmentation. To bolster robustness in handling a broader spectrum of prompts, rather than being limited to specific prompts tailored for particular tasks, researchers often employ data augmentation on the question component of the instruction stream. A popular strategy involves leveraging existing large language models to rephrase the same question in multiple ways. For example, consider the original question: “What is written on the sign in the image?” It can be rephrased as: “Can you read the text displayed on the sign shown in the image?” “Identify the sign in the image.” “Please examine the image and list the words that appear within the sign.” By utilizing such varied templates, researchers can train MLLMs to better interpret and respond to a wide range of prompts, thereby enhancing their flexibility and accuracy in real-world applications.

3) Level 3: Zero-shot Generalization. To enhance the generalization ability to handle unseen tasks, several strategies typically are employed:

Chain of Thought (CoT) (Wei et al., 2022) reasoning involves breaking down complex problems into a series of intermediate steps or sub-tasks, allowing a model to tackle each part systematically. Some studies have demonstrated improvements by incorporating text-level CoT reasoning (Zhang et al., 2024c) or box-level visual CoT supervision (Shao et al., 2025). To better illustrate the process, consider the prompt: “What is the average of the last four countries’ data?”, the CoT reasoning unfolds as follows: i) Identify the data for the last four countries; ii) Calculate the sum of these values; iii) Calculate the average by dividing the sum by the number of countries.

Another strategy is Retrieval-Augmented Generation (RAG). RAG (Arslan et al., 2024) combines the strengths of retrieval-based and generation-based approaches by integrating an information retrieval component with a generative model. This method allows the model to access a vast external knowledge base, retrieving pertinent information to inform and enhance the generation process.

3.3 Preference Alignment

In the modality and instruction alignment stages, the model predicts the next token based on previous ground-truth tokens during training, and on its own prior outputs during inference. If errors occur in the outputs, this can lead to a distribution shift in inference. The more output the model has, the more serious this phenomenon becomes. In previous natural language processing (NLP) works (Lai et al., 2024; Pang et al., 2025), a series of preference alignment techniques (Rafailov et al., 2024; Ouyang et al., 2022; Shao et al., 2024; Wang et al., 2024a) have been proposed to optimize the output of the model to make it more consistent with human values and expectations. Benefiting from the success of preference alignment applied to NLP, InternVL2-MPO (Wang et al., 2024d) introduces preference alignment to the multimodal field and proposes a Mixed Preference Optimization (MPO) to improve multimodal reasoning. Specifically, they propose a continuation-based Dropout Next Token Prediction (DropoutNTP) pipeline for samples lacking clear ground truth and a correctness-based pipeline for samples with clear ground truth. This strategy improves the performance of the model on OCRBench (Fu et al., 2024). Nevertheless, its potential to enhance document multimodal reasoning remains under-explored.

4 Datasets and Benchmarks

The rapid advancements in TIU tasks have been fundamentally driven by the proliferation of specialized datasets and standardized benchmarks. As illustrated in Table 2, we systematically categorize TIU-related datasets into two types: *domain-specific* (Document, Chart, Scene, Table, and GUI) and *comprehensive* scenarios.

Specifically, some datasets are derived by converting training data from traditional tasks into Visual Question Answering (VQA) formats, such as text detection, text spotting, table recognition, and *etc.*. These datasets are typically utilized for modal-

Domain	Dataset	Language	Scene Sources	#Images	#Q&A pairs	Train/Test
Document	DocVQA (Mathew et al., 2021b)	English	Industry document	12,767	50,000	Train + Test
	Docmatix (Laurençon et al., 2024)	English	Industry document	2.4M	9.5M	Train
	InfoVQA (Mathew et al., 2022)	English	Infographics	5,485	30,035	Train + Test
	MP-DocVQA (Tito et al., 2023)	English	Industry documents	47,952	46,176	Train + Test
	DocGenome (Xia et al., 2024a)	English	Scientific document	6.8M	3,000	Train
	IIT-CDIP (Xu et al., 2020)	English	Multi-domain	11M	-	Train
	synthdog (Kim et al., 2022)	English	Multi-domain	2M	-	Train
	CCPdf (Turski et al., 2023)	Multilingual	Multi-domain	1.1M	-	Train
	RVL-CDIP (Harley et al., 2015)	English	Industry document	159,418	-	Train
	VisualMRC (Tanaka et al., 2021)	English	Webpage Document	10,197	30,562	Train + Test
	KLC (Stanislawek et al., 2021)	English	Industry document	2463	22,224	Train + Test
	OCREval (Lv et al., 2023)	English	Multi-domain	2,297	-	Test
	MMLongBench-Doc (Ma et al., 2024)	English	Multi-domain Long Documents	135	1k	Test
	Do-GOOD (He et al., 2023a)	English	Industry document	410k	50k	Test
	OCR-VQA (Mishra et al., 2019)	English	Book covers	207,572	>1M	Train + Test
	SlideVQA (Tanaka et al., 2023)	English	Slide decks	52,480	14,484	Train + Test
	PDF-VQA (Ding et al., 2023)	English	Scientific document	13,484	140,610	Train + Test
	BenthamQA (Mathew et al., 2021a)	English	Handwritten document	338	200	Train + Test
	FinanceQA (Sujet AI, 2024)	English	Financial reports	9,801	100k	-
	Ureader (Ye et al., 2023a)	English	Multi-domain	24.5k	24.5k	Train
	ColPali (Faysse et al., 2024)	English	Multi-domain	118,695	118,695	Train + Test
	FUNSD (Jaume et al., 2019)	English	Scanned forms	199	5312	Train + Test
	SROIE (Huang et al., 2019)	English	Multi-domain	973	52,316	Train + Test
	POIE (Kuang et al., 2023)	English	Multi-domain	3,000	111,155	Train + Test
	IAM (Marti and Bunke, 2002)	English	Lancaster-Oslo/Bergen	1066	-	Train
Chart	ChartQA (Masry et al., 2022)	English	Charts and Plots	20,882	32,719	Train + Test
	PlotQA (Methani et al., 2020)	English	Plots (Real world data source)	224,377	28.9M	Train
	FigureQA (Kahou et al., 2017)	English	Science style image	>100,000	>1.3M	Train
	DVQA (Kafle et al., 2018)	English	Data Visualizations	300,000	3,487,194	Train
	Unichart (Masry et al., 2023)	English	Multi-domain	290,736	300,000	Train
	LRV-Instruction (Liu et al., 2023a)	English	Multi-domain	400k	400k	Train + Test
	VisText (Tang et al., 2023)	English	Financial reports	12,441	12,441	Train + Test
	Chart2Text (Obeid and Hoque, 2020)	English	Financial reports	8,305	8,305	Train + Test
	ArxivQA (Li et al., 2024c)	English	Scientific Chart	35,000	100,000	Train + Test
	ChartY (Chen et al., 2024a)	Multilingual	Charts and Plots	6k	6k	Test
	ChartX (Xia et al., 2024b)	English	Charts and Plots	6k	6k	Test
	MMC (Liu et al., 2024c)	English	Plots (Real world data source)	1.7k	2.9k	Test
Scene	ChartBench (Xu et al., 2024)	English	Plots (Real world data source)	68k	549k	Test
	TextCaps (Sidorov et al., 2020)	English	Scene Text	28,408	142,040	Train + Test
	TextVQA (Singh et al., 2019)	English	Scene Text	28,408	45,336	Train + Test
	ST-VQA (Biten et al., 2019)	English	Scene Text	23,038	31,791	Train + Test
	MT-VQA (Wen et al., 2024)	Multilingual	20 fine-grained scenes	8,794	28,607	Train + Test
	OCRVQA (Mishra et al., 2019)	English	Scene Text	207,572	1M	Train
	ICDAR13 (Karatzas et al., 2013)	English	Scene Text	229	-	Train
	ICDAR15 (Karatzas et al., 2015)	English	Scene Text	1000	-	Train
	TotalText (Ch'ng and Chan, 2017)	English	Scene Text	1000	-	Train
	CTW1500 (Yuliang et al., 2017)	English	Scene Text	1255	-	Train
	LSVT (Sun et al., 2019)	Chinese	Scene Text	30,000	-	Train
	RCTW (Shi et al., 2017)	Chinese	Scene Text	8,034	-	Train
Table	LAION-OCR (Schuhmann et al., 2022)	English	Scene Text	-	-	Train
	Wukong-OCR (Gu et al., 2022)	Chinese	Scene Text	-	-	Train
	TableQA (Sun et al., 2020)	English	Financial reports	6,000	64,891	Train
	WikiTableQuestions (Pasupat et al., 2015)	English	Multi-domain	2,108	22,033	Train + Test
	DeepForm (Svetlichnaya, 2020)	English	Political campaign finance receipts	1100	5500	Train + Test
	TabFact (Chen et al., 2019)	English	Wikipedia tables	14,922	117,273	Train + Test
	TabMWP (Lu et al., 2022)	English	Educational documents	38,431	38,431	Train
	TURL (Deng et al., 2022)	English	Wikipedia	200,000	-	Train
	PubTabNet (Zhong et al., 2020)	English	Scientific articles	200,000	-	Train
	TableVQA-Bench (Kim et al., 2024)	English	Scientific and Financial Reports	0.9k	1.5k	Test
	MMTAB-eval (Zheng et al., 2024)	English	Scientific and Financial Reports	23k	49k	Test
	ComTQA (Zhao et al., 2024)	English	Scientific and Financial Reports	1.6k	9k	Test
GUI	TAT-DQA (Zhu et al., 2022)	English	Financial reports	3,067	16,558	Train + Test
	VQAonBD (Raja et al., 2023)	English	Financial reports	48,895	1,531,455	Train + Test
Comprehensive	MultiHiertt (Zhao et al., 2022)	English	Financial reports	89,646	10,440	Train + Test
	ScreenQA (Hsiao et al., 2022)	English	Mobile app screenshots	35,352	85,984	Train + Test
Comprehensive	Screen2Words (Wang et al., 2021)	English	Android app screenshot	22,417	112,085	Train + Test
	OCRBench (Liu et al., 2024d)	English	Multi-domain	0.9k	1k	Test
Comprehensive	Seed-bench-2-plus (Li et al., 2024b)	English	Multi-domain	0.6k	2.3k	Test
	CONTEXTUAL (Wadhawan et al., 2024)	English	Multi-domain	0.5k	0.5k	Test
	OCRBench v2 (Fu et al., 2024)	English	Multi-domain	9.5k	10k	Test
	FOX (Liu et al., 2024b)	Multilingual	Scientific document	0.7k	2.2k	Test
	DocLocal4K (Hu et al., 2024c)	English	Multi-domain	4.2k	4.2k	Test
	CC-OCR (Yang et al., 2024b)	Multilingual	Multi-domain	7k	-	Test
	MMDocBench (Zhu et al., 2024)	English	Multi-domain	2.4k	4.3k	Test

Table 2: Representative datasets and benchmarks for Text-rich Image Understanding. Each dataset is marked for training and testing typically according to its content, functions, and user requirements.

ity alignment in the first stage of training, enabling models to bridge the gap between textual and visual information effectively. Other datasets are specifically designed in VQA formats for certain scenarios, such as DocVQA (Mathew et al., 2021b), InfoVQA (Mathew et al., 2022), ChartQA (Masry

et al., 2022), and TextVQA (Singh et al., 2019). These datasets have played a pivotal role in advancing the field of TIU by providing structured and domain-specific challenges. Their introduction has significantly accelerated progress in tasks like document understanding, chart interpretation, and

natural scene text comprehension. Consequently, published papers frequently report these metrics, as they not only contribute to instruction alignment in the second stage of training but also serve as essential benchmarks for evaluating model performance.

In addition to training datasets, there is a distinct category of datasets that are exclusively designed for evaluating specific capabilities of MLLMs. Examples include TableVQA-Bench (Kim et al., 2024), ChartBench (Xu et al., 2024), and MMLongBench-Doc (Ma et al., 2024). These datasets are tailored to assess advanced functionalities such as long-context understanding, cross-modal reasoning, and domain-specific comprehension. By providing targeted evaluation frameworks, they enable researchers to identify strengths and weaknesses in MLLMs, driving further innovation and refinement in the field.

5 Challenges and Trends

As shown in Table 1, we calculated the average scores from four popular and widely used evaluation datasets, which can basically reflect the performance of MLLMs on TIU tasks. The top five models are Qwen2-VL-72B (88.70), InternVL2.5-78B (87.73), InternVL2.5-38B (87.45), InternVL2.5-26B (85.85), and DeepSeek-VL2-27B (85.40). This indicates that the most state-of-the-art (SOTA) MLLMs currently employ OCR-free encoders, which avoids redundant tokens and complex model architectures. Despite the promising and significant progress made by current MLLMs, the field still faces considerable challenges that require further research and innovation:

Computational Efficiency and Model Compression. The computational demands of current MLLMs remain a critical bottleneck, primarily due to two factors: (1) the necessity of processing high-resolution document images, which imposes substantial computational resource requirements, and (2) the prevalent use of 7-billion-parameter architectures, while delivering state-of-the-art performance, incur high deployment costs and latency. These challenges underscore the importance of developing more efficient MLLM architectures that balance performance with reduced computational overhead. Encouragingly, recent advancements, as illustrated in Table 1, demonstrate promising trends toward model miniaturization. For instance, Mini-monkey (Huang et al., 2024c) achieves performance comparable to 7B-parameter models on

multiple TIU tasks while utilizing only 2B parameters, highlighting the potential for lightweight yet powerful architectures.

Optimization of Visual Feature Representation. A persistent challenge in MLLMs is the disproportionate length of image tokens compared to text tokens, which significantly increases computational complexity and degrades inference efficiency. Addressing this issue requires innovative approaches to compress image tokens without sacrificing model performance. Promising directions include the development of efficient visual encoders, adaptive token compression mechanisms, and advanced techniques for cross-modal feature fusion. Crucially, these methods must preserve the semantic richness of document content during compression. As shown in Table 1, recent architectural innovations, such as mPLUG-DocOwl2’s (Hu et al., 2024d) visual token compression, have made strides in this direction by enabling the processing of larger input images while maintaining benchmark performance.

Long Document Understanding Capability. While MLLMs excel at single-page document understanding, their performance on multi-page or long-document tasks remains suboptimal. Key challenges include modeling long-range dependencies, maintaining contextual coherence across pages, and efficiently processing extended sequences. The emergence of specialized benchmarks for long-document understanding (Ma et al., 2024), as highlighted in Table 2, is expected to drive significant progress in this field by providing standardized evaluation frameworks and fostering targeted research efforts.

Multilingual Document Understanding. Current MLLMs are predominantly optimized for English and a limited set of high-resource languages, resulting in inadequate performance in multilingual and low-resource language scenarios. Addressing this limitation requires the development of comprehensive multilingual datasets that encompass diverse linguistic and cultural contexts. Recent initiatives, such as MT-VQA (Tang et al., 2024) and CC-OCR (Yang et al., 2024b) (referenced in Table 2), represent important steps forward by introducing TIU tasks specifically designed to evaluate multilingual capabilities. These efforts, coupled with advances in cross-lingual transfer learning, are expected to significantly enhance the inclusivity and applicability of MLLMs in global contexts.

6 Limitation

This paper provides a systematic review of multi-modal large language models (MLLMs) in the field of Text-rich Image Understanding (TIU). While the research team has conducted comprehensive retrieval and integration of core literature prior to the submission deadline, certain minor studies may still remain uncovered. It should be particularly noted that due to publisher formatting requirements, the exposition of existing technical approaches and benchmark datasets in this work maintains essential conciseness. For complete algorithmic implementation details and experimental parameter configurations, researchers are strongly recommended to consult the original publications.

References

Abdelrahman Abdallah, Daniel Eberharter, Zoe Pfister, and Adam Jatowt. 2024. Transformers and language models in form understanding: A comprehensive review of scanned document analysis. *arXiv preprint arXiv:2403.04080*.

Mirna Al-Shetairy, Hanan Hindy, Dina Khattab, and Mostafa M. Aref. 2024. [Transformers utilization in chart understanding: A review of recent advances future trends](#). *Preprint*, arXiv:2410.13883.

Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katie Millican, Malcolm Reynolds, Roman Ring, Eliza Rutherford, Serkan Cabi, Tengda Han, Zhitao Gong, Sina Samangooei, Marianne Monteiro, Jacob Menick, Sebastian Borgeaud, Andrew Brock, Aida Nematzadeh, Sahand Sharifzadeh, Mikolaj Binkowski, Ricardo Barreira, Oriol Vinyals, Andrew Zisserman, and Karen Simonyan. 2022. [Flamingo: a visual language model for few-shot learning](#). *Preprint*, arXiv:2204.14198.

Muhammad Arslan, Hussam Ghanem, Saba Munawar, and Christophe Cruz. 2024. A survey on rag with llms. *Procedia Computer Science*, 246:3781–3790.

Jinze Bai, Shuai Bai, Yunfei Chu, Zeyu Cui, Kai Dang, Xiaodong Deng, Yang Fan, Wenbin Ge, Yu Han, Fei Huang, et al. 2023a. Qwen technical report. *arXiv preprint arXiv:2309.16609*.

Jinze Bai, Shuai Bai, Shusheng Yang, Shijie Wang, Sinan Tan, Peng Wang, Junyang Lin, Chang Zhou, and Jingren Zhou. 2023b. Qwen-vl: A frontier large vision-language model with versatile abilities. *arXiv preprint arXiv:2308.12966*.

Ali Furkan Biten, Ruben Tito, Andres Mafra, Lluís Gomez, Marçal Rusinol, Ernest Valveny, CV Jawahar, and Dimosthenis Karatzas. 2019. Scene text visual question answering. In *Proceedings of the*

IEEE/CVF international conference on computer vision, pages 4291–4301.

Zheng Cai, Maosong Cao, Haojiong Chen, Kai Chen, Keyu Chen, Xin Chen, Xun Chen, Zehui Chen, Zhi Chen, Pei Chu, et al. 2024. Internlm2 technical report. *arXiv preprint arXiv:2403.17297*.

Junbum Cha, Wooyoung Kang, Jonghwan Mun, and Byungseok Roh. 2024. Honeybee: Locality-enhanced projector for multimodal llm. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13817–13827.

Jiaqi Chen, Jianheng Tang, Jinghui Qin, Xiaodan Liang, Lingbo Liu, Eric P Xing, and Liang Lin. 2021. Geoqa: A geometric question answering benchmark towards multimodal numerical reasoning. *arXiv preprint arXiv:2105.14517*.

Jinyue Chen, Lingyu Kong, Haoran Wei, Chenglong Liu, Zheng Ge, Liang Zhao, Jianjian Sun, Chunrui Han, and Xiangyu Zhang. 2024a. Onechart: Purify the chart structural extraction via one auxiliary token. In *Proceedings of the 32nd ACM International Conference on Multimedia*, pages 147–155.

Wenhu Chen, Hongmin Wang, Jianshu Chen, Yunkai Zhang, Hong Wang, Shiyang Li, Xiyu Zhou, and William Yang Wang. 2019. Tabfact: A large-scale dataset for table-based fact verification. *arXiv preprint arXiv:1909.02164*.

Xi Chen, Josip Djolonga, Piotr Padlewski, Basil Mustafa, Soravit Changpinyo, Jialin Wu, Carlos Riquelme Ruiz, Sebastian Goodman, Xiao Wang, Yi Tay, et al. 2023. Pali-x: On scaling up a multilingual vision and language model. *arXiv preprint arXiv:2305.18565*.

Zhe Chen, Weiyun Wang, Yue Cao, Yangzhou Liu, Zhangwei Gao, Erfei Cui, Jinguo Zhu, Shenglong Ye, Hao Tian, Zhaoyang Liu, et al. 2024b. Expanding performance boundaries of open-source multimodal models with model, data, and test-time scaling. *arXiv preprint arXiv:2412.05271*.

Zhe Chen, Weiyun Wang, Hao Tian, Shenglong Ye, Zhangwei Gao, Erfei Cui, Wenwen Tong, Kongzhi Hu, Jiapeng Luo, Zheng Ma, et al. 2024c. How far are we to gpt-4v? closing the gap to commercial multimodal models with open-source suites. *Science China Information Sciences*, 67(12):220101.

Zhe Chen, Jiannan Wu, Wenhai Wang, Weijie Su, Guo Chen, Sen Xing, Muyan Zhong, Qinglong Zhang, Xizhou Zhu, Lewei Lu, Bin Li, Ping Luo, Tong Lu, Yu Qiao, and Jifeng Dai. 2024d. [Internvl: Scaling up vision foundation models and aligning for generic visual-linguistic tasks](#). *Preprint*, arXiv:2312.14238.

Chee Kheng Ch'ng and Chee Seng Chan. 2017. Total-text: A comprehensive dataset for scene text detection and recognition. In *2017 14th IAPR international conference on document analysis and recognition (ICDAR)*, volume 1, pages 935–942. IEEE.

654	Wenliang Dai, Junnan Li, Dongxu Li, Anthony	Adam W Harley, Alex Ufkes, and Konstantinos G Der-	710
655	Meng Huat Tiong, Junqi Zhao, Weisheng Wang,	panis. 2015. Evaluation of deep convolutional nets	711
656	Boyang Li, Pascale Fung, and Steven Hoi.	for document image classification and retrieval. In	712
657	2023. Instructblip: Towards general-purpose vision-	<i>2015 13th International Conference on Document</i>	713
658	language models with instruction tuning . <i>Preprint</i> ,	<i>Analysis and Recognition (ICDAR)</i> , pages 991–995.	714
659	arXiv:2305.06500.	IEEE.	715
660	Xiang Deng, Huan Sun, Alyssa Lees, You Wu, and Cong	Jiabang He, Yi Hu, Lei Wang, Xing Xu, Ning Liu, Hui	716
661	Yu. 2022. Turl: Table understanding through repre-	Liu, and Heng Tao Shen. 2023a. Do-GOOD: towards	717
662	sensation learning. <i>ACM SIGMOD Record</i> , 51(1):33–	distribution shift evaluation for pre-trained visual doc-	718
663	40.	ument understanding models. <i>arXiv</i> , 2306.02623.	719
664	Yihao Ding, Jean Lee, and Soyeon Caren Han. 2024.	Jiabang He, Lei Wang, Yi Hu, Ning Liu, Hui Liu, Xing	720
665	Deep learning based visually rich document con-	Xu, and Heng Tao Shen. 2023b. Icl-d3ie: In-context	721
666	tent understanding: A survey. <i>arXiv preprint</i>	learning with diverse demonstrations updating for	722
667	<i>arXiv:2408.01287</i> .	document information extraction. In <i>Proceedings</i>	723
668	Yihao Ding, Siwen Luo, Hyunsuk Chung, and	<i>of the IEEE/CVF International Conference on Com-</i>	724
669	Soyeon Caren Han. 2023. Pdf-vqa: A new dataset	<i>puter Vision</i> , pages 19485–19494.	725
670	for real-world vqa on pdf documents. In <i>Machine</i>	Wenyi Hong, Weihan Wang, Qingsong Lv, Jiazheng	726
671	<i>Learning and Knowledge Discovery in Databases:</i>	Xu, Wenmeng Yu, Junhui Ji, Yan Wang, Zihan Wang,	727
672	<i>Applied Data Science and Demo Track</i> , pages 585–	Yuxiao Dong, Ming Ding, and Jie Tang. 2023. Co-	728
673	601, Cham. Springer Nature Switzerland.	gagent: A visual language model for gui agents .	729
674	Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey,	<i>Preprint</i> , arXiv:2312.08914.	730
675	Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman,	Yu-Chung Hsiao, Fedir Zubach, Gilles Baechler, Victor	731
676	Akhil Mathur, Alan Schelten, Amy Yang, Angela	Carbune, Jason Lin, Maria Wang, Srinivas Sunkara,	732
677	Fan, et al. 2024. The llama 3 herd of models. <i>arXiv</i>	Yun Zhu, and Jindong Chen. 2022. Screenqa: Large-	733
678	<i>preprint arXiv:2407.21783</i> .	scale question-answer pairs over mobile app screen-	734
679	Manuel Faysse, Hugues Sibille, Tony Wu, Bilel Om-	shots. <i>arXiv preprint arXiv:2209.08199</i> .	735
680	rani, Gautier Viaud, Céline Hudelot, and Pierre	Anwen Hu, Haiyang Xu, Jiabo Ye, Ming Yan, Liang	736
681	Colombo. 2024. Colpali: Efficient document re-	Zhang, Bo Zhang, Chen Li, Ji Zhang, Qin Jin, Fei	737
682	trieval with vision language models. <i>arXiv preprint</i>	Huang, and Jingren Zhou. 2024a. mPLUG-DocOwl	738
683	<i>arXiv:2407.01449</i> .	1.5: unified structure learning for OCR-free document	739
684	Hao Feng, Qi Liu, Hao Liu, Jingqun Tang, Wengang	understanding. <i>arXiv</i> , 2403.12895.	740
685	Zhou, Houqiang Li, and Can Huang. 2024. DocPe-	Anwen Hu, Haiyang Xu, Jiabo Ye, Ming Yan, Liang	741
686	dia: unleashing the power of large multimodal model	Zhang, Bo Zhang, Chen Li, Ji Zhang, Qin Jin, Fei	742
687	in the frequency domain for versatile document un-	Huang, et al. 2024b. mplug-docowl 1.5: Unified	743
688	derstanding. <i>arXiv</i> , 2311.11810.	structure learning for ocr-free document understand-	744
689	Hao Feng, Zijian Wang, Jingqun Tang, Jinghui Lu, Wen-	ing. <i>arXiv preprint arXiv:2403.12895</i> .	745
690	gang Zhou, Houqiang Li, and Can Huang. 2023.	Anwen Hu, Haiyang Xu, Jiabo Ye, Ming Yan, Liang	746
691	Unidoc: A universal large multimodal model for si-	Zhang, Bo Zhang, Ji Zhang, Qin Jin, Fei Huang, and	747
692	multaneous text detection, recognition, spotting and	Jingren Zhou. 2024c. mPLUG-DocOwl 1.5: Uni-	748
693	understanding. <i>arXiv preprint arXiv:2308.11592</i> .	fied structure learning for OCR-free document under-	749
694	Ling Fu, Biao Yang, Zhebin Kuang, Jiajun Song, Yuzhe	<i>standing</i> . In <i>Findings of the Association for Computa-</i>	750
695	Li, Linghao Zhu, Qidi Luo, Xinyu Wang, Hao Lu,	<i>tional Linguistics: EMNLP 2024</i> , pages 3096–3120,	751
696	Mingxin Huang, et al. 2024. Ocrbench v2: An im-	Miami, Florida, USA. Association for Computational	752
697	proved benchmark for evaluating large multimodal	Linguistics.	753
698	models on visual text localization and reasoning.	Anwen Hu, Haiyang Xu, Liang Zhang, Jiabo Ye, Ming	754
699	<i>arXiv preprint arXiv:2501.00321</i> .	Yan, Ji Zhang, Qin Jin, Fei Huang, and Jingren Zhou.	755
700	Masato Fujitake. 2024. LayoutLLM: large language	2024d. mPLUG-DocOwl2: high-resolution com-	756
701	model instruction tuning for visually rich document	pressing for OCR-free multi-page document under-	757
702	understanding. In <i>International Conference on Lan-</i>	standing. <i>arXiv</i> , 2409.03420.	758
703	<i>guage Resources and Evaluation (LREC)</i> .	Jiani Huang, Haihua Chen, Fengchang Yu, and Wei Lu.	759
704	Jiaxi Gu, Xiaojun Meng, Guansong Lu, Lu Hou, Niu	2024a. From detection to application: Recent ad-	760
705	Minzhe, Xiaodan Liang, Lewei Yao, Runhui Huang,	vances in understanding scientific tables and figures.	761
706	Wei Zhang, Xin Jiang, et al. 2022. Wukong: A 100	<i>ACM Computing Surveys</i> .	762
707	million large-scale chinese cross-modal pre-training	Kung-Hsiang Huang, Hou Pong Chan, Yi R. Fung,	763
708	benchmark. <i>Advances in Neural Information Pro-</i>	Haoyi Qiu, Mingyang Zhou, Shafiq Joty, Shih-Fu	764
709	<i>cessing Systems</i> , 35:26418–26431.	Chang, and Heng Ji. 2024b. From pixels to insights:a	765

766	survey on automatic chart understanding in the era of	Geewook Kim, Hodong Lee, Daehee Kim, Haeji Jung,	823
767	large foundation models. <i>arXiv</i> , 2403.12027.	Sanghee Park, Yoonsik Kim, Sangdoo Yun, Taeho	824
768	Mingxin Huang, Yuliang Liu, Dingkan Liang,	Kil, Bado Lee, and Seunghyun Park. 2023. Visually-	825
769	Lianwen Jin, and Xiang Bai. 2024c. Mini-	situated natural language understanding with con-	826
770	monkey:alleviating the semantic sawtooth effect for	trastive reading model and frozen large language	827
771	lightweight MLLMs via complementary image pyra-	models. <i>arXiv preprint arXiv:2305.15080</i> .	828
772	mid. <i>arXiv</i> , 2408.02034.	Yoonsik Kim, Moonbin Yim, and Ka Yeon Song.	829
773	Zheng Huang, Kai Chen, Jianhua He, Xiang Bai, Di-	2024. TableVQA-Bench: a visual question answer-	830
774	mosthenis Karatzas, Shijian Lu, and CV Jawahar.	ing benchmark on multiple table domains. <i>arXiv</i> ,	831
775	2019. Icdar2019 competition on scanned receipt ocr	2404.19205.	832
776	and information extraction. In <i>2019 International</i>	Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi	833
777	<i>Conference on Document Analysis and Recognition</i>	Mao, Chloe Rolland, Laura Gustafson, Tete Xiao,	834
778	(ICDAR), pages 1516–1520. IEEE.	Spencer Whitehead, Alexander C Berg, Wan-Yen Lo,	835
779	Guillaume Jaume, Hazim Kemal Ekenel, and Jean-	et al. 2023. Segment anything. In <i>Proceedings of the</i>	836
780	Philippe Thiran. 2019. Funsd: A dataset for form	<i>IEEE/CVF International Conference on Computer</i>	837
781	understanding in noisy scanned documents. In <i>2019</i>	<i>Vision</i> , pages 4015–4026.	838
782	<i>International Conference on Document Analysis and</i>	Jianfeng Kuang, Wei Hua, Dingkan Liang, Mingkun	839
783	<i>Recognition Workshops (ICDARW)</i> , volume 2, pages	Yang, Deqiang Jiang, Bo Ren, and Xiang Bai. 2023.	840
784	1–6. IEEE.	Visual information extraction in the wild: practical	841
785	Kushal Kafle, Brian Price, Scott Cohen, and Christo-	dataset and end-to-end solution. In <i>International</i>	842
786	pher Kanan. 2018. Dvqa: Understanding data visual-	<i>Conference on Document Analysis and Recognition</i> ,	843
787	izations via question answering. In <i>Proceedings of</i>	pages 36–53. Springer.	844
788	<i>the IEEE conference on computer vision and pattern</i>	Xin Lai, Zhuotao Tian, Yukang Chen, Senqiao Yang, Xi-	845
789	<i>recognition</i> , pages 5648–5656.	angru Peng, and Jiaya Jia. 2024. Step-dpo: Step-wise	846
790	Samira Ebrahimi Kahou, Vincent Michalski, Adam	preference optimization for long-chain reasoning of	847
791	Atkinson, Ákos Kádár, Adam Trischler, and Yoshua	llms. <i>arXiv preprint arXiv:2406.18629</i> .	848
792	Bengio. 2017. Figureqa: An annotated figure	Hugo Laurençon, Andrés Marafioti, Victor Sanh, and	849
793	dataset for visual reasoning. <i>arXiv preprint</i>	Léo Tronchon. 2024. Building and better understand-	850
794	<i>arXiv:1710.07300</i> .	ing vision-language models: insights and future direc-	851
795	Dimosthenis Karatzas, Lluís Gomez-Bigorda, Angue-	tions. In <i>Workshop on Responsibly Building the Next</i>	852
796	los Nicolaou, Suman Ghosh, Andrew Bagdanov,	<i>Generation of Multimodal Foundational Models</i> .	853
797	Masakazu Iwamura, Jiri Matas, Lukas Neumann, Vi-	Hugo Laurençon, Léo Tronchon, Matthieu Cord,	854
798	jay Ramaseshan Chandrasekhar, Shijian Lu, et al.	and Victor Sanh. 2024. What matters when	855
799	2015. Icdar 2015 competition on robust reading.	building vision-language models? <i>Preprint</i> ,	856
800	In <i>2015 13th international conference on document</i>	<i>arXiv:2405.02246</i> .	857
801	<i>analysis and recognition (ICDAR)</i> , pages 1156–1160.	Bo Li, Yuanhan Zhang, Dong Guo, Renrui Zhang,	858
802	IEEE.	Feng Li, Hao Zhang, Kaichen Zhang, Peiyuan	859
803	Dimosthenis Karatzas, Faisal Shafait, Seiichi Uchida,	Zhang, Yanwei Li, Ziwei Liu, et al. 2024a. Llava-	860
804	Masakazu Iwamura, Lluís Gomez i Bigorda,	onevision: Easy visual task transfer. <i>arXiv preprint</i>	861
805	Sergi Robles Mestre, Joan Mas, David Fernan-	<i>arXiv:2408.03326</i> .	862
806	dez Mota, Jon Almazan Almazan, and Lluís Pere	Bohao Li, Yuying Ge, Yi Chen, Yixiao Ge, Ruimao	863
807	De Las Heras. 2013. Icdar 2013 robust reading	Zhang, and Ying Shan. 2024b. Seed-bench-2-plus:	864
808	competition. In <i>2013 12th international conference</i>	Benchmarking multimodal large language models	865
809	<i>on document analysis and recognition</i> , pages 1484–	with text-rich visual comprehension. <i>arXiv preprint</i>	866
810	1493. IEEE.	<i>arXiv:2404.16790</i> .	867
811	Geewook Kim, Teakgyu Hong, Moonbin Yim,	Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi.	868
812	JeongYeon Nam, Jinyoung Park, Jinyeong Yim, Won-	2023. Blip-2: Bootstrapping language-image pre-	869
813	seok Hwang, Sangdoo Yun, Dongyoon Han, and	training with frozen image encoders and large lan-	870
814	Seunghyun Park. 2022. Ocr-free document under-	guage models . <i>Preprint</i> , <i>arXiv:2301.12597</i> .	871
815	standing transformer. In <i>European Conference on</i>	Lei Li, Yuqi Wang, Runxin Xu, Peiyi Wang, Xiachong	872
816	<i>Computer Vision</i> , pages 498–517. Springer.	Feng, Lingpeng Kong, and Qi Liu. 2024c. Mul-	873
817	Geewook Kim, Teakgyu Hong, Moonbin Yim,	timodal ArXiv: A dataset for improving scientific	874
818	Jinyoung Park, Jinyeong Yim, Wonseok	comprehension of large vision-language models . In	875
819	Hwang, Sangdoo Yun, Dongyoon Han, and	<i>Proceedings of the 62nd Annual Meeting of the As-</i>	876
820	Seunghyun Park. 2021. Donut:document under-	<i>sociation for Computational Linguistics (Volume 1:</i>	877
821	standing transformer without OCR. <i>arXiv</i> ,	<i>Long Papers)</i> , pages 14369–14387, Bangkok, Thai-	878
822	30daa51cae78df53563f436a5f1cd2107655df43.	land. Association for Computational Linguistics.	879

880	Wentong Li, Yuqian Yuan, Jian Liu, Dongqi Tang, Song	Haotian Liu, Chunyuan Li, Yuheng Li, and Yong Jae	938
881	Wang, Jie Qin, Jianke Zhu, and Lei Zhang. 2024d.	Lee. 2023b. Improved baselines with visual instruc-	939
882	Tokenpacker: Efficient visual projector for multi-	tion tuning . <i>Preprint</i> , arXiv:2310.03744.	940
883	modal llm . <i>Preprint</i> , arXiv:2407.02392.		
884	Zhang Li, Biao Yang, Qiang Liu, Zhiyin Ma, Shuo	Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae	941
885	Zhang, Jingxu Yang, Yabo Sun, Yuliang Liu, and	Lee. 2023c. Visual instruction tuning.	942
886	Xiang Bai. 2024e. Monkey:image resolution and		
887	text label are important things for large multi-modal	Yuliang Liu, Zhang Li, Mingxin Huang, Biao Yang,	943
888	models. In <i>Computer Vision and Pattern Recognition</i>	Wenwen Yu, Chunyuan Li, Xu-Cheng Yin, Cheng-	944
889	(CVPR).	Lin Liu, Lianwen Jin, and Xiang Bai. 2024d. Ocr-	945
		bench: on the hidden mystery of ocr in large multi-	946
890	Zhiqi Li, Guo Chen, Shilong Liu, Shihao Wang,	modal models. <i>Science China Information Sciences</i> ,	947
891	Vibashan VS, Yishen Ji, Shiyi Lan, Hao Zhang, Yilin	67(12):220102.	948
892	Zhao, Subhashree Radhakrishnan, Nadine Chang,		
893	Karan Sapra, Amala Sanjay Deshmukh, Tuomas Rin-	Yuliang Liu, Biao Yang, Qiang Liu, Zhang Li, Zhiyin	949
894	tamaki, Matthieu Le, Iliia Karmanov, Lukas Voegtle,	Ma, Shuo Zhang, and Xiang Bai. 2024e. TextMon-	950
895	Philipp Fischer, De-An Huang, Timo Roman, Tong	key: an OCR-Free large multimodal model for under-	951
896	Lu, Jose M. Alvarez, Bryan Catanzaro, Jan Kautz,	standing document. <i>arXiv</i> , 2403.04473.	952
897	Andrew Tao, Guilin Liu, and Zhiding Yu. 2025. Ea-		
898	gle 2: Building post-training data strategies from	Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei,	953
899	scratch for frontier vision-language models . <i>Preprint</i> ,	Zheng Zhang, Stephen Lin, and Baining Guo. 2021.	954
900	arXiv:2501.14818.	Swin transformer: Hierarchical vision transformer	955
		using shifted windows. In <i>Proceedings of the</i>	956
901	Minghui Liao, Zhisheng Zou, Zhaoyi Wan, Cong Yao,	<i>IEEE/CVF international conference on computer vi-</i>	957
902	and Xiang Bai. 2022. Real-time scene text detection	<i>sion</i> , pages 10012–10022.	958
903	with differentiable binarization and adaptive scale		
904	fusion. <i>IEEE transactions on pattern analysis and</i>	Pan Lu, Liang Qiu, Kai-Wei Chang, Ying Nian Wu,	959
905	<i>machine intelligence</i> , 45(1):919–931.	Song-Chun Zhu, Tanmay Rajpurohit, Peter Clark,	960
		and Ashwin Kalyan. 2022. Dynamic prompt learning	961
906	Wenhui Liao, Jiapeng Wang, Hongliang Li, Chengyu	via policy gradient for semi-structured mathematical	962
907	Wang, Jun Huang, and Lianwen Jin. 2024a. Do-	reasoning. <i>arXiv preprint arXiv:2209.14610</i> .	963
908	cLayLLM: an efficient and effective multi-modal ex-		
909	extension of large language models for text-rich docu-	Chuwei Luo, Yufan Shen, Zhaoqing Zhu, Qi Zheng,	964
910	ment understanding. <i>arXiv</i> , 2408.15045.	Zhi Yu, and Cong Yao. 2024a. LayoutLLM: layout	965
		instruction tuning with large language models for	966
911	Wenhui Liao, Jiapeng Wang, Hongliang Li, Chengyu	document understanding. In <i>Computer Vision and</i>	967
912	Wang, Jun Huang, and Lianwen Jin. 2024b. Do-	<i>Pattern Recognition (CVPR)</i> .	968
913	clayllm: An efficient and effective multi-modal exten-		
914	sion of large language models for text-rich document	Chuwei Luo, Yufan Shen, Zhaoqing Zhu, Qi Zheng,	969
915	understanding. <i>arXiv preprint arXiv:2408.15045</i> .	Zhi Yu, and Cong Yao. 2024b. Layoutllm: Lay-	970
		out instruction tuning with large language models	971
916	Chaohu Liu, Kun Yin, Haoyu Cao, Xinghua Jiang, Xin	for document understanding. In <i>Proceedings of the</i>	972
917	Li, Yinsong Liu, Deqiang Jiang, Xing Sun, and Linli	<i>IEEE/CVF Conference on Computer Vision and Pat-</i>	973
918	Xu. 2024a. Hrvda: High-resolution visual document	<i>tern Recognition</i> , pages 15630–15640.	974
919	assistant. In <i>Proceedings of the IEEE/CVF Confer-</i>		
920	<i>ence on Computer Vision and Pattern Recognition</i> ,	Tengchao Lv, Yupan Huang, Jingye Chen, Yuzhong	975
921	pages 15534–15545.	Zhao, Yilin Jia, Lei Cui, Shuming Ma, Yaoyao Chang,	976
922	Chenglong Liu, Haoran Wei, Jinyue Chen, Lingyu	Shaohan Huang, Wenhui Wang, Li Dong, Weiyao	977
923	Kong, Zheng Ge, Zining Zhu, Liang Zhao, Jianjian	Luo, Shaoxiang Wu, Guoxin Wang, Cha Zhang, and	978
924	Sun, Chunrui Han, and Xiangyu Zhang. 2024b. Fo-	Furu Wei. 2024. KOSMOS-2.5: a multimodal liter-	979
925	cus anywhere for fine-grained multi-page document	ate model. <i>arXiv</i> , 2309.11419.	980
926	understanding. <i>arXiv preprint arXiv:2405.14295</i> .		
927	Fuxiao Liu, Kevin Lin, Linjie Li, Jianfeng Wang, Yaser	Tengchao Lv, Yupan Huang, Jingye Chen, Yuzhong	981
928	Yacoob, and Lijuan Wang. 2023a. Mitigating hal-	Zhao, Yilin Jia, Lei Cui, Shuming Ma, Yaoyao Chang,	982
929	lucination in large multi-modal models via robust	Shaohan Huang, Wenhui Wang, et al. 2023. Kosmos-	983
930	instruction tuning. In <i>The Twelfth International Con-</i>	2.5: A multimodal literate model. <i>arXiv preprint</i>	984
931	<i>ference on Learning Representations</i> .	<i>arXiv:2309.11419</i> .	985
932	Fuxiao Liu, Xiaoyang Wang, Wenlin Yao, Jianshu Chen,	Yubo Ma, Yuhang Zang, Liangyu Chen, Meiqi Chen,	986
933	Kaiqiang Song, Sangwoo Cho, Yaser Yacoob, and	Yizhu Jiao, Xinze Li, Xinyuan Lu, Ziyu Liu, Yan Ma,	987
934	Dong Yu. 2024c. MMC: advancing multimodal chart	Xiaoyi Dong, Pan Zhang, Liangming Pan, Yu-Gang	988
935	understanding with large-scale instruction tuning. In	Jiang, Jiaqi Wang, Yixin Cao, and Aixin Sun. 2024.	989
936	<i>North American Chapter of the Association for Com-</i>	MMLongBench-Doc: benchmarking long-context	990
937	<i>putational Linguistics (NAACL)</i> .	document understanding with visualizations. In <i>Con-</i>	991
		<i>ference on Neural Information Processing Systems</i>	992
		(<i>NeurIPS</i>).	993

- U-V Marti and Horst Bunke. 2002. The iam-database: an english sentence database for offline handwriting recognition. *International journal on document analysis and recognition*, 5:39–46.
- Ahmed Masry, Parsa Kavehzadeh, Xuan Long Do, Enamul Hoque, and Shafiq Joty. 2023. Unichart: A universal vision-language pretrained model for chart comprehension and reasoning. *arXiv preprint arXiv:2305.14761*.
- Ahmed Masry, Do Xuan Long, Jia Qing Tan, Shafiq Joty, and Enamul Hoque. 2022. Chartqa: A benchmark for question answering about charts with visual and logical reasoning. *arXiv preprint arXiv:2203.10244*.
- Minesh Mathew, Viraj Bagal, Rubèn Tito, Dimosthenis Karatzas, Ernest Valveny, and CV Jawahar. 2022. Infographicvqa. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 1697–1706.
- Minesh Mathew, Lluís Gomez, Dimosthenis Karatzas, and CV Jawahar. 2021a. Asking questions on handwritten document collections. *International Journal on Document Analysis and Recognition (IJDR)*, 24(3):235–249.
- Minesh Mathew, Dimosthenis Karatzas, and CV Jawahar. 2021b. Docvqa: A dataset for vqa on document images. In *Proceedings of the IEEE/CVF winter conference on applications of computer vision*, pages 2200–2209.
- Nitesh Methani, Pritha Ganguly, Mitesh M Khapra, and Pratyush Kumar. 2020. Plotqa: Reasoning over scientific plots. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 1527–1536.
- Anand Mishra, Shashank Shekhar, Ajeet Kumar Singh, and Anirban Chakraborty. 2019. Ocr-vqa: Visual question answering by reading text in images. In *2019 international conference on document analysis and recognition (ICDAR)*, pages 947–952. IEEE.
- Mor Shpigel Nacson, Aviad Aberdam, Roy Ganz, Elad Ben Avraham, Alona Golts, Yair Kittenplon, Shai Mazor, and Ron Litman. 2024. [Docvlm: Make your vlm an efficient reader](#). *Preprint*, arXiv:2412.08746.
- Jason Obeid and Enamul Hoque. 2020. [Chart-to-text: Generating natural language descriptions for charts by adapting the transformer model](#). In *Proceedings of the 13th International Conference on Natural Language Generation*, pages 138–147, Dublin, Ireland. Association for Computational Linguistics.
- Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, et al. 2023. Dinov2: Learning robust visual features without supervision. *arXiv preprint arXiv:2304.07193*.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. 2022. Training language models to follow instructions with human feedback. *Advances in neural information processing systems*, 35:27730–27744.
- Richard Yuanzhe Pang, Weizhe Yuan, He He, Kyunghyun Cho, Sainbayar Sukhbaatar, and Jason Weston. 2025. Iterative reasoning preference optimization. *Advances in Neural Information Processing Systems*, 37:116617–116637.
- Jaeyoo Park, Jin Young Choi, Jeonghyung Park, and Bohyung Han. 2024. Hierarchical visual feature aggregation for ocr-free document understanding. *arXiv preprint arXiv:2411.05254*.
- Panupong Pasupat, Percy Liang, and etc. 2015. Compositional semantic parsing on semi-structured tables. *arXiv preprint arXiv:1508.00305*.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR.
- Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D Manning, Stefano Ermon, and Chelsea Finn. 2024. Direct preference optimization: Your language model is secretly a reward model. *Advances in Neural Information Processing Systems*, 36.
- Sachin Raja, Ajoy Mondal, and CV Jawahar. 2023. Icdar 2023 competition on visual question answering on business document images. In *International Conference on Document Analysis and Recognition*, pages 454–470. Springer.
- C Schuhmann, A Köpf, R Vencu, T Coombes, and R Beaumont. 2022. Laion coco: 600m synthetic captions from laion2b-en. URL <https://laion.ai/blog/laion-coco>.
- Hao Shao, Shengju Qian, Han Xiao, Guanglu Song, Zhuofan Zong, Letian Wang, Yu Liu, and Hongsheng Li. 2025. Visual cot: Advancing multi-modal language models with a comprehensive dataset and benchmark for chain-of-thought reasoning. *Advances in Neural Information Processing Systems*, 37:8612–8642.
- Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang, Mingchuan Zhang, YK Li, Y Wu, et al. 2024. Deepseekmath: Pushing the limits of mathematical reasoning in open language models. *arXiv preprint arXiv:2402.03300*.
- Baoguang Shi, Cong Yao, Minghui Liao, Mingkun Yang, Pei Xu, Linyan Cui, Serge Belongie, Shijian Lu, and Xiang Bai. 2017. Icdar2017 competition on reading chinese text in the wild (rctw-17). In *2017 14th iapr international conference on document analysis and recognition (ICDAR)*, volume 1, pages 1429–1434. IEEE.

1106	Min Shi, Fuxiao Liu, Shihao Wang, Shijia Liao, Subhashree Radhakrishnan, De-An Huang, Hongxu Yin, Karan Sapra, Yaser Yacoob, Humphrey Shi, Bryan Catanzaro, Andrew Tao, Jan Kautz, Zhiding Yu, and Guilin Liu. 2024. Eagle: Exploring the design space for multimodal llms with mixture of encoders . <i>Preprint</i> , arXiv:2408.15998.	
1107		
1108		
1109		
1110		
1111		
1112		
1113	Oleksii Sidorov, Ronghang Hu, Marcus Rohrbach, and Amanpreet Singh. 2020. Textcaps: a dataset for image captioning with reading comprehension. In <i>Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part II 16</i> , pages 742–758. Springer.	
1114		
1115		
1116		
1117		
1118		
1119	Amanpreet Singh, Vivek Natarajan, Meet Shah, Yu Jiang, Xinlei Chen, Dhruv Batra, Devi Parikh, and Marcus Rohrbach. 2019. Towards vqa models that can read. In <i>Proceedings of the IEEE/CVF conference on computer vision and pattern recognition</i> , pages 8317–8326.	
1120		
1121		
1122		
1123		
1124		
1125	Tomasz Stanisławek, Filip Graliński, Anna Wróblewska, Dawid Lipiński, Agnieszka Kaliska, Paulina Rosalska, Bartosz Topolski, and Przemysław Biecek. 2021. Kleister: key information extraction datasets involving long documents with complex layouts. In <i>International Conference on Document Analysis and Recognition</i> , pages 564–579. Springer.	
1126		
1127		
1128		
1129		
1130		
1131		
1132	N Subramani, A Matton, M Greaves, and A Lam. A survey of deep learning approaches for ocr and document understanding. arxiv 2020. <i>arXiv preprint arXiv:2011.13534</i> .	
1133		
1134		
1135		
1136	Hamed Rahimi Sujet AI, Allaa Boutaleb. 2024. Sujet-finance-qa-vision-100k: A large-scale dataset for financial document vqa .	
1137		
1138		
1139	Ningyuan Sun, Xuefeng Yang, and Yunfeng Liu. 2020. Tableqa: a large-scale chinese text-to-sql dataset for table-aware sql generation. <i>arXiv preprint arXiv:2006.06434</i> .	
1140		
1141		
1142		
1143	Yipeng Sun, Zihan Ni, Chee-Kheng Chng, Yuliang Liu, Canjie Luo, Chun Chet Ng, Junyu Han, Errui Ding, Jingtuo Liu, Dimosthenis Karatzas, et al. 2019. Icdar 2019 competition on large-scale street view text with partial labeling-rrc-lsvt. In <i>2019 International Conference on Document Analysis and Recognition (ICDAR)</i> , pages 1557–1562. IEEE.	
1144		
1145		
1146		
1147		
1148		
1149		
1150	S Svetlichnaya. 2020. Deepform: Understand structured documents at scale.	
1151		
1152	Ryota Tanaka, Taichi Iki, Kyosuke Nishida, Kuniko Saito, and Jun Suzuki. 2024a. InstructDoc: a dataset for zero-shot generalization of visual document understanding with instructions. In <i>AAAI Conference on Artificial Intelligence (AAAI)</i> .	
1153		
1154		
1155		
1156		
1157	Ryota Tanaka, Taichi Iki, Kyosuke Nishida, Kuniko Saito, and Jun Suzuki. 2024b. Instructdoc: A dataset for zero-shot generalization of visual document understanding with instructions. In <i>Proceedings of the AAAI Conference on Artificial Intelligence</i> , volume 38, pages 19071–19079.	
1158		
1159		
1160		
1161		
1162		
	Ryota Tanaka, Kyosuke Nishida, Kosuke Nishida, Taku Hasegawa, Itsumi Saito, and Kuniko Saito. 2023. Slidevqa: A dataset for document visual question answering on multiple images. In <i>Proceedings of the AAAI Conference on Artificial Intelligence</i> , volume 37, pages 13636–13645.	1163 1164 1165 1166 1167 1168
	Ryota Tanaka, Kyosuke Nishida, and Sen Yoshida. 2021. Visualmrc: Machine reading comprehension on document images. In <i>Proceedings of the AAAI Conference on Artificial Intelligence</i> , volume 35, pages 13878–13888.	1169 1170 1171 1172 1173
	Benny Tang, Angie Boggust, and Arvind Satyanarayan. 2023. VisText: A benchmark for semantically rich chart captioning . In <i>Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)</i> , pages 7268–7298, Toronto, Canada. Association for Computational Linguistics.	1174 1175 1176 1177 1178 1179 1180
	Jingqun Tang, Qi Liu, Yongjie Ye, Jinghui Lu, Shu Wei, Chunhui Lin, Wanqing Li, Mohamad Fitri Faiz Bin Mahmood, Hao Feng, Zhen Zhao, et al. 2024. Mtvqa: Benchmarking multilingual text-centric visual question answering. <i>arXiv preprint arXiv:2405.11985</i> .	1181 1182 1183 1184 1185
	Rubèn Tito, Dimosthenis Karatzas, and Ernest Valveny. 2023. Hierarchical multimodal transformers for multipage docvqa. <i>Pattern Recognition</i> , 144:109834.	1186 1187 1188
	Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. 2023a. Llama: Open and efficient foundation language models. <i>arXiv preprint arXiv:2302.13971</i> .	1189 1190 1191 1192 1193 1194
	Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. 2023b. Llama 2: Open foundation and fine-tuned chat models. <i>arXiv preprint arXiv:2307.09288</i> .	1195 1196 1197 1198 1199 1200
	Thanh-Nghia Truong, Cuong Tuan Nguyen, Richard Zanibbi, Harold Mouchère, and Masaki Nakagawa. 2024. A survey on handwritten mathematical expression recognition: The rise of encoder-decoder and gnn models . <i>Pattern Recognition</i> , 153:110531.	1201 1202 1203 1204 1205
	Michał Turski, Tomasz Stanisławek, Karol Kaczmarek, Paweł Dyda, and Filip Graliński. 2023. Ccpdf: Building a high quality corpus for visually rich documents from web crawl data. In <i>International Conference on Document Analysis and Recognition</i> , pages 348–365. Springer.	1206 1207 1208 1209 1210 1211
	Rohan Wadhawan, Hritik Bansal, Kai-Wei Chang, and Nanyun Peng. 2024. Contextual: Evaluating context-sensitive text-rich visual reasoning in large multimodal models . In <i>Forty-first International Conference on Machine Learning</i> .	1212 1213 1214 1215 1216

1217	Bryan Wang, Gang Li, Xin Zhou, Zhourong Chen, Tovi Grossman, and Yang Li. 2021. Screen2words: Automatic mobile ui summarization with multimodal learning. In <i>The 34th Annual ACM Symposium on User Interface Software and Technology</i> , pages 498–510.	1274
1218		1275
1219		1276
1220		1277
1221		1278
1222		1279
1223	Dongsheng Wang, Natraj Raman, Mathieu Sibue, Zhiqiang Ma, Petr Babkin, Simerjot Kaur, Yulong Pei, Armineh Nourbakhsh, and Xiaomo Liu. 2023. DocLLM: a layout-aware generative language model for multimodal document understanding. <i>arXiv</i> , 2401.00908.	1280
1224		1281
1225		1282
1226		1283
1227		1284
1228		1285
1229	Fei Wang, Wenxuan Zhou, James Y Huang, Nan Xu, Sheng Zhang, Hoifung Poon, and Muhao Chen. 2024a. mdpo: Conditional preference optimization for multimodal large language models. <i>arXiv preprint arXiv:2406.11839</i> .	1286
1230		1287
1231		1288
1232		1289
1233		1290
1234	Peng Wang, Shuai Bai, Sinan Tan, Shijie Wang, Zhihao Fan, Jinze Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, et al. 2024b. Qwen2-vl: Enhancing vision-language model’s perception of the world at any resolution. <i>arXiv preprint arXiv:2409.12191</i> .	1291
1235		
1236		
1237		
1238		
1239	Weiyun Wang, Zhe Chen, Wenhai Wang, Yue Cao, Yangzhou Liu, Zhangwei Gao, Jinguo Zhu, Xizhou Zhu, Lewei Lu, Yu Qiao, and Jifeng Dai. 2024c. Enhancing the reasoning ability of multimodal large language models via mixed preference optimization. <i>Preprint</i> , arXiv:2411.10442.	1292
1240		1293
1241		1294
1242		1295
1243		1296
1244		
1245	Weiyun Wang, Zhe Chen, Wenhai Wang, Yue Cao, Yangzhou Liu, Zhangwei Gao, Jinguo Zhu, Xizhou Zhu, Lewei Lu, Yu Qiao, et al. 2024d. Enhancing the reasoning ability of multimodal large language models via mixed preference optimization. <i>arXiv preprint arXiv:2411.10442</i> .	1297
1246		1298
1247		1299
1248		1300
1249		
1250		
1251	Haoran Wei, Lingyu Kong, Jinyue Chen, Liang Zhao, Zheng Ge, Jinrong Yang, Jianjian Sun, Chunrui Han, and Xiangyu Zhang. 2024. Vary: Scaling up the vision vocabulary for large vision-language model. In <i>European Conference on Computer Vision</i> , pages 408–424. Springer.	1301
1252		1302
1253		1303
1254		1304
1255		1305
1256		1306
1257	Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. 2022. Chain-of-thought prompting elicits reasoning in large language models. <i>Advances in neural information processing systems</i> , 35:24824–24837.	1307
1258		1308
1259		1309
1260		1310
1261		
1262	Shijie Wen, Minglang Qiao, Lai Jiang, Mai Xu, Xin Deng, and Shengxi Li. 2024. Mt-vqa: A multi-task approach for quality assessment of short-form videos. In <i>Proceedings of the 3rd Workshop on Quality of Experience in Visual Multimedia Applications</i> , pages 30–38.	1311
1263		1312
1264		1313
1265		1314
1266		
1267		
1268	Sanghyun Woo, Shoubhik Debnath, Ronghang Hu, Xinlei Chen, Zhuang Liu, In So Kweon, and Saining Xie. 2023. Convnext v2: Co-designing and scaling convnets with masked autoencoders. In <i>Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition</i> , pages 16133–16142.	1315
1269		1316
1270		1317
1271		1318
1272		1319
1273		1320
	Zhiyu Wu, Xiaokang Chen, Zizheng Pan, Xingchao Liu, Wen Liu, Damai Dai, Huazuo Gao, Yiyang Ma, Chengyue Wu, Bingxuan Wang, et al. 2024. Deepseek-vl2: Mixture-of-experts vision-language models for advanced multimodal understanding. <i>arXiv preprint arXiv:2412.10302</i> .	1321
		1322
		1323
		1324
		1325
		1326
		1327
		1328
		1329
	Renqiu Xia, Song Mao, Xiangchao Yan, Hongbin Zhou, Bo Zhang, Haoyang Peng, Jiahao Pi, Daocheng Fu, Wenjie Wu, Hancheng Ye, et al. 2024a. Docgenome: An open large-scale scientific document benchmark for training and testing multi-modal large language models. <i>arXiv preprint arXiv:2406.11633</i> .	
	Renqiu Xia, Bo Zhang, Hancheng Ye, Xiangchao Yan, Qi Liu, Hongbin Zhou, Zijun Chen, Min Dou, Botian Shi, Junchi Yan, et al. 2024b. Chartx & chartvlm: A versatile benchmark and foundation model for complicated chart reasoning. <i>arXiv preprint arXiv:2402.12185</i> .	
	Xudong Xie, Liang Yin, Hao Yan, Yang Liu, Jing Ding, Minghui Liao, Yuliang Liu, Wei Chen, and Xiang Bai. 2024. PDF-WuKong: a large multimodal model for efficient long PDF reading with end-to-end sparse sampling. <i>arXiv</i> , 2410.05970.	
	Yiheng Xu, Minghao Li, Lei Cui, Shaohan Huang, Furu Wei, and Ming Zhou. 2019. LayoutLM: pre-training of text and layout for document image understanding. In <i>Knowledge Discovery and Data Mining</i> .	
	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 <i>Proceedings of the 26th ACM SIGKDD international conference on knowledge discovery & data mining</i> , pages 1192–1200.	
	Zhengzhuo Xu, Sinan Du, Yiyan Qi, Chengjin Xu, Chun Yuan, and Jian Guo. 2024. ChartBench: a benchmark for complex visual reasoning in charts. <i>arXiv</i> , 2312.15915.	
	An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Haoran Wei, et al. 2024a. Qwen2. 5 technical report. <i>arXiv preprint arXiv:2412.15115</i> .	
	Zhibo Yang, Jun Tang, Zhaohai Li, Pengfei Wang, Jianqiang Wan, Humen Zhong, Xuejing Liu, Mingkun Yang, Peng Wang, Yuliang Liu, et al. 2024b. Cc-ocr: A comprehensive and challenging ocr benchmark for evaluating large multimodal models in literacy. <i>arXiv preprint arXiv:2412.02210</i> .	
	Jiabo Ye, Anwen Hu, Haiyang Xu, Qinghao Ye, Ming Yan, Guohai Xu, Chenliang Li, Junfeng Tian, Qi Qian, Ji Zhang, Qin Jin, Liang He, Xin Lin, and Fei Huang. 2023a. UReader: Universal OCR-free visually-situated language understanding with multimodal large language model. In <i>Findings of the Association for Computational Linguistics: EMNLP 2023</i> , pages 2841–2858, Singapore. Association for Computational Linguistics.	

1330	1384	Yilun Zhao, Yunxiang Li, Chenying Li, and Rui Zhang.
1331	1385	2022. MultiHiertt: Numerical reasoning over multi
1332	1386	hierarchical tabular and textual data . In <i>Proceedings</i>
1333	1387	<i>of the 60th Annual Meeting of the Association for</i>
1334	1388	<i>Computational Linguistics (Volume 1: Long Papers)</i> ,
1335	1389	pages 6588–6600, Dublin, Ireland. Association for
	1390	Computational Linguistics.
1336	1391	Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan
1337	1392	Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin,
1338	1393	Zhuohan Li, Dacheng Li, Eric Xing, et al. 2023.
1339	1394	Judging llm-as-a-judge with mt-bench and chatbot
1340	1395	arena. <i>Advances in Neural Information Processing</i>
	1396	<i>Systems</i> , 36:46595–46623.
1341	1397	Mingyu Zheng, Xinwei Feng, Qingyi Si, Qiaoqiao She,
1342	1398	Zheng Lin, Wenbin Jiang, and Weiping Wang. 2024.
1343	1399	Multimodal table understanding. <i>arXiv preprint</i>
1344	1400	<i>arXiv:2406.08100</i> .
1345	1401	Xu Zhong, Elaheh ShafieiBavani, and Antonio Ji-
1346	1402	meno Yepes. 2020. Image-based table recognition:
1347	1403	data, model, and evaluation. In <i>European conference</i>
1348	1404	<i>on computer vision</i> , pages 564–580. Springer.
1349	1405	Fengbin Zhu, Wenqiang Lei, Fuli Feng, Chao Wang,
1350	1406	Haozhou Zhang, and Tat-Seng Chua. 2022. Towards
1351	1407	complex document understanding by discrete reason-
1352	1408	ing. In <i>Proceedings of the 30th ACM International</i>
	1409	<i>Conference on Multimedia</i> , pages 4857–4866.
1353	1410	Fengbin Zhu, Ziyang Liu, Xiang Yao Ng, Haohui
1354	1411	Wu, Wenjie Wang, Fuli Feng, Chao Wang, Huanbo
1355	1412	Luan, and Tat Seng Chua. 2024. MMDocBench:
1356	1413	benchmarking large vision-language models for fine-
1357	1414	grained visual document understanding. <i>arXiv</i> ,
	1415	2410.21311.
1358		
1359		
1360		
1361		
1362		
1363		
1364		
1365		
1366		
1367		
1368		
1369		
1370		
1371		
1372		
1373		
1374		
1375		
1376		
1377		
1378		
1379		
1380		
1381		
1382		
1383		