

MangaVQA and MangaLMM: A Benchmark and Specialized Model for Multimodal Manga Understanding

Jeonghun Baek* Kazuki Egashira* Shota Onohara* Atsuyuki Miyai*

Yuki Imajuku Hikaru Ikuta Kiyoharu Aizawa

The University of Tokyo

baek@hal.t.u-tokyo.ac.jp

https://manga109.github.io/MangaVQA_LMM/

Abstract

Manga, or Japanese comics, is a richly multimodal narrative form that blends images and text in complex ways. Teaching large multimodal models (LMMs) to understand such narratives at a human-like level could help manga creators reflect on and refine their stories. To this end, we introduce two benchmarks for multimodal manga understanding: MangaOCR, which targets in-page text recognition, and MangaVQA, a novel benchmark designed to evaluate contextual understanding through visual question answering. MangaVQA consists of 526 high-quality, manually constructed question-answer pairs, enabling reliable evaluation across diverse narrative and visual scenarios. Building on these benchmarks, we develop MangaLMM, a manga-specialized model finetuned from the open-source LMM Qwen2.5-VL to jointly handle both tasks. Through extensive experiments, including comparisons with proprietary models such as GPT-4o and Gemini 2.5, we assess how well LMMs understand manga. Our benchmark and model provide a comprehensive foundation for evaluating and advancing LMMs in the richly narrative domain of manga.

1 Introduction

Manga is a rich and distinctive form of multimodal narrative, combining complex panel layouts, expressive visual elements, and text embedded directly within images. As large multimodal models (LMMs) continue to advance in vision-language understanding, enabling them to understand manga presents an exciting opportunity, not only as a technical milestone, but also as a way to support human creativity. Such models could assist manga creators in reflecting on and refining their stories. To provide meaningful assistance, an LMM would need to function like a skilled editor or assistant, capable of reading and understanding manga in a way

human does. This calls for evaluating models' abilities to process visual-textual content and follow the context in a coherent and human-like manner.

Although recent efforts such as Magi (Sachdeva and Zisserman, 2024; Sachdeva et al., 2024; Sachdeva and Zisserman, 2025) and CoMix (Vivoli et al., 2024) have tackled comic understanding, they primarily focus on generating transcriptions from comic pages – they do not evaluate to what extent models can accurately read in-page text using optical character recognition (OCR), or understand the content based on that text through visual question answering (VQA). As a result, it remains unclear to what extent models truly comprehend manga content in a human-like manner based on the embedded textual information.

To pave a reliable path toward comprehensive manga understanding in LMMs, we believe it is essential to evaluate two core capabilities: OCR and VQA. To address these needs, we propose two benchmarks: MangaOCR and MangaVQA. **MangaOCR** focuses on detecting and recognizing textual content such as dialogue and sound effects. We consolidate existing annotations from the well-known Manga109 dataset (Matsui et al., 2017; Aizawa et al., 2020) and the manga onomatopoeia dataset (Baek et al., 2022) to construct this benchmark. Further, as our primary contribution, we propose **MangaVQA**, a novel benchmark designed to evaluate an LMM's ability to accurately answer targeted, factual questions grounded in both visual and textual context. It consists of 526 high-quality, manually constructed question-answer pairs covering a diverse range of scenarios, enabling assessment of a model's narrative understanding. Together, these benchmarks provide a comprehensive framework for evaluating a model's ability to understand manga as a multimodal narrative medium, with MangaVQA playing a central role in assessing deeper semantic and contextual comprehension.

Furthermore, truly human-like understanding of

*Equal contribution.

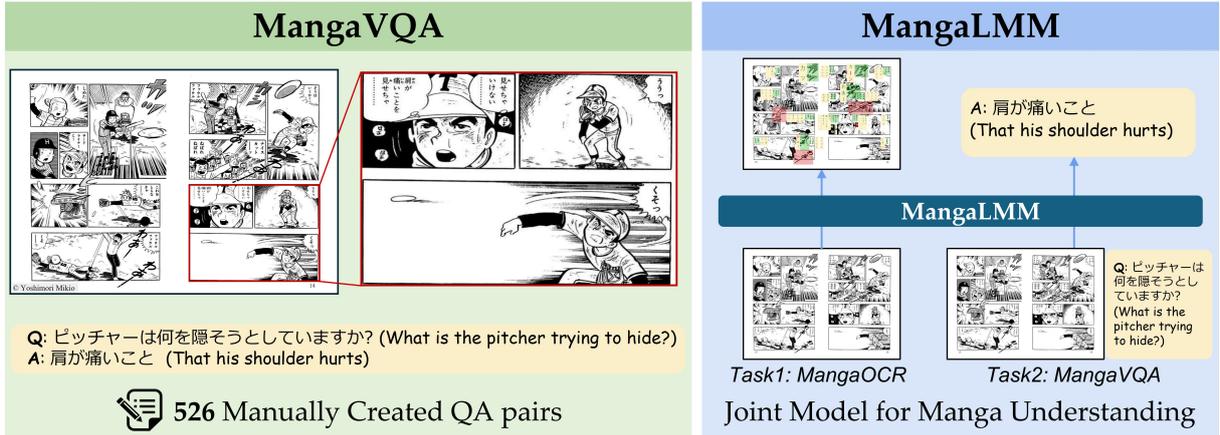


Figure 1: **Overview of MangaVQA and MangaLMM.** We present MangaVQA, a newly proposed benchmark for multimodal context understanding, consisting of 526 manually constructed question–answer pairs. We also develop MangaLMM, a manga-specialized model jointly trained to handle both MangaOCR and MangaVQA tasks.

manga requires the ability to jointly perform both OCR and VQA, rather than treating them as isolated tasks. Therefore, building on our two proposed benchmarks, we finetune an open-source LMM (Qwen2.5-VL (Bai et al., 2025)) to develop **MangaLMM**, a manga-specialized model designed to jointly address both OCR and VQA tasks. MangaLMM serves as a practical baseline for human-like manga understanding. We conduct comprehensive experiments, including analyses on model and dataset size, and compare MangaLMM with state-of-the-art proprietary models such as GPT-4o (Hurst et al., 2024) and Gemini 2.5 (Google DeepMind, 2025) to evaluate the current landscape of multimodal manga understanding. Our results show that even the proprietary models struggle on our two benchmarks, while MangaLMM jointly handle OCR and VQA, achieving promising performance on both.

An overview of our MangaVQA benchmark and the MangaLMM model is shown in Figure 1. Our contributions are summarized as follows:

- We present **MangaVQA**, a benchmark of 526 manually constructed question–answer pairs, and **MangaOCR**, focusing on in-page text detection and recognition. Together, they enable comprehensive evaluation of multimodal manga understanding.
- We develop **MangaLMM**, a manga-specialized version of Qwen2.5-VL, finetuned to jointly address VQA and OCR.
- We perform extensive analysis and evaluate MangaLMM against proprietary models, highlighting the limitations of general-purpose LMMs in stylized visual domains.

2 Related Work: Comic Data and Tasks

Recent work, CoMix (Vivoli et al., 2024), has unified various comic-related tasks by analyzing existing datasets, including French comics (eBDtheque (Guérin et al., 2013)), American comics (COMICS (Iyer et al., 2017) and DCM772 (Nguyen et al., 2018)), and Japanese comics (Manga109 (Matsui et al., 2017) and PopManga (Sachdeva and Zisserman, 2024)). CoMix primarily focuses on transcript generation-related tasks, including object detection, speaker identification, character re-identification, reading order prediction, and character naming prediction. Similarly, the recent Magi series (v1 (Sachdeva and Zisserman, 2024), v2 (Sachdeva et al., 2024), and v3 (Sachdeva and Zisserman, 2025)) also centers on transcript generation. Notably, Magi v3 extends this pipeline by generating image captions from transcriptions and further producing prose based on those captions.

Although recent studies such as CoMix and the Magi series have addressed a wide range of tasks, the evaluation of OCR has often been underexplored, particularly in detecting the locations of texts within an image and recognizing their content. One exception is COMICS TEXT+ (Soykan et al., 2024), which evaluates OCR performance at the panel level, but it does not address page-level evaluation. However, humans typically perceive and interpret text at the page level, integrating visual and textual cues across the entire layout. To reflect this human reading process, we evaluate OCR performance on two-page spreads using MangaOCR. A detailed discussion is provided in Appendix A.

Existing studies have also largely overlooked the visual question answering (VQA) task in the context of comics. Among prior datasets, the Manga Understanding Benchmark (MangaUB (Ikuta et al., 2025)) is the most closely related to our proposed MangaVQA. While MangaUB can be considered a simple VQA benchmark, it contains only eight predefined question types—such as identifying the number of characters, the weather, or the time of day—thus offering limited question diversity. As a result, MangaUB does not address a broad spectrum of VQA problems centered on text understanding in manga. Furthermore, its scope is restricted to the panel level.

In contrast, MangaVQA goes beyond individual panels and focuses on two-page spreads, reflecting how humans naturally read manga. It features diverse VQA questions grounded in textual content at the spread level, aiming to approximate the reading experience of human readers. In this regard, MangaVQA is conceptually aligned with TextVQA (Singh et al., 2019) and DocVQA (Mathew et al., 2021), as it requires models to understand and reason over text embedded in images.

3 The Manga109 Dataset and Our Consolidated MangaOCR Dataset

This section presents the widely used manga dataset Manga109 (Matsui et al., 2017) and our MangaOCR Benchmark.

3.1 Manga109: A Widely Used Dataset for Manga Research

Among the many comic datasets introduced in the Related Work, we selected Manga109 for its open-access license, diverse manga titles, and rich annotations and meta-information. It has also been widely used in previous comic-related research (Sachdeva et al., 2024; Sachdeva and Zisserman, 2025; Baek et al., 2022; Li et al., 2024b; Ikuta et al., 2025), making it a reliable and practical dataset for our study.

Manga109 is a dataset composed of 109 volumes of Japanese comics (manga). Manga is a unique visual storytelling medium characterized by spatially arranged panels and artistic expression. The Manga109 dataset captures many distinctive features of manga, including its predominantly black-and-white artwork, two-page spreads, right-to-left reading order, vertical text layout, and the frequent

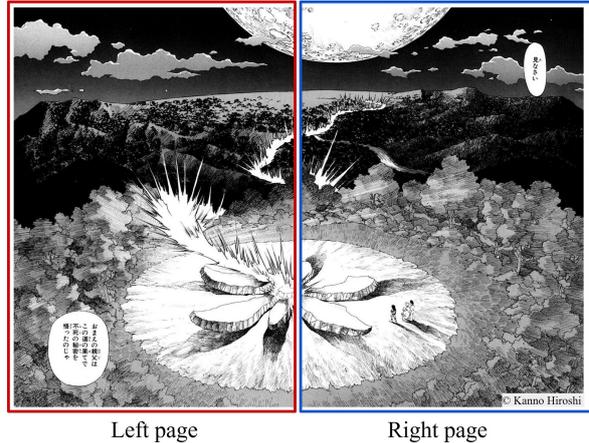


Figure 2: Illustration of a two-page spread from the Manga109 dataset.

use of stylized onomatopoeia (e.g., Boom, Bang) integrated into the illustrations. It also contains culturally specific dialogue, often incorporating honorifics and idiomatic expressions. Although these characteristics are not explicitly annotated, they present unique challenges for manga understanding tasks. Given these characteristics, Manga109 serves as a representative dataset for developing and evaluating manga understanding models. Figure 2 shows an example of two-page spreads from the Manga109 dataset.

3.2 MangaOCR: A Consolidated Dataset for Manga Text Recognition

Text in manga carries essential narrative information, appearing as speech balloons and stylized onomatopoeia integrated into the artwork. Recognizing such text is crucial for machine understanding of manga, as humans also rely on this information to comprehend the story. MangaOCR addresses this challenge by targeting two key categories of embedded text: dialogue and onomatopoeia. We construct the MangaOCR dataset by consolidating existing annotations from the Manga109 dataset and the manga onomatopoeia dataset (Baek et al., 2022). It contains approximately 209K narrative text instances, spanning a wide variety of visual styles and layouts. Training with MangaOCR can improve the ability of LMMs to extract and interpret textual information in manga, contributing to better overall understanding. The MangaOCR task is performed on two-page spreads and primarily consists of two sub-tasks: text detection, which localizes textual regions, and text recognition, which reads the localized text.

Author-Aware Dataset Split. We adopt the

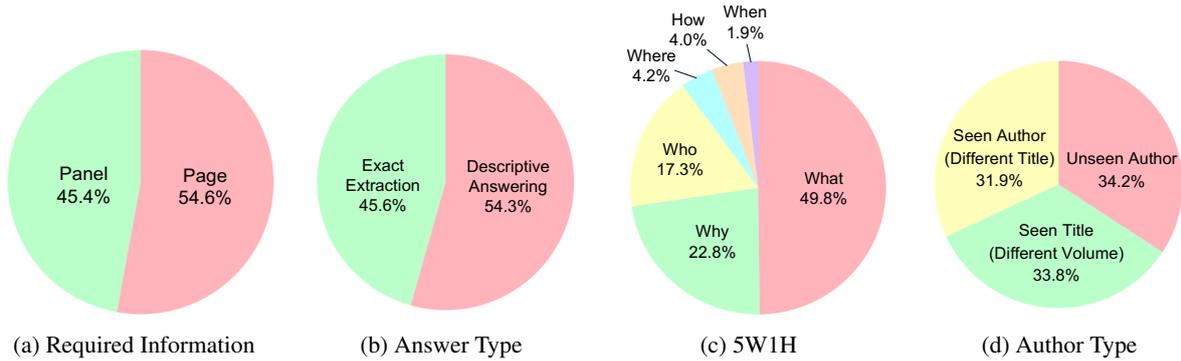


Figure 3: **Distributions in MangaVQA.** The dataset is structured along four key axes: (a) Required Information, (b) Answer Type, (c) 5W1H, and (d) Author Type.

Count type	Total	Train	Valid	Test
Comic volumes	109	89	7	13
Images	10,602	8,763	673	1,166
MangaOCR				
Dialogue	148K	120K	9K	18K
Onomatopoeia	61K	50K	4K	7K
Total	209K	170K	13K	26K
MangaVQA				
QA pairs	40,363	39,837	—	526

Table 1: **Statistics of manga datasets.** More details about MangaVQA are presented in §4 and §5.

dataset split protocol from prior work (Baek et al., 2022), with a few modifications. In the original split, the 109 volumes were divided into training, validation, and test sets based on author information. To evaluate intra-series generalization, five of the ten test volumes belong to the same series as those in the training set, where the first volume is included in the training set and the last volume is in the test set. This setting tests whether a model trained on the beginning of a series can generalize to its later volumes. To evaluate intra-author generalization, the remaining five test volumes are titles by authors who also have other works in the training set. This allows us to assess whether a model can generalize across different works by the same author.

To further evaluate out-of-distribution generalization with respect to author identity, we move three volumes from the validation set to the test set. These volumes are authored by individuals who did not contribute to any works in the training set. Table 1 shows the dataset statistics after the split.

4 MangaVQA: A Novel Benchmark for Multimodal Context Understanding

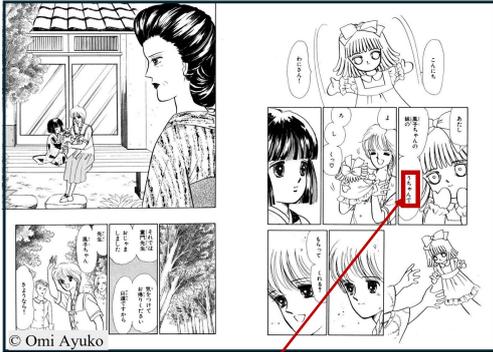
To evaluate model performance under realistic conditions, we manually constructed a set of question–answer (QA) pairs based on images from Manga109. Five annotators among the authors carefully and meticulously developed a high-quality evaluation set for MangaVQA, incorporating thorough human inspection and verification. To ensure robust and unambiguous evaluation, we focused on questions with definite answers, avoiding those that could be inferred merely from vague visual impressions. The 526 well-curated samples offer a comprehensive evaluation signal that sufficiently covers various aspects for reliable model comparison.

As shown in Figure 3, the question types are designed based on four key axes: (a) whether solving the question requires information from a single panel or multiple panels at the page level, (b) the answer type, distinguishing between exact extraction (word-level answers) and descriptive answering (sentence-level or explanatory answers), (c) 5W1H: whether the question asks about a person (who), an object or action (what), a time (when), a place (where), a reason (why), or a method or condition (how), and (d) inclusion of the author / title in the training split.

We illustrate examples along axis (b), the answer type, in Fig. 4. The categorization of (b) the answer type is as follows: **(1) Exact Extraction (240 questions): Questions that Require Extracting Answer Words from the Image.** These questions necessitate accurately retrieving the answer word from the manga page. We include one example in the left of Fig. 4. The question is “What is the name of the doll that Fuko-chan received?” and the answer is “Fu-chan”, which is directly written in

(1) Exact Extraction

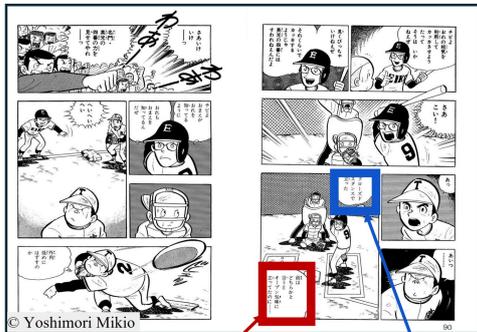
Q. 風子ちゃんがおもらったお人形の名前は何ですか？
(What is the name of the doll that Fuko-chan received?)



A. ふうちゃん (Fu-chan)

(2) Descriptive Answering

Q. 捕手は、打者のどのような変化に気づきましたか？
(What changes did the catcher notice in the batter?)



A. 以前はオープン気味に立っていたが、今はクローズドスタンドで立っている (He used to stand with an open stance, but now he stands with a closed stance.)

Figure 4: **Main categorization of MangaVQA: Answer type.** MangaVQA consists of (1) Exact Extraction, where the answer is directly extracted from the image; and (2) Descriptive Answering, where the answer requires explanatory or contextual responses beyond simple word extraction.

the dialogue. This category assesses the LMM’s basic comprehension ability to identify and extract the correct answer part from the manga panels.

(2) Descriptive Answering (286 questions): Questions that Require Contextual or Explanatory Responses. These questions go beyond simple answer word extraction and require comprehending the context within the manga. We include one example in the middle of Fig. 4. The question is “What changes did the catcher notice in the batter?”. The correct answer is “He used to stand with an open stance, but now he stands with a closed stance”. This category allows us to evaluate whether the LMM can not only recognize the dialogue but also understand its underlying meaning

in the context of the narrative.

5 MangaLMM: A Specialized Model for MangaOCR and MangaVQA

We develop MangaLMM, a specialized model designed to read and understand manga in a human-like manner. To build MangaLMM, we finetune the open-source LMM Qwen2.5-VL (Bai et al., 2025) on the MangaOCR and MangaVQA datasets, resulting in a joint model for both tasks. In this section, we describe the training data construction and training details for MangaLMM.

5.1 Training Data Construction

OCR Training Set T_{OCR} . For the OCR task, we use the MangaOCR training set, as described in §3.2. For each image, we format the sequence of text annotations as $\{“bbox_2d” : coordinates_i, “text_content” : text_i\}$, for each i , where $coordinates_i$ corresponds to the location of the text i in the image represented as $x_{top_left}, y_{top_left}, x_{bottom_right}, y_{bottom_right}$.

Synthetic VQA Training Set T_{VQA} . For the VQA task, we generate synthetic training data using GPT-4o (Hurst et al., 2024) (gpt-4o-2024-11-20). Following the synthetic data construction used in LLaVA (Liu et al., 2023), we generate five questions per image using both the image and its annotation from the OCR training set T_{OCR} . Here we exclude $< 0.1\%$ of the images where the text annotation is not included or GPT-4o refused to respond (e.g., due to violent content). Although we requested GPT-4o to generate five questions per image, it occasionally returned fewer than five. As a result, we created a total of 39,837 synthetic VQA samples from 8,379 images. The prompt used for question generation is shown in Table A in the Appendix. We plan to release this as a training split of our MangaVQA.

5.2 Training Details

LMM Selection. Our tasks require an open-source multilingual LMM that can handle Japanese and also has strong Japanese OCR capabilities, which are important for understanding manga. Several powerful multilingual LMMs have been proposed recently (Yue et al., 2025a; Wang et al., 2024; Bai et al., 2025; Maaz et al., 2024; Cohere Labs, 2025; Microsoft, 2025). Among them, the Qwen series (Wang et al., 2024; Bai et al., 2025) and Phi-4 (Microsoft, 2025) are especially notable for their

Japanese OCR performance. In this work, we build MangaLMM based on Qwen2.5-VL (Bai et al., 2025), which is one of the strongest open-source models in this category.

Training Strategy. We perform continual finetuning on both T_{OCR} and T_{VQA} using the pretrained Qwen2.5-VL 7B (Qwen2.5-VL-7B-Instruct). Most hyperparameters follow the original Qwen2.5-VL configuration, with a few modifications. For Manga109 images (1654×1170 resolution), we follow Qwen2.5-VL’s image resizing mechanism, which is based on pixel count thresholds, where the minimum and maximum number of input pixels are 3,136 and 2,116,800, respectively. We train for one epoch with a batch size of 32.

Elapsed Time for Training. Each dataset is trained for one epoch. Training Qwen2.5-VL 7B using four NVIDIA A100 GPUs took about 1 hour when using T_{OCR} or T_{VQA} , and about 2 hours when using both T_{OCR} and T_{VQA} .

6 Experiments

Evaluation Protocol for MangaOCR. We follow the evaluation protocols from prior OCR studies (Ye et al., 2023b; Huang et al., 2023) and IC-DAR 2019 multilingual OCR competitions (Chng et al., 2019; Zhang et al., 2019; Sun et al., 2019; Nayef et al., 2019). First, a predicted bounding box is considered a correct detection if its intersection over union (IoU) with a ground truth box exceeds 0.5. Based on the matched boxes, we compute precision (P), recall (R), and the harmonic mean (Hmean). Second, for each matched box, we calculate the normalized edit distance (NED) between the predicted and ground truth texts as a character-level metric. NED ranges from 0 to 1, with higher values indicating better performance; details are in the supplementary materials.

Since LMMs sometimes output the same word repeatedly, we apply post-processing to exclude repeated text segments that appear more than ten times, treating them as noise. Except for the analysis in §6.2, we report only the end-to-end Hmean for simplicity.

Evaluation Protocol for MangaVQA. Following LLaVA-Bench (Liu et al., 2023), we adopt the LLM-as-a-judge approach (Zheng et al., 2024) as our evaluation metric. We select a single LLM as the judge, which assigns scores to model responses. Specifically, we provide the judge with the question, a human-written answer, and the model’s re-

Method	MangaOCR Hmean (%)	MangaVQA LLM (/10.0)
GPT-4o	0.0	6.00
Gemini 2.5 Flash	0.0	7.26
Claude Sonnet 4.5	0.0	5.84
Phi-4-Multimodal-5.6B	0.0	3.39
Pangea-7B	0.0	3.23
LLaVA-OV-1.5-8B	0.0	3.46
Sarashina2-Vision-8B	0.0	4.45
Gemma-3-12B	0.0	4.13
Heron-NVILA-Lite-15B	0.0	3.76
Qwen2.5-VL-7B	0.9	5.65
MangaLMM (Ours)	71.5	6.68

Table 2: Comparison of LMMs on MangaOCR and MangaVQA.

sponse. Based on the human-written answer, the judge evaluates whether the model’s response is appropriate and relevant to the question on a 1–10 scale. To avoid circular bias with GPT-4o, which was used in constructing our VQA training dataset, we employ Gemini 2.5 Flash (Google DeepMind, 2025) (gemini-2.5-flash) as the judge LLM. The prompt used for LLM-as-a-judge is shown in Table C in the Appendix.

LMMs Used for Comparison. We evaluate three proprietary LMMs, gpt-4o-2024-11-20 (Hurst et al., 2024), gemini-2.5-flash (Google DeepMind, 2025), and claude-sonnet-4-5-20250929 (Anthropic, 2025), and seven open-source LMMs with Japanese capability: Qwen2.5-VL-7B-Instruct (Bai et al., 2025), Phi-4-multimodal-instruct (Abouelenin et al., 2025), Pangea-7B (Yue et al., 2025b), LLaVA-OneVision-1.5-8B-Instruct (An et al., 2025), sarashina2-vision-8b (Intuitions, 2025), gemma-3-12b-it (Team et al., 2025), and Heron-NVILA-Lite-15B (Inc., 2025).

6.1 Main Results

Table 2 compares LMMs for both MangaOCR and MangaVQA tasks. Overall, MangaLMM can handle both tasks effectively: it achieves over 70% OCR score and shows competitive VQA performance. While it falls short of the proprietary model Gemini, it outperforms the other proprietary models GPT-4o and Claude Sonnet 4.5. MangaLMM achieves the best performance among the open-source models by a clear margin.

Analysis of Low Performance on MangaOCR. As shown in Table 2, all LMMs

FT data	MangaOCR	MangaVQA
	Hmean (%)	LLM (/10.0)
None	0.9	5.65
T _{OCR}	74.5 ± 1.3	1.20 ± 0.32
T _{VQA}	0.0 ± 0.0	6.57 ± 0.06
T _{OCR} +T _{VQA}	71.2 ± 0.6	6.68 ± 0.03

Table 3: **Effect of finetuning (FT)**. FT is performed on the OCR training set T_{OCR}, the VQA training set T_{VQA}, or both.

except MangaLMM show near-zero scores on the MangaOCR benchmark. Most of their predictions consist of meaningless repetitions or short repeated tokens. The extremely low OCR score before finetuning is likely due to two main factors: (1) these models are not familiar with manga data, and (2) their weak detection capabilities may limit OCR performance. Prior work (Wu et al., 2024) has shown that GPT-4o, for example, exhibits poor detection ability, which may also apply to the other models.

Despite the near-zero OCR score—where not only position information is missing but even the correct text content is not generated—these models still manage to answer certain VQA questions that require interpreting text within the image. This is somewhat *counterintuitive*. Although the models fail to explicitly output the correct OCR results, they appear to capture some textual semantics from the image. This suggests that they are able to extract relevant information needed for answering VQA questions, even without performing OCR correctly. We provide more analysis and discussion in Appendices E.3 and A.

Analysis of the Effect of Finetuning. Table 3 shows the effect of finetuning. Each experiment was conducted three times, and the mean and standard deviation are reported. Finetuning Qwen2.5-VL on T_{OCR} and T_{VQA} enables the model to specialize in each respective task. On MangaOCR, the finetuned model achieves a significant improvement to a score of 74.5±1.3, which we discuss further in §6.2. On MangaVQA, the model, initially underperforming GPT-4o, surpasses it after finetuning. These results highlight the effectiveness of our synthetic VQA training set T_{VQA}, which we further analyze in §6.3.

Analysis from the Perspective of Task Interference. MangaLMM, a Qwen2.5-VL model fine-

Stage	Prec.	Recall	Hmean
Detection	82.2	75.3	78.6
End-to-end	74.8	68.5	71.5

Table 4: MangaLMM’s detection and end-to-end performance on MangaOCR.

tuned jointly on both T_{OCR} and T_{VQA}, shows a slight drop in OCR performance compared to using T_{OCR} alone, but achieves a small gain in VQA score over using T_{VQA} alone. A common issue in multi-task learning is *task interference* (Maninis et al., 2019; Yu et al., 2020; Ding et al., 2023; Chen et al., 2024), where models jointly trained on multiple tasks (e.g., *A* and *B*) often perform worse on task *A* than models trained solely on *A*. Under this assumption, one might expect the VQA performance of a jointly trained OCR+VQA model to degrade relative to a VQA-only model. Interestingly, we instead observe a slight improvement in VQA performance under joint training, contrary to typical interference expectations. This suggests that although task interference may exist, the enhanced OCR capability likely provides helpful textual cues that marginally improve VQA performance.

6.2 Performance Analysis of MangaOCR

Table 4 presents MangaOCR performance at both the detection and end-to-end stages. The Hmean of detection is 78.6%, while that of end-to-end reaches 71.5%, implying that once text regions are detected, the model reads them with approximately 91.0% ($=71.5 / 78.6$) accuracy. Some false positives occur when the model predicts text that actually appears in the manga but is not included in the annotations—for instance, page numbers or editorial marks outside the narrative content such as dialogue or onomatopoeia. Consequently, precision is unlikely to reach 100%. In contrast, recall is relatively low (68.5%), suggesting that about 31.5% of the ground-truth narrative text remains undetected, leaving room for improvement in capturing all semantically relevant content.

6.3 Performance Analysis of MangaVQA

Category-wise VQA Performance. Figure 5 presents a breakdown of model performance across the annotated categories in MangaVQA. We observe consistent performance gains across all tags, indicating that our training contributes to stable improvement in VQA capability without favoring

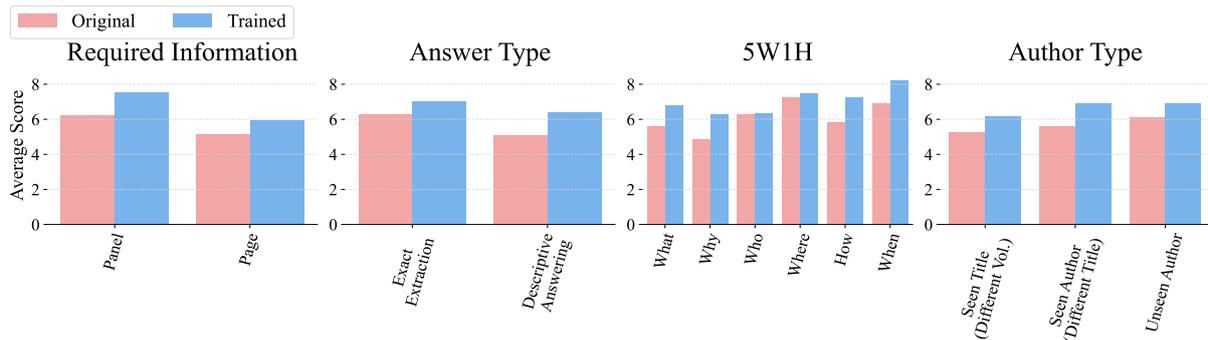


Figure 5: **Category-wise score breakdown.** Compared to the original model (Qwen2.5-VL-7B-Instruct), our trained MangaLMM improves scores across every tag in every category.

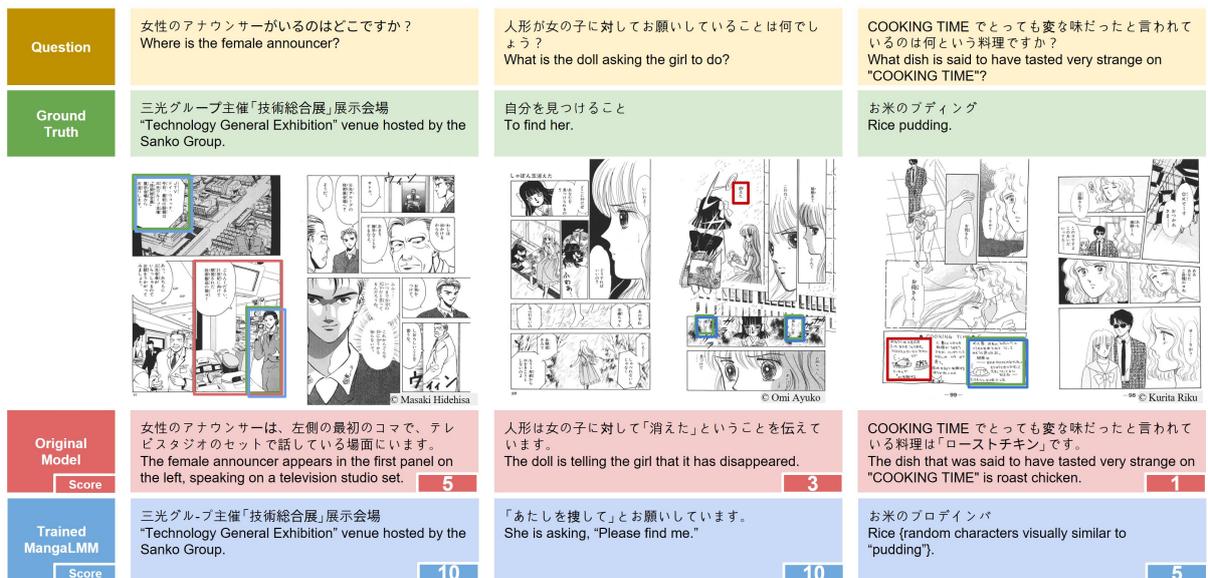


Figure 6: **Qualitative analysis on MangaVQA.** Regions relevant to the question or models' answer are highlighted with boxes in corresponding colors. In the left and middle examples, performance improves significantly after training, whereas in the right example, the trained model still struggles to produce an accurate answer.

OCR Annot.	LLM (/10.0)
	5.64
✓	6.68

Table 5: Effect of OCR Annotation on VQA Generation.

specific categories. Interestingly, the model also generalizes well to questions from unseen authors.

Effect of OCR Annotation when Generating VQA Data. When generating synthetic QA pairs for training, we include the OCR annotations in the prompt provided to GPT-4o. To examine their impact, we compare VQA data generated with and without text information. As shown in Table 5, the model trained on VQA data generated without OCR annotations achieves a score of 5.64, which does not exceed GPT-4o's own performance (6.00).

In contrast, using OCR-guided VQA data significantly improves the score to 6.68, even surpassing GPT-4o. These findings suggest that incorporating OCR annotations helps GPT-4o generate higher-quality QA pairs that enable the trained model to surpass GPT-4o's own performance.

Qualitative Analysis of MangaVQA. In Figure 6, we compare outputs of the original Qwen model and our trained model: **Left:** The original model provides a general answer, whereas the trained model leverages text-bubble content to produce a more specific one, improving the score (5 → 10). **Middle:** The original model extracts irrelevant text, while the trained model identifies the correct text, yielding a higher score (3 → 10). **Right:** The original model outputs an unrelated dish name, while the trained model identifies the correct one but makes character-level recognition

errors, yielding a partial score increase (1 → 5).

7 Conclusion

We presented MangaVQA, a benchmark for evaluating how well LMMs can understand manga in a human-like manner through contextual visual question answering, and MangaOCR, a consolidated benchmark for in-page text recognition. Together, they cover both the textual and narrative aspects of multimodal manga understanding. To establish a strong baseline, we developed MangaLMM, a specialized model jointly finetuned on OCR and VQA tasks. Experiments show that even state-of-the-art proprietary LMMs struggle with the unique complexity of manga, while MangaLMM performs well across both tasks. By releasing open benchmarks, synthetic data, and a strong open-source baseline, we aim to foster future research on multimodal manga understanding. Our work provides a concrete example of building and evaluating context-aware, domain-specialized LMMs, serving as a practical reference for similar research in other domains.

Limitation

One limitation of our model lies in its relatively slow inference speed for OCR. LMMs are inherently slower than dedicated OCR models; for instance, processing 1,166 test images containing 25,651 text instances takes several hours on an A100 GPU. In contrast, a dedicated OCR model such as DeepSolo (Ye et al., 2023b), which runs at over 10 FPS, can complete the same task in about two minutes. This slowdown primarily results from the large number of output tokens and occasional repetition or looping in the generated outputs during inference.

Acknowledgments

This work was supported by JSPS KAKENHI Grant Number 24K23882 and by the NVIDIA Academic Grant Program. This research utilized NVIDIA Saturn Cloud.

References

Abdelrahman Abouelenin, Atabak Ashfaq, Adam Atkinson, Hany Awadalla, Nguyen Bach, Jianmin Bao, Alon Benhaim, Martin Cai, Vishrav Chaudhary, Congcong Chen, and 1 others. 2025. Phi-4-mini technical report: Compact yet powerful multimodal lan-

guage models via mixture-of-loras. *arXiv preprint arXiv:2503.01743*.

Kiyoharu Aizawa, Azuma Fujimoto, Atsushi Otsubo, Toru Ogawa, Yusuke Matsui, Koki Tsubota, and Hikaru Ikuta. 2020. Building a manga dataset “manga109” with annotations for multimedia applications. *IEEE MultiMedia*.

Xiang An, Yin Xie, Kaicheng Yang, Wenkang Zhang, Xiuwei Zhao, Zheng Cheng, Yirui Wang, Songcen Xu, Changrui Chen, Chunsheng Wu, and 1 others. 2025. Llava-onevision-1.5: Fully open framework for democratized multimodal training. *arXiv preprint arXiv:2509.23661*.

Anthropic. 2025. Introducing claude sonnet 4.5. <https://www.anthropic.com/news/claude-sonnet-4-5>. Accessed: 2025-10-05.

Jeonghun Baek, Yusuke Matsui, and Kiyoharu Aizawa. 2022. Coo: Comic onomatopoeia dataset for recognizing arbitrary or truncated texts. In *ECCV*.

Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibao Song, Kai Dang, Peng Wang, Shijie Wang, Jun Tang, Humen Zhong, Yuanzhi Zhu, Mingkun Yang, Zhaohai Li, Jianqiang Wan, Pengfei Wang, Wei Ding, Zheren Fu, Yiheng Xu, and 8 others. 2025. Qwen2.5-vl technical report. *arXiv preprint arXiv:2502.13923*.

Zeren Chen, Ziqin Wang, Zhen Wang, Huayang Liu, Zhenfei Yin, Si Liu, Lu Sheng, Wanli Ouyang, and Jing Shao. 2024. Octavius: Mitigating task interference in MLLMs via loRA-moe. In *ICLR*.

Chee Kheng Chng, Yuliang Liu, Yipeng Sun, Chun Chet Ng, Canjie Luo, Zihan Ni, ChuanMing Fang, Shuaitao Zhang, Junyu Han, Errui Ding, and 1 others. 2019. Icdar2019 robust reading challenge on arbitrary-shaped text-rrc-art. In *ICDAR*.

Cohere Labs. 2025. Aya vision 8b. <https://huggingface.co/CohereLabs/aya-vision-8b>. Accessed: 2025-10-04.

Chuntao Ding, Zhichao Lu, Shangguang Wang, Ran Cheng, and Vishnu Naresh Boddeti. 2023. Mitigating task interference in multi-task learning via explicit task routing with non-learnable primitives. In *CVPR*.

Google DeepMind. 2025. Gemini 2.5: Our most intelligent ai model. <https://blog.google/technology/google-deepmind/gemini-model-thinking-updates-march-2025/>. Accessed: 2025-10-04.

Clément Guérin, Christophe Rigaud, Antoine Mercier, Farid Ammar-Boudjelal, Karell Bertet, Alain Bouju, Jean-Christophe Burie, Georges Louis, Jean-Marc Ogier, and Arnaud Revel. 2013. ebdtheque: a representative database of comics. In *ICDAR*.

- Mingxin Huang, Jiaxin Zhang, Dezhi Peng, Hao Lu, Can Huang, Yuliang Liu, Xiang Bai, and Lianwen Jin. 2023. Estextspotter: Towards better scene text spotting with explicit synergy in transformer. In *ICCV*.
- Aaron Hurst, Adam Lerer, Adam P Goucher, Adam Perelman, Aditya Ramesh, Aidan Clark, AJ Ostrow, Akila Welihinda, Alan Hayes, Alec Radford, and 1 others. 2024. Gpt-4o system card. *arXiv preprint arXiv:2410.21276*.
- Hikaru Ikuta, Leslie Wohler, and Kiyoharu Aizawa. 2025. Mangaub: A manga understanding benchmark for large multimodal models. *IEEE MultiMedia*.
- Turing Inc. 2025. Heron-nvila-lite-15b. <https://huggingface.co/turing-motors/Heron-NVILA-Lite-15B>. Accessed: 2025-10-04.
- SB Intuitions. 2025. Sarashina2-vision-8b. <https://huggingface.co/sbintuitions/sarashina2-vision-8b>. Accessed: 2025-10-04.
- Mohit Iyyer, Varun Manjunatha, Anupam Guha, Yogarshi Vyas, Jordan Boyd-Graber, Hal Daume, and Larry S Davis. 2017. The amazing mysteries of the gutter: Drawing inferences between panels in comic book narratives. In *CVPR*.
- Bo Li, Yuanhan Zhang, Dong Guo, Renrui Zhang, Feng Li, Hao Zhang, Kaichen Zhang, Yanwei Li, Ziwei Liu, and Chunyuan Li. 2024a. Llava-onevision: Easy visual task transfer. *arXiv preprint arXiv:2408.03326*.
- Yingxuan Li, Ryota Hinami, Kiyoharu Aizawa, and Yusuke Matsui. 2024b. Zero-shot character identification and speaker prediction in comics via iterative multimodal fusion. In *ACMMM*.
- Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. 2023. Visual instruction tuning. In *NeurIPS*.
- Yuan Liu, Haodong Duan, Yuanhan Zhang, Bo Li, Songyang Zhang, Wangbo Zhao, Yike Yuan, Jiaqi Wang, Conghui He, Ziwei Liu, and 1 others. 2024a. Mmbench: Is your multi-modal model an all-around player? In *ECCV*.
- Yuliang Liu, Biao Yang, Qiang Liu, Zhang Li, Zhiyin Ma, Shuo Zhang, and Xiang Bai. 2024b. Textmonkey: An ocr-free large multimodal model for understanding document. *arXiv preprint arXiv:2403.04473*.
- Muhammad Maaz, Hanoona Rasheed, Abdelrahman Shaker, Salman Khan, Hisham Cholakkal, Rao M Anwer, Tim Baldwin, Michael Felsberg, and Fahad S Khan. 2024. Palo: A polyglot large multimodal model for 5b people. In *WACV*.
- Kevis-Kokitsi Maninis, Ilija Radosavovic, and Iasonas Kokkinos. 2019. Attentive single-tasking of multiple tasks. In *CVPR*.
- Minesh Mathew, Dimosthenis Karatzas, and CV Jawahar. 2021. Docvqa: A dataset for vqa on document images. In *WACV*.
- Yusuke Matsui, Kota Ito, Yuji Aramaki, Azuma Fujimoto, Toru Ogawa, Toshihiko Yamasaki, and Kiyoharu Aizawa. 2017. Sketch-based manga retrieval using manga109 dataset. *MTAP*.
- Microsoft. 2025. Phi-4-multimodal-instruct. <https://huggingface.co/microsoft/Phi-4-multimodal-instruct>. Accessed: 2025-10-04.
- Nibal Nayef, Yash Patel, Michal Busta, Pinaki Nath Chowdhury, Dimosthenis Karatzas, Wafa Khelif, Jiri Matas, Umapada Pal, Jean-Christophe Burie, Chenglin Liu, and 1 others. 2019. Icdar2019 robust reading challenge on multi-lingual scene text detection and recognition—rrc-mlt-2019. In *ICDAR*.
- Nhu-Van Nguyen, Christophe Rigaud, and Jean-Christophe Burie. 2018. Digital comics image indexing based on deep learning. *J. Imaging*.
- Ragav Sachdeva, Gyungin Shin, and Andrew Zisserman. 2024. Tails tell tales: Chapter-wide manga transcriptions with character names. In *ACCV*.
- Ragav Sachdeva and Andrew Zisserman. 2024. The manga whisperer: Automatically generating transcriptions for comics. In *CVPR*.
- Ragav Sachdeva and Andrew Zisserman. 2025. From panels to prose: Generating literary narratives from comics. *arXiv preprint arXiv:2503.23344*.
- 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 *CVPR*.
- Gürkan Soykan, Deniz Yuret, and Tevfik Metin Sezgin. 2024. A comprehensive gold standard and benchmark for comics text detection and recognition. In *ICDAR*.
- Yipeng Sun, Zihan Ni, Chee-Kheng Chng, Yuliang Liu, Canjie Luo, Chun Chet Ng, Junyu Han, Errui Ding, Jingtuo Liu, Dimosthenis Karatzas, and 1 others. 2019. Icdar 2019 competition on large-scale street view text with partial labeling-rrc-lsvt. In *ICDAR*.
- Gemma Team, Aishwarya Kamath, Johan Ferret, Shreya Pathak, Nino Vieillard, Ramona Merhej, Sarah Perrin, Tatiana Matejovicova, Alexandre Ramé, Morgane Rivière, and 1 others. 2025. Gemma 3 technical report. *arXiv preprint arXiv:2503.19786*.
- Emanuele Vivoli, Marco Bertini, and Dimosthenis Karatzas. 2024. Comix: A comprehensive benchmark for multi-task comic understanding. In *NeurIPS*.

Peng Wang, Shuai Bai, Sinan Tan, Shijie Wang, Zhihao Fan, Jinze Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, and 1 others. 2024. Qwen2-vl: Enhancing vision-language model’s perception of the world at any resolution. *arXiv preprint arXiv:2409.12191*.

Yixuan Wu, Yizhou Wang, Shixiang Tang, Wenhao Wu, Tong He, Wanli Ouyang, Jian Wu, and Philip Torr. 2024. Dettolchain: A new prompting paradigm to unleash detection ability of mllm. In *ECCV*.

Zhibo Yang, Jun Tang, Zhaohai Li, Pengfei Wang, Jianqiang Wan, Humen Zhong, Xuejing Liu, Mingkun Yang, Peng Wang, Shuai Bai, and 1 others. 2024. Cc-ocr: A comprehensive and challenging ocr benchmark for evaluating large multimodal models in literacy. *arXiv preprint arXiv:2412.02210*.

Maoyuan Ye, Jing Zhang, Shanshan Zhao, Juhua Liu, Tongliang Liu, Bo Du, and Dacheng Tao. 2023a. DeepSolo++: Let transformer decoder with explicit points solo for multilingual text spotting. *arXiv preprint arXiv:2305.19957*.

Maoyuan Ye, Jing Zhang, Shanshan Zhao, Juhua Liu, Tongliang Liu, Bo Du, and Dacheng Tao. 2023b. DeepSolo: Let transformer decoder with explicit points solo for text spotting. In *CVPR*.

Tianhe Yu, Saurabh Kumar, Abhishek Gupta, Sergey Levine, Karol Hausman, and Chelsea Finn. 2020. Gradient surgery for multi-task learning. In *NeurIPS*.

Xiang Yue, Yuansheng Ni, Kai Zhang, Tianyu Zheng, Ruoqi Liu, Ge Zhang, Samuel Stevens, Dongfu Jiang, Weiming Ren, Yuxuan Sun, and 1 others. 2024. Mmmu: A massive multi-discipline multimodal understanding and reasoning benchmark for expert agi. In *CVPR*.

Xiang Yue, Yueqi Song, Akari Asai, Seungone Kim, Jean de Dieu Nyandwi, Simran Khanuja, Anjali Kantaruban, Lintang Sutawika, Sathyanarayanan Ramamoorthy, and Graham Neubig. 2025a. Pangea: A fully open multilingual multimodal llm for 39 languages. In *ICLR*.

Xiang Yue, Yueqi Song, Akari Asai, Seungone Kim, Jean de Dieu Nyandwi, Simran Khanuja, Anjali Kantaruban, Lintang Sutawika, Sathyanarayanan Ramamoorthy, and Graham Neubig. 2025b. Pangea: A fully open multilingual multimodal LLM for 39 languages. In *ICLR*.

Rui Zhang, Yongsheng Zhou, Qianyi Jiang, Qi Song, Nan Li, Kai Zhou, Lei Wang, Dong Wang, Minghui Liao, Mingkun Yang, and 1 others. 2019. Icdar 2019 robust reading challenge on reading chinese text on signboard. In *ICDAR*.

Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric Xing, and 1 others. 2024. Judging llm-as-a-judge with mt-bench and chatbot arena. In *NeurIPS*.

In this supplementary material, we provide additional details including (A) OCR evaluation in comics, (B) synthetic VQA examples, (C) additional MangaVQA examples, (D) setup details, and (E) additional results.

A OCR Evaluation in Comics

As described in §2, the evaluation of OCR has often been underexplored. Recent works such as Magi (Sachdeva and Zisserman, 2024) and CoMix (Vivoli et al., 2024) focus on transcription generation, which inherently includes OCR as a core component. CoMix, in particular, proposes a dedicated metric called the Hybrid Dialog Score for evaluating transcription tasks. However, this transcription-focused evaluation differs from direct OCR evaluation, which aims to assess whether the model accurately reads the text. First, transcription involves multiple subtasks beyond text detection and recognition, such as speaker identification, reading order prediction, and others. The quality of the final transcription output depends on the combined performance of these components, making it difficult to isolate and measure the accuracy of text recognition alone.

Second, transcription-based evaluations do not assess the positional accuracy of recognized text. Spatial information plays a crucial role in OCR, especially when the same text appears in multiple locations, as it helps identify which text instance is correct. For example, in Figure G(a), the word “わあー (waa-)” appears in four different locations, only one of which is correct. Without positional information, it becomes impossible to identify the correct instance. Moreover, spatial information is crucial for content understanding, as the interpretation of the same text can vary significantly depending on its location.

A proper evaluation of OCR in the manga domain allows us to better understand how well current LMMs can recognize text within manga. As described in the results section (§6.1), models such as GPT-4o exhibit near-zero OCR performance, yet are still able to answer VQA questions that rely on textual information. This result suggests that LMMs may be partially recognizing some text in the image. Our visualization of GPT-4o’s OCR output reveals that the detected text regions almost always appear in nonsensical locations, yet the model can still read certain parts of the text in the image. We provide a detailed analysis of this observation

in §E.3.

B Synthetic VQA Examples

For training our MangaLMM, we rely on synthetic VQAs generated by GPT-4o. In Figure A, we provide examples of these generated VQAs. As illustrated in the figure, GPT-4o is capable of producing accurate and diverse question–answer pairs.

We emphasize once again that providing GPT with text annotations is crucial for generating such high-quality VQAs. Without these annotations, GPT tends to produce unreliable outputs (e.g., misspelled extractions and factually incorrect questions) which significantly limit the performance of the MangaLMM trained on such data, as discussed in §6.3.

Human Validation. To validate the reliability of the synthetic VQA data generated by GPT, we conducted a manual evaluation. We randomly sampled 500 question–answer pairs and asked four human evaluators to assign scores to each item on a three-level scale: 0 (incorrect), 0.5 (partially correct), and 1 (correct). The average score is 0.80, suggesting that approximately 80% of the synthetic VQAs are judged to be appropriate by humans.

C Additional MangaVQA Examples

In this section, we provide additional examples from MangaVQA to illustrate how questions are categorized according to different aspects. Figure B shows examples categorized by required information, indicating whether the question can be answered from a single panel or requires multiple panels at the page level. Figure C presents examples categorized by 5W1H, showing the diversity of question types such as who, what, when, where, why, and how. These examples highlight the variety of question types and contextual understanding needed to answer MangaVQA questions.

Title: "AppareKappore"

© Hiroshi Kanno

Generated VQAs

Q1) 森魚は何を言いながら家を出発しましたか？
What did Morio say as he left the house?
A) 「学校見てくる」と言いながら家を出発しました。He said, "I'm going to check out the school," as he left the house.

Q2) 母親が森魚に指摘したことは何ですか？
What did his mother point out to Morio?
A) 「荷をほどきなさい、全然かたづかないじゃないの」と指摘しました。She pointed out, "Unpack your things, nothing's getting cleaned up this way!"

Q3) 森魚は学校までの道で何に戸惑っていますか？
What confused Morio on his way to school?
A) 同じ建物ばかりで方向がわからなくなり、迷っています。As all the buildings looked the same, he lost his sense of direction and got confused.

Q4) 森魚が最後に道を確信したきっかけは何ですか？
What helped Morio finally feel sure about the way?
A) 引っ越しの車が来るときの方向を思い出したことです。He remembered the direction from which the moving truck had come.

Q5) 母親が森魚の態度について言ったことは何ですか？
What did Morio's mother say about his attitude?
A) 「もう誰に似たんだか」と言いました。She said, "Good grief, who did you take after?"

(a) An example from the manga titled AppareKappore.

Title: "Garakutaya Manta"

© Tatsuki Nouda

Generated VQAs

Q1) この物語の登場人物たちは、何に向けて準備をしていますか？
What are the characters in this story preparing for?
A) 明日の試験に向けて準備しています。They are preparing for tomorrow's exam.

Q2) 主人公らが試験に不安を抱えるきっかけは何ですか？
What caused the main characters to feel anxious about the exam?
A) スマイリー若松が試験の前日に試験のことを告げたためです。Smiley Wakamatsu told them about the exam the day before it was to take place.

Q3) 主人公が外国語を勉強するのに納得できない理由は何ですか？
Why does the main character feel it's unreasonable to study a foreign language?
A) 日本人である自分が外国語を学ぶ必要がないと考えているからです。Because he believes that, as a Japanese person, there's no need for him to learn a foreign language.

Q4) 主人公は「国際相互理解」のためにどんな行動を取ることを提案しましたか？
What action did the main character suggest for promoting "international mutual understanding"?
A) 外国人の心を理解するために外国産ゲームをすることを提案しました。He suggested playing foreign-made video games to understand the hearts of foreigners.

Q5) 最後に登場人物が連れてきた「秘密兵器」とは何ですか？
What is the "secret weapon" that a character brings in at the end?
A) 具体的には明かされていませんが、主人公の試験に対応するために連れてきた存在です。It's not specified, but it is someone brought to help the main character with the exam.

(b) An example from the manga titled GarakutayaManta.

Figure A: Examples of synthetic VQA generation results. The most relevant part of the image for each question-answer pair is highlighted and translated in the corresponding color.

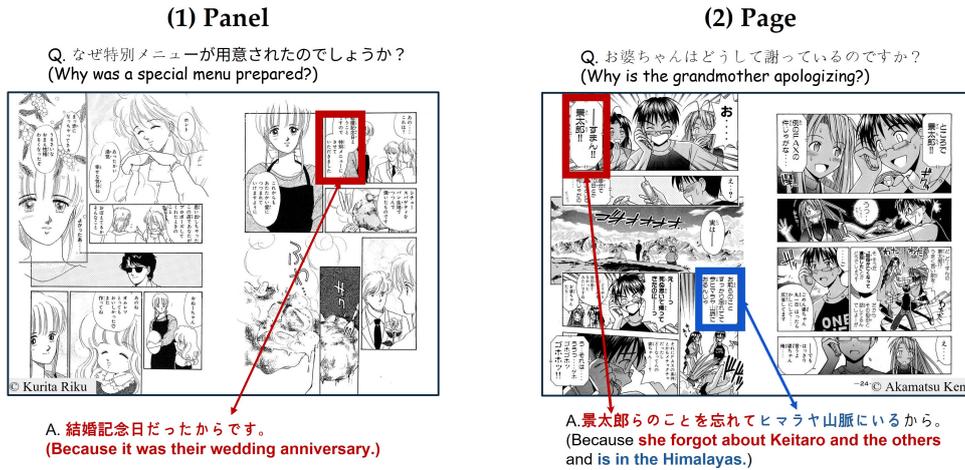


Figure B: **Categorization of MangaVQA: Required information.** MangaVQA questions are categorized by whether solving the question requires information from (1) a single panel or (2) multiple panels at the page level.

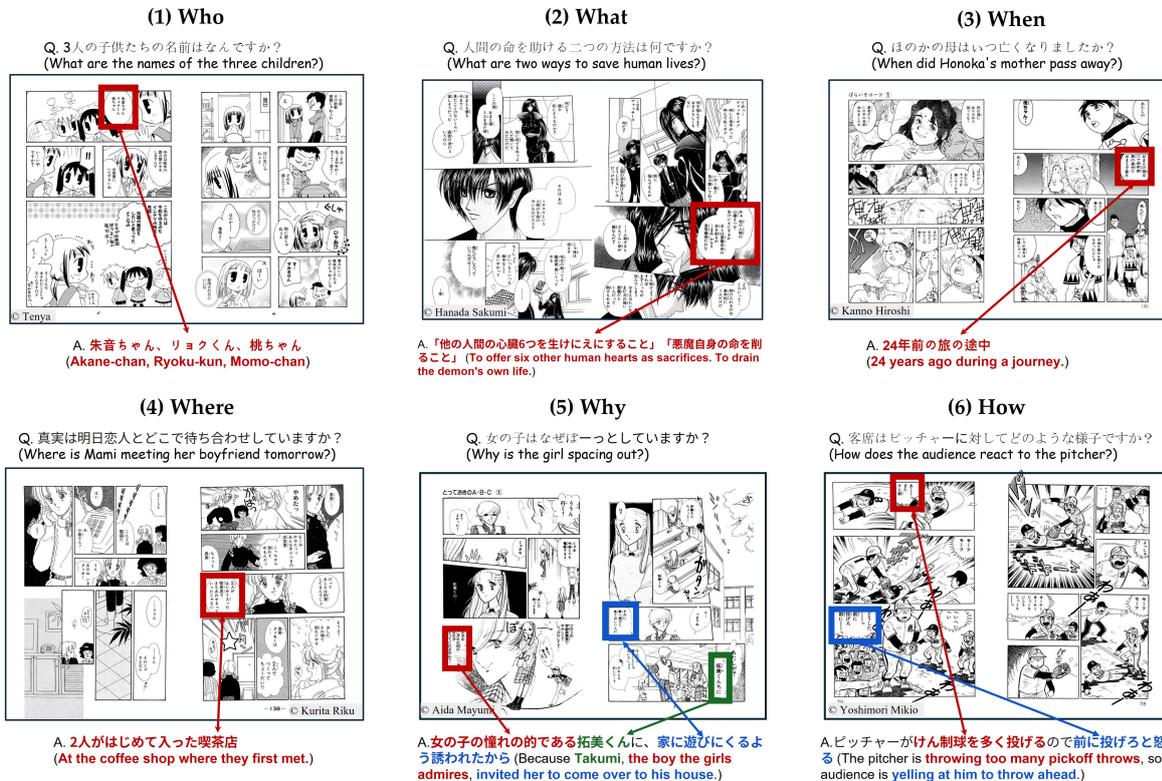


Figure C: **Categorization of MangaVQA: 5W1H.** MangaVQA questions are also categorized by 5W1H, that is, whether the question asks about a person (who), an object or action (what), a time (when), a place (where), a reason (why), or a method or condition (how).

D Setup Details

Evaluation Metric. We provide a detailed description of the normalized edit distance (NED, also referred to as 1-NED), which was used as the evaluation metric in MangaOCR. NED scales the standard edit distance to a range between 0 and 1, where higher values indicate better prediction. It is computed as follows:

$$\text{NED} = 1 - \sum_{i=1}^N \frac{\text{ED}(\text{GT}_i, \text{Pred}_i)}{\text{MaxLen}(\text{GT}_i, \text{Pred}_i)} \quad (1)$$

Here, GT_i and Pred_i denote the i -th ground truth and the model’s prediction, respectively. $\text{ED}(\cdot)$ calculates the edit distance between two strings, and $\text{MaxLen}(\cdot)$ returns the longer of the two string lengths. N indicates the total number of text instances.

D.1 Prompt

Prompt for Synthetic VQA Generation. For creating synthetic QA pairs for training, we provide GPT-4o with the prompt in Table A along with the corresponding image.

Prompt for Training and Evaluation. For training and inference, we use task-specific prompts. For the MangaOCR benchmark, we provide the prompt “Please perform OCR on this image and output the recognized Japanese text along with its position (grounding)” along with the input image. During training, the corresponding OCR annotations are included as supervision. When running OCR inference with models other than the Qwen2.5 VL series, the outputs varied in format unless explicitly specified. Therefore, we use the prompt in Table B to align their outputs with the OCR format used in the training data of MangaOCR.

For the MangaVQA benchmark, we use the prompt “あなたは日本語の漫画に関する質問に答えるAIです。与えられた画像に基づいて質問に答えてください。(You are an AI that answers questions about Japanese manga. Please answer the given question based on the provided image.)” together with the input image and a question. The ground-truth answer is given only during training. For MangaVQA evaluation, the prompt in Table C is used for LLM-as-a-judge.

Original Japanese

与えられる画像と、そこに書かれている文字情報を用いて、
 質問: [質問内容]
 回答: [回答内容]
 質問: [質問内容]
 回答: [回答内容]
 ...
 の形式でVQA問題を5問作ってください。解釈が曖昧になる主観的な問題ではなく、書かれている事実に基づいて客観的に判断できる問題を作ってください。またOCRのような文字の読み取り問題にはせず、内容理解を問う問題を作ってください。
 画像内の文字:
 {OCR ANNOTATION HERE}

Translated

Using the given image and the textual information written in it, create 5 VQA questions in the following format:
 Question: [Question content]
 Answer: [Answer content]
 Question: [Question content]
 Answer: [Answer content]
 ...
 Avoid subjective questions that could lead to ambiguous interpretations, and instead create questions that can be objectively answered based on the facts presented in the image. Also, do not include OCR-style text recognition questions; instead, create questions that test understanding of the image content.
 Text in the image:
 {OCR ANNOTATION HERE}

Table A: Prompt for the synthetic VQA generation.

Please perform OCR on this image and output the recognized Japanese text along with its position (grounding).

The output should be a JSON list. Each item in the list must follow the structure below:

```
\n{"bbox_2d": [x1, y1, x2, y2], "text_content": "..."}\n
```

The field "bbox_2d" must be a 2D bounding box that tightly encloses the text.

Use the format '[x1, y1, x2, y2]', where:

- 'x1', 'y1' are the coordinates of the top-left corner of the bounding box, and
- 'x2', 'y2' are the coordinates of the bottom-right corner.

Here is an example of the desired format:

```
\n{"bbox_2d": [1490, 138, 1546, 201], "text_content": "春休みです-"}\n
```

Please follow this format strictly.

Table B: OCR inference prompt for models other than the Qwen2.5 VL series.

System message

You are an evaluator. Your task is to rate how appropriate a model's response is to a question about a manga image. For each case, you will be given a question (based on a manga image), a human-written answer, and the model's response. The image is not shown, but the question and answer are based on it. Please evaluate as if the image were available.

Please rate how well the model's response answers the question, considering the intended image context and the human answer as a reference, using a scale from 1 to 10:

- 1 — Completely inappropriate or unrelated to the question or image context.
 - 2 — Mostly unrelated with major misunderstandings or incorrect information.
 - 3 — Slightly relevant, but largely incorrect or unhelpful.
 - 4 — Somewhat relevant, but contains significant errors or omissions.
 - 5 — Partially correct with noticeable inaccuracies, vagueness, or missing key points.
 - 6 — Generally okay, but missing core points or includes some incorrect interpretations.
 - 7 — Mostly correct and relevant, with only minor issues or small omissions.
 - 8 — Almost entirely accurate with only slight room for improvement.
 - 9 — Very appropriate, accurate, and well-aligned with the question and image context.
 - 10 — Perfectly appropriate, accurate, and fully answers the question as if the image were visible.
- Only return a single number (1–10). Do not include any explanations, justifications, or comments.

User prompt

Input:

```
"question": {question},  
"human-written answer": {answer},  
"model's response": {generated_answer},
```

Your score:

Table C: Prompt for MangaVQA evaluation.

Size	FT data	MangaOCR Hmean (%)	MangaVQA LLM (/10.0)
3B	None	0.1	4.66
	T _{OCR}	73.5	1.97
	T _{VQA}	0.0	5.77
	T _{OCR} +T _{VQA}	66.5	6.01
7B	None	0.9	5.65
	T _{OCR}	74.9	1.03
	T _{VQA}	0.0	6.54
	T _{OCR} +T _{VQA}	71.5	6.68

Table D: Effect of model size (3B and 7B).

Ratio (%)	MangaOCR Hmean (%)	MangaVQA LLM (/10.0)
25	59.0	6.20
50	64.9	6.20
75	68.4	6.48
100	71.5	6.68

Table E: Effect of dataset size.

E Additional Results

We provide additional analysis and experimental results on our benchmarks, MangaVQA and MangaOCR.

E.1 Effect of Model and Dataset Size

Appendix E.1 shows the performance of Qwen2.5-VL models of different sizes (3B and 7B) under various finetuning settings. Similar to the 7B model, the 3B model shows a slight drop in MangaOCR performance when finetuned on both T_{OCR} and T_{VQA}, while its MangaVQA performance improves slightly. Appendix E.1 shows the results of varying dataset size (25%, 50%, 75%, and 100%). We observe that performance generally improves as the dataset size increases.

E.2 More Analysis of MangaVQA

Comparison with Human Evaluation. To validate the reliability and consistency of the Gemini-judge employed in the MangaVQA evaluation, we conducted a comparative analysis between its evaluation scores and those provided by human annotators. Specifically, we asked two human evaluators to assign scores to all items in the benchmark dataset, following the same evaluation prompt used for the Gemini-judge.

The results of this comparison are illustrated in Figure D. We observe a small absolute differ-

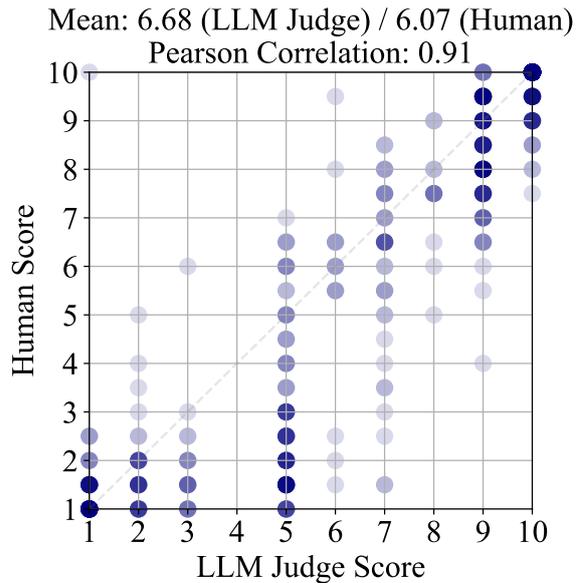


Figure D: Comparison between LLM-Judge (Gemini) and Human Evaluation. Darker points indicate a higher concentration of points.

ence in average scores ($\Delta = 0.61$). Additionally, there is a strong positive correlation between the scores assigned by the Gemini-judge and the human average ($r = 0.91$). These findings suggest that LLM-based evaluation can serve as a practical and consistent alternative to human judgment in our MangaVQA benchmark.

Comparison of Results Using GPT-4o as the Judge. Table F compares the results when Gemini 2.5 Flash and GPT-4o are used as the LLM judge. The overall trends remain consistent across both settings. Here, a potential circular bias may exist when the same LLM is used for both generating responses and judging them. In our case, the impact of such bias appears to be relatively minor. Specifically, the performance difference of Gemini 2.5 Flash between being judged by itself and by GPT-4o is only 0.12.

More Analysis of OCR Annotation when Generating VQA Data. As described in §6.3, OCR annotation plays a key role in generating high-quality QA pairs with GPT-4o. Here, we provide a more detailed analysis of the effect of OCR annotation. OCR annotation consists of both bounding box positions and their text content. We compare the synthetic VQA data generated by GPT-4o using only the text content with those generated using both bounding box positions and text content. Table G presents the results. Interestingly, our experiments show that using only the text content is more effective than including both text and positional

	Exact Extraction	Descriptive Answering	
Question	2人を待ちくたびれているのは誰だと言っていますか？ Who is said to be tired of waiting for the two of them?	グルメのお嬢さんは店の採点について何と言っていますか？ What does the young gourmet lady say about the restaurant's rating?	右側のページにおいて、男の子が怒っているのはどうしてですか？ On the right-hand page, why is the boy angry?
Ground Truth	サタン様 Lord Satan	まだ食べていないからわからないと言っています。 She says she doesn't know yet because she hasn't eaten there.	成瀬川が突然勝手に電話を切ったから。 Because Narusegawa suddenly hung up the phone without warning.
			
Original Model	漫画の内容から、2人を待ちくたびれているのは「いかげん」と言っています... Based on the manga's content, it says that "ikagen" is the one who is tired of waiting for the two...	グルメのお嬢さんは、「この店の採点は？」と尋ねています。 The young gourmet lady asks, "What's this restaurant's rating?"	右側のページでは、男の子が怒っている理由は、彼が何かを誤解している... On the right-hand page, the boy may be angry because he has misunderstood something...
Score	5	5	3
Trained MangaLMM	サタン様 Lord Satan	「まだ食べてないからわかりません」と言っています。 She says, "I don't know yet—I haven't eaten here!"	女の子が「思い出の女の子」について言及したため。 Because the girl mentioned the "girl from his memories".
Score	10	10	2

Figure E: **Category-wise analysis on MangaVQA**. The regions in the image relevant to the question or models' answer are highlighted with boxes in corresponding colors.

Method	Judge: Gemini	Judge: GPT-4o
GPT-4o	6.00	5.76
Gemini 2.5 Flash	7.26	7.14
Claude Sonnet 4.5	5.84	4.77
Phi-4-Multimodal-5.6B	3.39	3.08
Pangea-7B	3.23	2.98
LLaVA-OV-1.5-8B	3.46	3.15
Sarashina2-Vision-8B	4.45	4.13
Gemma-3-12B	4.13	3.47
Heron-NVILA-Lite-15B	3.76	3.32
Qwen2.5-VL 7B	5.65	5.36
MangaLMM (Ours)	6.68	6.57

Table F: **Comparison of LLM judge for MangaVQA**.

OCR Annot.	LLM (/10.0)
None	5.64
Text	6.68
Text + Pos.	6.46

Table G: **Effect of OCR Annotation on VQA Generation**.

information. Although our current approach did not benefit from positional information, leveraging it remains a promising direction for future work. Therefore, in our experiments, we use synthetic VQA examples generated using only the OCR text content.

Qualitative Analysis of MangaVQA. Figure E presents category-wise examples on MangaVQA.

For the categories on the left (Exact Extraction) and in the center (Descriptive Answering), the base Qwen 2.5-VL model often fails to locate the correct region and consequently extracts the wrong words. After finetuning, these issues are significantly improved in MangaLMM. On the other hand, for the category on the right (Descriptive Answering), MangaLMM tends to over-prioritise text extraction, leading to incorrect answers even after training.

E.3 More Analysis of MangaOCR

We present a qualitative analysis of MangaOCR results from GPT-4o and MangaLMM. As described in §6, text segments that appear more than ten times are considered noise and excluded from the results. Therefore, such repeated segments do not appear in the visualizations.

GPT-4o's Results on MangaOCR. Since previous studies have rarely conducted in-depth qualitative analysis of GPT-4o's OCR results, it is difficult to assess the model's actual performance on manga datasets. We address this gap by providing a detailed qualitative analysis of GPT-4o's MangaOCR outputs. Figure F shows GPT-4o's results on MangaOCR. These examples demonstrate the low zero-shot OCR performance of GPT-4o in the manga domain. The detected text regions almost always appear in incorrect or nonsensical locations, although the model can still read certain parts of the

text within the image. Because the predicted text positions are inaccurate, the outputs are considered entirely incorrect under OCR evaluation criteria. While some predicted text fragments correspond to actual text in the image, there are many cases—such as in Figure F(b)—where most of the text is not recognized at all. Even when text is recognized, it is often incorrect. While GPT-4o fails to correctly detect and recognize most of the text, it can still recognize partial text content, which may allow GPT-4o to answer some text-based VQA questions.

Interestingly, when performing OCR inference with GPT-4o, the model sometimes generates disclaimers such as: “The bounding box coordinates and text content are illustrative and may not perfectly match the actual image. For precise OCR and bounding box extraction, specialized OCR tools like Tesseract or Google Vision API should be used.” This suggests that GPT-4o itself acknowledges its limitations in precise OCR and recommends using dedicated OCR tools.

MangaLMM’s Results on MangaOCR. Figure G shows MangaLMM’s results on MangaOCR. As seen in the figure, most predictions appear correct, reflecting the model’s strong OCR capability across a wide range of text sizes, from large to small. The red regions indicate false negatives. Occasionally, even text that appears large and seemingly easy to detect is missed. According to our manual inspection, such cases are mostly onomatopoeia. This suggests that the model struggles more with onomatopoeic expressions, which are often written in non-standard fonts, sizes, or orientations, compared to regular text.

MangaOCR Evaluation without Positional Information. What if MangaOCR were evaluated without considering positional information? To further analyze the models’ text recognition ability, we evaluate them under a setting that does not depend on positional information. TextMonkey (Liu et al., 2024b) represents the first study to perform OCR evaluation without positional information, and several follow-up works (e.g., CC-OCR (Yang et al., 2024)) have adopted its evaluation approach. Following the evaluation method and code provided by TextMonkey, we adopt the so-called “Trans” mode, which ignores positional alignment and instead checks whether each ground-truth string appears anywhere within the predicted text.

In this evaluation, all predicted text strings from the image are concatenated into a single sequence, and each ground-truth instance is evaluated by

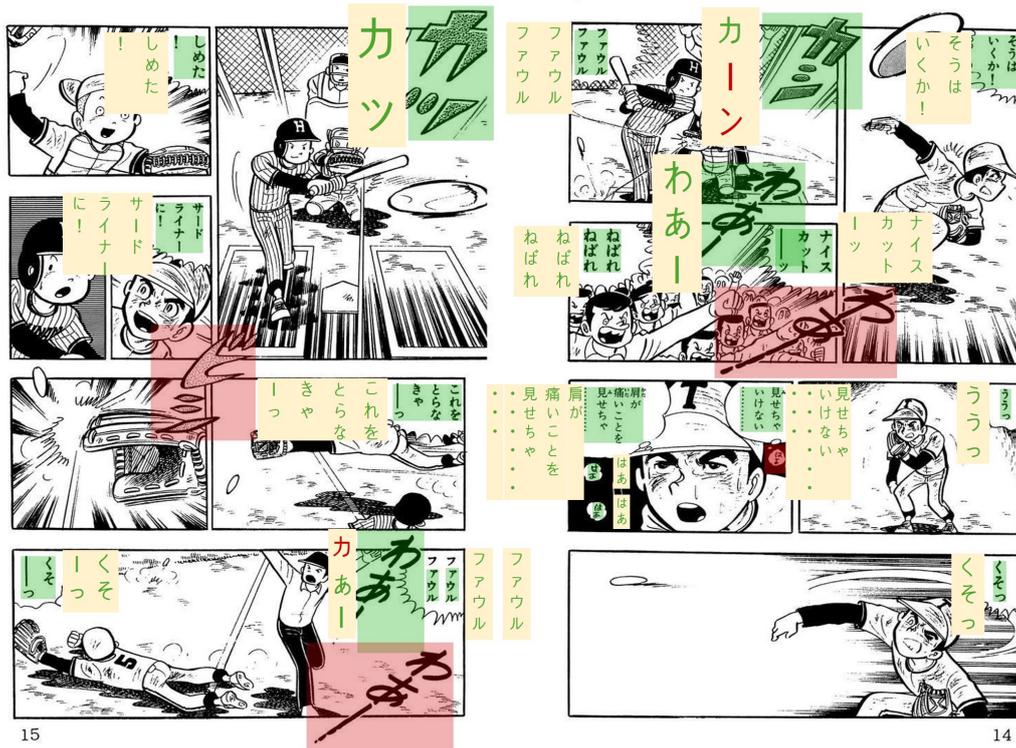
Method	Traditional OCR metric	TextMonkey-Trans
GPT-4o	0.0	12.9
Gemini 2.5 Flash	0.0	11.7
Claude Sonnet 4.5	0.0	7.3
Phi-4-Multimodal-5.6B	0.0	0.0
Pangea-7B	0.0	0.0
LLaVA-OV-1.5-8B	0.0	0.0
Sarashina2-Vision-8B	0.0	0.7
Gemma-3-12B	0.0	3.7
Heron-NVILA-Lite-15B	0.0	0.0
Qwen2.5-VL 7B	0.9	5.7
MangaLMM (Ours)	71.5	63.2

Table H: Model performance under two evaluation settings: the Traditional OCR metric and the TextMonkey “Trans” mode (ignoring positional information).

checking whether its text appears within the predicted string. Each ground-truth instance must be exactly and completely included in the prediction, and even a one-character mismatch results in an incorrect outcome. This stricter criterion generally yields lower scores than the edit distance–based evaluation, which assigns partial scores even when some characters in the text string are incorrect. For example, if the ground truth is “apple” and the concatenated prediction is “banana apple orange,” the instance is considered correct. However, if the concatenated prediction is “banana aple orange,” it is counted as incorrect due to the missing character.

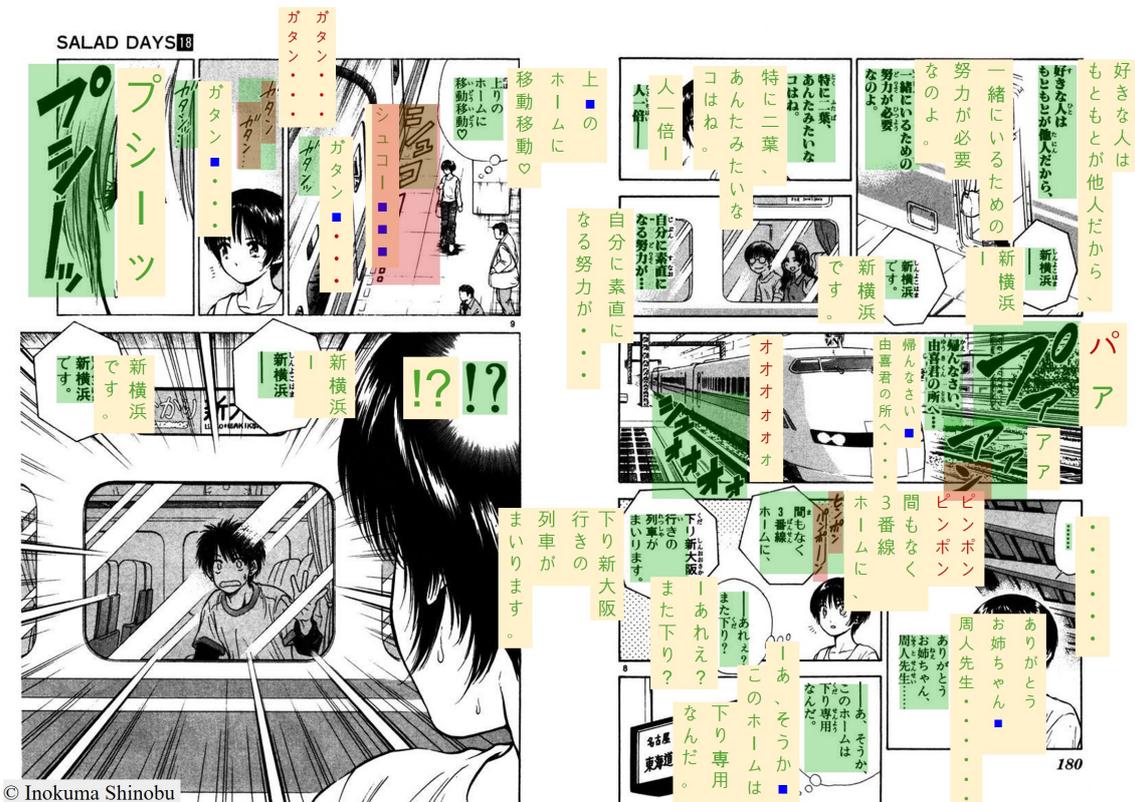
The results are summarized in Table H. Under this setting, several models such as GPT-4o, Gemini 2.5 Flash, Claude Sonnet 4.5, Gemma-3, and Qwen 2.5-VL achieved non-zero scores. However, models without finetuning still exhibited consistently low performance. This highlights the importance of finetuning on OCR-specific data.

Comparison with dedicated OCR system.. We evaluated DeepSolo++ (Ye et al., 2023a), a multilingual OCR system, on the MangaOCR task. The model achieved an Hmean of 5.4%, indicating very low performance in this domain. Since DeepSolo++ operates zero-shot on manga-style images, we observed that it often detects only small fragments of the text (e.g., one or two characters) rather than identifying the full text in the speech balloons, leading to many incorrect predictions. This behavior is consistent with what we observe in other LMMs when applied zero-shot to manga.



© Yoshimori Mikio

(a) An example from the manga titled Shimattelkouze.



© Inokuma Shinobu

(b) An example from the manga titled SaladDays.

Figure G: MangaLMM’s Results on MangaOCR. The green boxes indicate the detected text regions. The text shown near each green box is the predicted text for that detected region. The green text represents correctly predicted text, while the red text indicates incorrectly predicted text. Missing characters are marked with small blue squares. The red boxes show false negatives—text regions that should be detected but are missed. Most OCR results are correct.

Method	MangaOCR Hmean (%)	MangaVQA LLM (/10.0)	MMMU Acc. (%)	MMBench Acc. (%)
GPT-4o	0.0	6.00	70.7	89.0
Gemini 2.5 Flash	0.0	7.26	79.6	-
Phi-4-Multimodal-5.6B	0.0	3.39	55.1	86.7
Qwen2.5-VL 7B	0.9	5.65	58.6	87.8
MangaLMM	71.5	6.68	25.8	1.5

Table I: Comparison of LMMs on MangaOCR, MangaVQA, MMMU, and MMBench.

FT data	MangaOCR Hmean (%)	MangaVQA LLM (/10.0)	MMMU Acc. (%)	MMBench Acc. (%)
T _{OCR+TVQA} (MangaLMM)	71.5	6.68	25.8	1.5
T _{OCR+TVQA} + LO-50K	70.2	6.56	49.6	82.0

Table J: Finetuned results on a combined dataset including natural image understanding data. LO-50K denotes LLaVA Onevision 50K data.

E.4 Exploring the Capability for Manga and Natural Image Understanding

To further investigate the capability of LMMs across both manga and natural images, we conducted additional evaluations and experiments.

Evaluation on natural image benchmarks. We evaluated general-purpose baselines (GPT-4o, Gemini 2.5 Flash, Phi-4-Multimodal, and Qwen2.5-VL-7B) as well as our MangaLMM on two representative benchmarks commonly used for assessing LMM performance on natural images: MMMU (Yue et al., 2024) and MMBench (Liu et al., 2024a). The table below (an extended version of Table 2) shows the results. Scores for models other than MangaLMM are obtained from the official MMMU leaderboard or the Phi-4 technical report (Abouelenin et al., 2025). As expected, the general-domain performance of MangaLMM (finetuned from Qwen2.5-VL-7B) drops noticeably due to domain specialization, underscoring the challenge of maintaining broad visual understanding after task-specific finetuning.

Joint finetuning with natural image data. We also finetuned our model on a combined dataset consisting of manga data and 50K randomly sampled natural image examples from the LLaVA-OneVision dataset (Li et al., 2024a). As shown in the table below (an extended version of Table 3), this joint finetuning substantially restores the model’s performance on natural image benchmarks, while preserving strong performance on our MangaOCR and MangaVQA benchmarks.

These results demonstrate that although domain specialization may reduce general capability, it can

Method	Total (100 images)	JPN subset (6 images)	MangaOCR
Qwen2.5-VL 7B	2.8	0.3	0.9
MangaLMM	13.5	50.0	71.5

Table K: OCR evaluation on other comics from the eBDtheque dataset and MangaOCR.

be effectively recovered through joint training—highlighting the feasibility of developing LMMs that are both manga-capable and broadly applicable to natural image understanding.

Evaluation on Other Comics. We find that MangaLMM shows some degree of generalization to other comics, although its performance is naturally lower than on in-domain Japanese manga. As discussed in the related work, there is currently no established VQA dataset for the comic domain, making it difficult to use a standard VQA benchmark in this setting. To examine cross-domain performance, we therefore conducted an additional OCR evaluation on the well-known comic dataset eBDtheque (Guérin et al., 2013), which contains 100 images from European, American, and Japanese comics. Table K shows the results.

The performance of MangaLMM on eBDtheque is lower than on MangaOCR, reflecting the domain differences between the datasets. Notably, MangaLMM still outperforms Qwen2.5-VL, suggesting that training on manga contributes positively to OCR performance even in other comic styles. We attribute the lower performance on eBDtheque to two main factors: (1) Page format mismatch: MangaOCR consists of two-page spreads, whereas eBDtheque images are single-page. This mismatch frequently leads to shifted or misaligned predicted bounding boxes. (2) Color vs. grayscale: MangaOCR is entirely black-and-white, while more than half of eBDtheque images are in color. MangaLMM, trained on grayscale manga, shows degraded performance on colored pages. We expect both issues to be addressable by incorporating single-page layouts and color comic images during training. While MangaLMM is intentionally specialized for manga, extending it toward a general comic-capable LMM is a promising direction for future work.