

Global-Local Dual Perception for MLLMs in High-Resolution Text-Rich Image Translation

Junxin Lu¹, Tengfei Song², Zhanglin Wu², Pengfei Li², Xiaowei Liang², Hui Yang², Kun Chen², Ning Xie², Yunfei Lu², Jing Zhao³, Shiliang Sun⁴, and Daimeng Wei²

¹ School of Communication and Information Engineering, Shanghai University,
Shanghai 200444, China
junxinlu.ecnu@gmail.com

² Huawei Translation Services Center, China
songtf00@gmail.com, lipengfei203@h-partners.com
{luyunfeicx, moonlight-007}@163.com,
{wuzhanglin2, liangxiaowei2, yanghui141, chenkun49, nicolas.xie}@huawei.com

³ School of Computer Science and Technology, East China Normal University,
Shanghai 200062, China
jzhao2011@gmail.com

⁴ State Key Laboratory of Submarine Geoscience, School of Automation and
Intelligent Sensing, Shanghai Jiao Tong University, Shanghai 200240, China
shiliangsun@gmail.com *

Abstract. Text Image Machine Translation (TIMT) aims to translate text embedded in images in the source-language into target-language, requiring synergistic integration of visual perception and linguistic understanding. Existing TIMT methods, whether cascaded pipelines or end-to-end multimodal large language models (MLLMs), struggle with high-resolution text-rich images due to cluttered layouts, diverse fonts, and non-textual distractions, resulting in text omission, semantic drift, and contextual inconsistency. To address these challenges, we propose GLoTran, a global-local dual visual perception framework for MLLM-based TIMT. GLoTran integrates a low-resolution global image with multi-scale region-level text image slices through pre-local region translation replay and an instruction-guided alignment strategy, conditioning MLLMs to maintain scene-level contextual consistency while faithfully capturing fine-grained textual details. Moreover, to realize this dual-perception paradigm, we construct GLoD, a large-scale text-rich TIMT dataset comprising 510K high-resolution global-local image-text pairs covering diverse real-world scenarios. Extensive experiments demonstrate that GLoTran substantially improves translation completeness and accuracy over state-of-the-art MLLMs.

Keywords: Text Image Machine Translation · MLLMs · global-local dual visual perception

* Corresponding author: Shiliang Sun




Fig. 1: Comparison of (a) cascade methods, (b) traditional end-to-end models, (c) MLLMs-based models, and (d) our proposed GLoTran for TIMT. Through a dual visual perception strategy integrating global contextual understanding and local textual focus, we enable more complete and accurate TIMT.

1 Introduction

Text image machine translation (TIMT) [37,45] aims to translate source-language text embedded in images into a target-language, requiring a synergistic integration of visual perception and semantic understanding. Unlike conventional machine translation, which operates solely on text sequences, TIMT simultaneously handles text detection, recognition, and translation within complex visual contexts. Existing works primarily follow two paradigms: *cascade methods* [3,21,61], sequentially integrate optical character recognition (OCR) with neural machine translation, and *end-to-end models* [24,35,36,59], which jointly encode visual and linguistic modalities to directly map images into translated text. While both paradigms have driven progress, cascade methods suffer from error propagation and high computational cost, as illustrated in Fig.1(a), whereas conventional end-to-end models remain task-specific and struggle to generalize to diverse real-world scenarios, as shown in Fig.1(b).

Recently, multimodal large language models (MLLMs) have offered new avenues for TIMT. By integrating strong language priors with visual perception, MLLMs provide a unified framework for multimodal understanding, achieving promising results in image comprehension [14,32,39], document analysis [56,60], and visual question answering [13,47,62]. Inspired by these capabilities, recent studies have adapted MLLMs for TIMT through prompting and instruction tuning [19,30,58,66]. Despite impressive progress, current MLLM-based TIMT methods still struggle with high-resolution text-rich images. Real-world sources such as posters, screenshots, and documents often contain dense text interleaved with cluttered backgrounds, icons, and decorative elements, introducing non-textual distractions and producing an excessive number of visual tokens. This redundancy disperses attention from meaningful text regions, frequently causing local text omission, misrecognition, and fragmented semantic understanding, as illustrated in Fig.1(c). Furthermore, irregular layouts, heterogeneous fonts, and handwritten text disrupt spatial reasoning and compromise global-local semantic misalignment, resulting in translation omissions and hallucinations.

To address these limitations, we propose **GLoTran**, a global-local dual visual perception framework for TIMT. GLoTran equips MLLMs with the ability to capture global context while preserving fine-grained focus on text-intensive regions, effectively mitigating semantic drift, hallucinations, and omissions in visually complex images. Specifically, we first perform text region detection to extract local visual slices, while downsampling the original high-resolution image to a low-resolution global view for efficient context encoding. Both the global image and local slices are jointly fed into the MLLM, guided by an structured prompt that instructs the global image as a contextual reference and subsequent slices as its regional counterparts for per-slice translation. The global image provides scene-level layout and semantic priors, ensuring contextual consistency, while each local slice preserves detailed textual information critical for accurate text localization and translation, as shown in Fig.1(d). Furthermore, we construct **GLoD**, a large-scale TIMT dataset tailed for our global-local dual perception paradigm. GLoD contains over 510K global-local image-text pairs from diverse real-world scenarios, including menu, documents, photos, posters, and road signs, facilitating effective fine-tuning of MLLMs for high-resolution text-rich TIMT. Our main contributions can be summarized as follows:

- We propose GLoTran, a global-local dual visual perception framework for MLLMs in TIMT, which jointly models scene-level context understanding and fine-grained text regions perception, enabling accurate translation of high-resolution text-rich images.
- We construct a large-scale TIMT dataset with over 510K global-local image-text pairs, systematically curated to facilitate global-local dual perception training of MLLMs.
- Experimental results on multiple benchmarks demonstrate that GLoTran outperforms state-of-the-art MLLMs, achieving higher translation accuracy for high-resolution text-rich TIMT.

2 Related Work

Text Image Machine Translation. Different from conventional text machine translation [5, 27, 55, 63], text image machine translation (TIMT) aims to translate the source-language texts embedded in images into a target-language [24, 36]. Existing TIMT works can be broadly categorized into two paradigms: (i) *Cascade methods* sequentially perform text recognition and translation using multiple models, but are prone to error accumulation, structural redundancy, and high latency [3, 21, 61], and (ii) *End-to-end models* jointly model visual and linguistic modalities to directly translate images to text, mitigating error propagation, and improving efficiency [24, 33, 34, 37, 44]. However, cascade methods are inherently vulnerable to compound recognition and translation errors, while end-to-end models remain task-specific and lack the scalability and adaptability required for real-world scenarios [33, 58, 66].

MLLM-Based TIMT. Multimodal Large Language Models (MLLMs), such as Qwen3-VL [54] and InternVL3 [64], have demonstrated the potential for end-




Fig. 2: Overview of the proposed GLoTran framework. (a) The high-resolution input image is processed by a text region detector to identify candidate textual regions, which are subsequently sorted, merged, and cropped into localized slices. (b) The global image and the local slices are fed into GLoTran with a structured prompt, enabling global contextual understanding and local textual focus for TIMT.

to-end image-to-text translation. Inspired by these advances, some works have explored adapting MLLMs more effectively for TIMT [17, 29, 30, 40, 66]. MT³ [16] improves the translation quality of MLLM through multi-task reinforcement learning, while M4Doc [30] focuses on MLLM-based document image translation. InImageTrans [66] introduces a multi-condition direct preference optimization strategy [41] to fine-tune MLLMs, reducing hallucinations and improving translation performance. Despite their comparative performance, existing MLLM-based TIMT methods remain largely limited to low-resolution images with relatively simple visual content. The application of MLLMs to end-to-end TIMT for real-world high-resolution text-rich images is still underexplored.

For high-resolution images, although some works [4, 39, 43, 51, 53] combine global views with local patches, they primarily leverage local patches as auxiliary visual evidence for global queries to support holistic VQA and reasoning tasks. Moreover, their random partitioning [39, 43, 51] or search-based strategies [4, 53] often lead to fragmented text, redundant regions, and missing text instances. However, TIMT requires comprehensive, ordered, and semantically consistent translation of all textual content in an image, rendering them not directly applicable. High-resolution images, such as menus and posters, typically contain cluttered backgrounds, diverse layouts, and non-textual visual distractions. These characteristics make it challenging for existing models to maintain fine-grained textual focus and global contextual consistency, resulting in hallucinations, semantic drift, and cross-region inconsistencies. In contrast, GLoTran adopts a text-region-centered extraction and perception paradigm. It translates each local region while actively querying global context for relevant semantic cues (e.g., layout and theme), and enforces translation coherence through pre-local region translation replay and a structured global-local prompting strategy.

3 Methodology

In this section, we introduce GLoTran, an efficient and scalable framework tailored for TIMT. GLoTran adopts a global-local dual visual perception strategy, enabling comprehensive contextual understanding while maintaining fine-grained attention to local textual regions.

3.1 Architecture Overview

TIMT translating source-language text embedded within an image into a target-language sequence. Formally, given a TIMT dataset $\mathcal{D} = \{(I, Y)\}$, where I denotes a source-language image and Y represents its corresponding target-language translation. TIMT seeks to learn a translator $\mathcal{F} : I \rightarrow Y$, such that the predicted sequence $\hat{Y} = \mathcal{F}(I)$ faithfully preserves the semantics of Y .

Unlike existing MLLM-based methods that perform translation over an entire high-resolution image in a single forward pass, GLoTran adopts a global-local dual visual perception strategy to balance holistic contextual understanding with fine-grained textual focus. As illustrated in Fig.2, GLoTran first detects translatable text regions within I and extracts the cropped local slices $\mathcal{I} = \{I_1, I_2, \dots, I_{N_s}\}$. The original image I is simultaneously resized into a low-resolution global view I_g that captures the overall layout and semantics of the scene, while each local slice I_i preserves detailed text content for fine-grained translation. The translation is then performed regressively across regions, guided by a structured instruction prompt $\mathcal{P}$ (see Sec. 3.3). At each translation step i , GLoTran predicts $\hat{Y}_i$ for I_i , conditioned on I_g , I_i and $\mathcal{P}_{<i}$. $\mathcal{P}_{<i}$ incorporates the translated outputs $\{\hat{Y}_{i-\eta}, \dots, \hat{Y}_{i-1}\}$ from the preceding η regions, where η is termed as pre-region translation replay window. This replay mechanism provides contextual cues such as sentence continuation, coreference grounding, and terminology consistency, thereby improving the accuracy and discourse coherence of local translations within complex, multi-text images.

3.2 Global-Local Dual Visual Perception

TIMT requires both scene-level semantic understanding and fine-grained textual recognition. However, pretrained MLLMs are constrained by limited input resolutions and the rigid square-patch design of their visual encoders. When MLLMs directly process high-resolution text-rich images, the superlinear computational overhead and dense non-text visual distractions overload visual attention, thereby weakening fine-grained text recognition in TIMT and significantly degrading translation accuracy. Consequently, we propose a global-local dual visual perception strategy that jointly models a low-resolution global view and multi-scale local slices, achieving a balance between global contextual understanding and local textual fidelity.

Specifically, a self-trained text region detector (see Appendix 10 for details) identifies candidate bounding boxes $\mathcal{B} = \{b_i\}_{i=1}^{N_b}$ from image I , where each box

corresponds to a block- or paragraph-level region that groups sentences from the same coherent text segment. Each detected region is cropped and normalized to form an initial local slice set $\hat{\mathcal{I}} = \{\hat{I}_i = \text{Crop}(I, b_i)\}_{i=1}^{N_b}$. Adjacent or semantically related regions are further merged according to geometric and typographic cues (e.g., overlap, line height, alignment), resulting in a compact grouped slice set $\tilde{\mathcal{I}} = \mathcal{M}(\hat{\mathcal{I}}) = \{\tilde{I}_i\}_{i=1}^{N_s}$, with $N_s \ll N_b$. To ensure translation coherence, the slices are deterministically arranged using a content ordering model (e.g., MinerU [48]) $\text{ord}(\cdot)$, yielding $\mathcal{I} = \text{ord}(\tilde{\mathcal{I}})$. Meanwhile, the original image I is downsampled into a low-resolution global view I_g , which captures holistic scene layout and contextual priors, whereas each local slice I_i retains fine-grained textual details.

Both I_g and I_i are independently encoded by a shared visual encoder (e.g., ViT [12]) to extract visual features $\tilde{V}_g$ and $\tilde{V}_i$:

$$\tilde{V}_g = \text{Enc}(I_g) \in \mathbb{R}^{n_g \times d_v}, \tilde{V}_i = \text{Enc}(I_i) \in \mathbb{R}^{n_s \times d_v}, \quad (1)$$

where d_v denotes the visual feature dimension, and n_g and n_s correspond to the number of patch tokens for the global and local views, respectively. Subsequently, a vision projector aligns visual features to the textual feature space, producing $V_g \in \mathbb{R}^{n_g \times d_t}$ and $V_i \in \mathbb{R}^{n_s \times d_t}$, where d_t denotes the text embedding dimension. Overall, the decoding of the LLM for global-local dual visual input is defined as:

$$\hat{Y}_i = \text{LLM}([E_g; V_g; E_l; V_i; \mathcal{P}_{<i}] \in \mathbb{R}^{L \times d_t}), \quad (2)$$

where E_g and E_l are textual embeddings of source-aware identifier tokens $\langle \text{image}_g \rangle$ and $\langle \text{image}_l \rangle$, which help the LLM distinguish features from global and local views.

To effectively integrate global context into local perception while preserving local text details, we introduce hierarchical cross-attention between V_g and V_i at early Transformer layers [0, 8, 16, 24]. This allows each local token to selectively attend to semantically relevant global tokens, thereby enhancing contextual grounding and maintaining coherence across the visual hierarchy [52, 57]. Formally, for a query q_i from a local token and a key k_j from a global token, the attention weight is computed as:

$$\text{Attn}(q_i, k_j) = \frac{q_i k_j^\top}{\sqrt{d_t}} + b_{ij}, \quad (3)$$

where b_{ij} is a learnable scalar bias conditioned on source type (global vs. local) and token spatial proximity. This adaptive bias encourages semantically grounded and spatially aware interactions, enabling model to prioritize relevant global cues for local disambiguation (e.g., resolving ambiguous text via scene context) while preserving intra-line coherence and ignoring irrelevant regions.

The training objective of GLoTran, parameterized by θ , is defined as:

$$\begin{aligned} \theta^* &= \arg \min_{\theta} \mathcal{L}(\mathcal{F}(I, \mathcal{P}); \theta) \\ &= \arg \min_{\theta} \frac{1}{|\mathcal{D}|} \sum_{(I, Y) \in \mathcal{D}} \left[- \sum_{i=1}^{N_s} Y_i \log p(\hat{Y}_i \mid I_g, I_i, \mathcal{P}_{<i}; \theta) \right], \end{aligned} \quad (4)$$




Fig. 3: An illustration of the structured prompt design in GLoTran.

where $\mathcal{L}$ denotes the loss of cross-entropy on target-language tokens. During training, teacher forcing [6, 46] is applied to replace previously predicted translations $\hat{Y}_{<i}$ with ground-truth $Y_{<i}$ within the replay window η , ensuring stable optimization and effective cross-region contextual learning.

3.3 Prompt Construction of GLoTran

Within GLoTran, Translation is performed in a regressive regional manner, where each local translation step is conditioned on the global image, the current local slice, and previously regions translations through a replay mechanism, thereby ensuring translation consistency and coherence across regions. As illustrated in Fig.3, the structured prompt $\mathcal{P}$ is composed of four interdependent components:

- **Global Understanding Instruction** encodes holistic scene semantics and spatial layout from the global image, establishing a contextual foundation for subsequent region-level translations.
- **Local Focus Instruction** emphasizes the active text region in current local slice, steering MLLMs attention toward fine-grained textual understanding while preserving alignment with the global contextual and visual layout.
- **Global-Local Consistency Rule** enforces semantic coherence between regional translations and the global view. Before generating outputs, MLLM validates that local translations are consistent with global scene semantics, mitigating contextual drift and contextual deviation. Furthermore, translated outputs from preceding regions $\{\hat{Y}_{i-\eta}, \dots, \hat{Y}_{i-1}\}$ are incorporated into the prompt via a translation replay window of size η , providing discourse cues for terminology uniformity and narrative continuity across slices.
- **Translation Instruction** specifies the explicit TIMT objective, e.g., *Translate the English text in this local slice into Chinese, using the global image for contextual grounding. Maintain linguistic consistency and coherence with previous slices translation:* { 斯克里布纳五月刊 } is provided when translating

the second slice I_2 , where the previous translation 斯克里布纳五月刊 originates from the slice I_1 . This explicit directive reinforces task fidelity and maintains linguistic fluency across regional decoding steps.

4 GLoD Dataset

4.1 Dataset Statistics

Tab.1 provides a comparative overview of our GLoD dataset against existing TIMT datasets, including OCRMT30K [25], DoTA [28], IIMT [24], and MIT-10M [26]. While most prior datasets offer full-image-level pairs, where each image aligns with a single holistic translation, such coarse-grained structures limit training for our global-local dual visual perception frameworks. To address this, we construct GLoD, a large-scale multimodal dataset tailored for global-local dual-input TIMT in text-rich scenarios, comprising 510K images and 517,354 image-text pairs, spanning over 40 real-world scenes in 5 languages. Curated via a rigorous pipeline, it features high-resolution, text-rich content, irregular layouts, complex backgrounds, and diverse fonts, as shown in Fig.4.

Table 1: The statistics of GLoD and other TIMT dataset.

Dataset	Text-Rich	Global-Local	Language	Images	Image-Text
ECOIT [65]	✗	✗	2	479,490	479,490
OCRMT30K [25]	✗	✗	2	30,000	30,000
DoTA [28]	✓	✗	2	126,345	126,000
IIMT [24]	✗	✗	1	89,033	89,033
MIT-10M [26]	✗	✗	14	840,855	10,931,115
GLoD	✓	✓	5	517,354	517,354

4.2 Curation Pipeline of GLoD

The GLoD curation pipeline encompasses five key stages: scene conceptualization, data collection and pre-filtering, text region detection and grouping, global-local translation, and quality control.

Scene Conceptualization. We define over 40 practical translation scenarios, including documents, posters, menus, road signs, and receipts, ensuring broad coverage of visual and content diversity.

Data Collection and Pre-Filtering. Images are sourced via large-scale web crawling and integration with existing text-image datasets. Strict filters exclude privacy-sensitive, confidential, or copyright-protected content, as well as low-resolution, blurred, distorted, or poor OCR confidence (using PaddleOCR [8]), uniform visual styles, and low text richness images.

Text Region Detection and Grouping. Filtered images undergo automated text region detection using complementary detectors, PaddleOCR and Qwen3-VL-Plus [54], to capture diverse text layouts through cross-region verification and complementary fusion. Detected regions are hierarchically grouped into semantically coherent blocks (phrases, sentences, or paragraphs) based on




Fig. 4: Overview of GLoD curation pipeline. The construction pipeline systematically generates high-quality global and local translation pairs through five core stages: conceptualization, data collection and pre-filtering, text region detection and grouping, global-local translation, and quality control.

reading order, spatial adjacency, and visual heuristics, while ambiguous or underspecified regions are excluded.

Global-Local Translation. Each global-local pair is processed through dual OCR recognition and bi-directional translation fusion. PaddleOCR captures fine-grained local content, while Qwen3-VL-Plus provides global context for cross-line disambiguation, enabling robust recognition under diverse fonts, print qualities, and image sizes. The recognized text is then translated using GPT-4o [23] and DeepSeek-R1 [20] under a multi-round mutual back-translation and fusion process, ensuring linguistic completeness and contextual alignment. Human experts conduct final verification and correction, focusing on semantics, grammar, and layout correspondence. This pipeline produces approximately 2M global-local image-text pairs with high linguistic and structural fidelity.

Quality Control. To guaranty reliability, we employ a two-stage validation strategy. Low-confidence OCR segments, ultra-short translations, and non-linguistic symbols are first removed. Subsequently, semantic faithfulness and consistency between global and local translations are quantified using multilingual embedding similarity (e.g., LaBSE [15] or SimCSE [18]) and round-trip translation consistency. Local regions showing semantic drift or contextual mismatch with their global references are excluded.

5 Experiments

5.1 Training and Test datasets

We conduct full-parameter fine-tuning of existing MLLMs, e.g., InternVL2.5 4B [64] and Qwen3-VL 8B [54], on GLOD using global-local dual visual perception strategy. To ensure a fair comparison, we evaluate the English-to-Chinese translation performance on the MCiTon [66] dataset, which contains 1,450 image-text pairs from documents, nature scenes and posters. Additionally, we evaluate multilingual translation capabilities on the MTIT6 dataset [40], which includes five

Table 2: Performance comparison of translating English into Chinese with open-source and close-source MLLMs on the MCiTon [66] test set. We report **BLEU**/**COMET** for translation quality of each real-world scenarios. $\uparrow$ indicates higher is better. The **bold** represents the best results, and the underline represents the second best results.

Method	Size	Document			Scene			Poster		Avg↑
		Paper ↑	News ↑	Novel ↑	Title ↑	Sign ↑	Introduction ↑	Cover ↑	Leaflet ↑	
Close-Source MLLMs										
GPT-4o [23]	~ 200B	62.1 / 86.3	51.6 / 85.6	44.2 / 83.4	51.3 / 83.4	46.1 / 83.3	40.1 / 75.3	34.9 / 69.1	43.9 / 71.7	46.8 / 79.9
Qwen-VL-Max [1]	~ 72B	60.5 / 85.7	58.3 / 85.7	44.8 / 83.5	50.1 / 81.6	43.2 / 80.5	37.8 / 74.6	34.2 / 67.3	43.7 / 69.9	46.5 / 78.6
Open-Source MLLMs										
Qwen2.5-VL-Instruct [2]	7B	60.7 / 85.8	53.4 / 85.5	27.4 / 78.1	47.9 / 80.1	44.0 / 82.2	36.2 / 72.8	30.3 / 64.4	43.6 / 72.4	42.9 / 77.7
Qwen2-VL [49]	7B	60.0 / 85.6	46.4 / 81.7	31.8 / 73.9	27.9 / 67.8	25.9 / 66.4	24.0 / 70.9	19.4 / 70.3	30.4 / 73.3	36.9 / 73.1
DeepSeek-VL-Chat [31]	7B	44.7 / 80.8	8.7 / 57.4	4.9 / 50.6	5.4 / 47.6	13.5 / 59.9	10.8 / 52.1	9.7 / 45.8	13.3 / 52.9	13.9 / 55.9
InternVL2 [7]	8B	44.0 / 83.0	35.8 / 78.9	23.7 / 71.3	27.2 / 65.3	28.2 / 61.1	25.4 / 68.7	19.4 / 62.0	21.0 / 62.1	28.1 / 68.5
ImageTrans [66]	8B	59.0 / 85.0	46.5 / 80.9	33.5 / 76.5	37.1 / 70.6	28.9 / 66.9	32.0 / 71.7	19.0 / 62.3	32.1 / 70.8	35.9 / 72.7
InternVL3 [64]	8B	59.8 / 86.0	48.9 / 83.9	31.2 / 79.8	50.5 / 82.1	41.8 / 81.6	32.8 / 72.6	29.5 / 66.1	35.8 / 70.9	41.4 / 77.9
Qianfan-VL [9]	8B	55.4 / 83.3	52.9 / 81.2	32.0 / 80.7	47.5 / 78.4	43.0 / 80.9	34.4 / 71.7	25.1 / 61.7	34.5 / 69.3	40.6 / 75.9
SAIL-VL2 [10]	8B	61.3 / 85.4	53.6 / 84.9	43.4 / 81.6	48.8 / 78.8	43.5 / 82.1	36.9 / 73.5	32.9 / 66.4	40.1 / 71.5	45.1 / 78.0
Qwen3-VL-Instruct [54]	8B	60.9 / 86.1	54.3 / 86.2	43.1 / 82.1	52.3 / 81.8	44.2 / 83.0	33.6 / 75.8	36.2 / 70.4	48.2 / 75.1	46.7 / 80.2
InternLM-XComposer2 [11]	11B	37.3 / 78.4	26.7 / 60.1	8.50 / 45.2	24.1 / 60.2	22.4 / 60.0	20.2 / 61.1	16.6 / 57.3	22.0 / 60.0	22.7 / 60.6
InternVL3 [64]	14B	63.2 / 86.4	53.8 / 85.4	36.7 / 81.5	49.1 / 81.5	42.4 / 81.8	34.7 / 73.8	30.1 / 66.2	35.3 / 70.9	41.7 / 78.2
CogVLM [50]	17B	59.2 / 84.7	44.8 / 80.6	30.7 / 73.7	36.4 / 71.1	27.7 / 66.7	30.5 / 70.6	25.7 / 67.8	30.1 / 71.8	35.5 / 72.7
Qwen2.5-VL-Instruct [2]	32B	63.2 / 86.3	54.4 / 86.2	28.5 / 79.2	39.1 / 77.3	46.2 / 82.6	39.9 / 75.7	33.9 / 68.1	45.5 / 72.8	42.2 / 78.6
InternVL2 [7]	40B	48.8 / 83.9	40.1 / 79.8	26.6 / 72.8	27.1 / 67.1	28.4 / 62.5	28.7 / 69.9	19.7 / 64.1	23.4 / 66.1	30.3 / 70.3
InternVL3* [64]	8B	58.3 / 86.3	50.5 / 85.5	32.1 / 81.4	52.4 / 83.4	38.1 / 82.3	34.6 / 73.6	31.0 / 68.5	33.8 / 69.3	41.4 / 78.8
Qwen3-VL-Instruct* [54]	8B	60.8 / 86.2	56.0 / 86.4	43.8 / 83.9	46.1 / 80.5	44.7 / 83.8	35.5 / 77.4	36.2 / 72.6	48.5 / 75.9	46.5 / 80.8
GLoTran-LoRA (Qwen3-VL 8B)	8B	64.8 / 85.9	56.0 / 86.7	45.5 / 84.8	53.1 / 81.2	45.0 / 82.7	40.0 / 74.5	36.1 / 69.9	49.5 / 75.8	42.2 / 78.6
GLoTran (Qwen3-VL 8B)	8B	66.3 / 87.9	57.5 / 88.6	46.8 / 86.6	54.3 / 83.2	46.3 / 84.8	41.3 / 76.4	37.7 / 72.0	51.1 / 77.9	50.2 / 82.2

translation tasks: Japanese-to-Chinese (jp $\rightarrow$ zh), Korean-to-Chinese (ko $\rightarrow$ zh), Chinese-to-English (zh $\rightarrow$ en), Chinese-to-Japanese (zh $\rightarrow$ jp) and Chinese-to-Korean (zh $\rightarrow$ ko), with approximately 200 high-resolution images per language pair.

5.2 Baselines and Implementation Details

We compare GLoTran with a wide range of open-source MLLMs, including Qwen2.5-VL-Instruct [2], Qwen3-VL [54], and InternVL3. We also include comparisons with state-of-the-art commercial MLLMs, namely GPT-4o [23] and Qwen-VL-Max [1]. For evaluation of translation quality, we adopt the metrics BLUE [38] and COMET [42], following the same evaluation procedures as ImageTrans [66]. All reported results are obtained either from officially released benchmarks or via re-inference using publicly available checkpoints to ensure fairness and consistency. GLoD contains 512,180 training pairs and 5,174 validation pairs under a 99:1 split, with no image overlap with MCiTon [66] or MTIT6 [40]. All original input images are resized to a fixed resolution of 224×224 . During full-parameter fine-tuning, the batch size is set to 128 with gradient accumulation over 8 steps. The AdamW optimizer is used with an initial learning rate of 1.0×10^{-5} , and the model is trained for 10,000 iterations in total. All experiments are conducted on 8 NVIDIA A100 (80GB) GPUs.

5.3 Performance Comparison

Multi-Scenes Translation Evaluation. Tab.2 reports the performance of various MLLMs on the MCiTon [66] benchmark. As shown in Tab.2, GLoTran, built




Fig. 5: (a)-(b) performance comparison of multilingual translating task with open-source MLLMs on MTIT6 [40]. (c)-(d) performance analysis with respect to global image resolution and the pre-region translation replay window size η .

on Qwen3-VL 8B, achieves improvements across all eight scenarios, outperforming most open-source and closed-source MLLMs of comparable or larger scales. GLoTran surpasses Qwen3-VL 8B, achieving gains of 7.49% and 2.49% in BLEU and COMET, highlighting the effectiveness of GLoTran in improving sensitivity to fine-grained textual details while preserving global contextual comprehension. In dense and heterogeneous layouts (e.g., documents and posters), GLoTran improves BLEU by 4.6% and 5.4% average gains over all open-source MLLMs. In scenes with small and scattered texts (e.g., introductions and leaflets), GLoTran achieves larger relative gains, outperforming Qwen3-VL 8B by approximately 5.3 BLEU and 1.7 COMET points, while achieving comparable improvements with Qwen2.5-VL-Instruct requires scaling its parameters from 7B to 32B. We further perform supervised fine-tuning (SFT) on our GLoD dataset for InternVL3 8B and Qwen3-VL 8B, denoted as InternVL3 $\star$ and Qwen3-VL $\star$ in Tab.2. The results indicate that directly fine-tuning these models on TIMT with GLoD, without tailored adaptation for high-resolution text-rich image translation, brings only limited gains and may even degrade performance in complex scenarios. These consistent improvements confirm that GLoTran effectively captures subtle textual variations in visually complex regions, alleviating fine-grained translation errors that commonly challenge MLLMs.

Furthermore, the results in Tab.2 indicate that the performance of MLLMs in TIMT tasks does not exhibit a strong positive correlation with the parameter scale. For models built on the same backbone, increasing the parameters beyond 7B yields only marginal or negligible gains. Specifically, InternVL2 8B and InternVL2 40B differ by merely 2.2 BLEU and 1.8 COMET points (28.1/68.5



Fig. 6: Qualitative comparison. The results in □ indicates omissions, those in □ indicate mistranslations, and those in □ indicate hallucinations caused by contextual misunderstanding.

vs. 30.3/70.3). In contrast, Qwen2.5-VL-Instruct 32B even underperforms its 7B variant in BLEU (42.2 vs. 42.9). These observations suggest a potential scaling saturation effect, which implies that the larger model capacity alone may not adequately address the intrinsic challenges of TIMT.

Multilingual Translation Evaluation. To further evaluate cross-lingual performance, we conduct multilingual translation experiments on the MTIT6 [40] dataset, covering six language pairs. As shown in Fig.5(a)-(b), GLoTran consistently achieves the best performance in all tasks, surpassing Qwen3-VL 8B and InternVL3 8B. Although InternVL3 and Qwen3-VL demonstrate strong multilingual understanding and text-image translation, they are constrained by two critical limitations: (i) insufficient suppression of non-textual visual distractions, which can produce hallucinatory or contextually inconsistent translations, and (ii) inadequate attention to small or fine-grained textual regions, leading to omissions or mistranslations of marginal or low-contrast text. These limitations contribute to the performance plateau observed in existing MLLMs, even as their parameter counts increase, underscoring the need for strategies that jointly capture global contextual coherence and fine-grained local perception.

Qualitative Comparison. Fig.6 provides a visual comparison of the translation outputs of Qwen3-VL 8B, InternVL3, and GLoTran. Translated texts are overlaid on the original images using a rule-based mapping strategy across challenging scenarios, including complex visual backgrounds, dense document,

Table 3: Ablation study on the effect of GLoTran across different backbone models. We report **BLEU** and **COMET** scores for each translation task. The Gain row denotes the relative improvement (%) of GLoTran over its baseline, calculated as $(\text{GLoTran} - \text{Baseline}) / \text{Baseline} \times 100\%$.

Method	zh → en	zh → ko	ko → zh	jp → zh	Novel	Title	Introduction	Leaflet
InternVL2.5 4B	4.02 / 70.78	1.19 / 60.44	2.57 / 48.20	28.37 / 76.66	18.48 / 75.67	31.44 / 76.10	26.16 / 71.09	21.22 / 63.92
GLoTran (InternVL2.5 4B)	6.68 / 71.46	2.26 / 60.86	5.97 / 58.89	29.93 / 78.33	25.80 / 77.51	38.20 / 79.82	29.47 / 71.52	23.27 / 65.50
Gain Δ	+66.17 / +0.96	+89.92 / +0.69	+132.29 / +22.18	+5.50 / +2.18	+39.61 / +2.43	+21.56 / +4.88	+12.65 / +0.61	+9.66 / +2.47
Qwen2.5-VL 7B	9.54 / 74.45	4.13 / 66.76	20.23 / 70.27	35.83 / 79.09	27.43 / 78.09	47.91 / 80.09	36.18 / 72.78	43.66 / 72.40
GLoTran (Qwen2.5-VL 7B)	10.99 / 76.78	5.45 / 68.49	22.39 / 70.99	43.81 / 81.84	44.59 / 83.60	53.15 / 82.92	38.93 / 74.12	45.64 / 73.74
Gain Δ	+15.20 / +3.13	+31.96 / +2.59	+10.68 / +1.02	+22.27 / +3.48	+62.59 / +7.06	+10.93 / +3.53	+7.60 / +1.84	+4.54 / +1.85
Qwen3-VL 8B	11.27 / 76.57	8.24 / 72.56	17.29 / 67.09	37.71 / 77.73	43.56 / 83.11	52.30 / 81.80	33.63 / 75.84	48.21 / 75.08
GLoTran (Qwen3-VL 8B)	11.87 / 79.27	9.62 / 75.46	20.69 / 69.73	43.54 / 81.92	46.84 / 86.61	54.33 / 83.21	41.32 / 76.40	51.18 / 77.91
Gain Δ	+5.33 / +3.53	+16.75 / +4.00	+19.66 / +3.94	+15.46 / +5.39	+7.53 / +4.21	+3.88 / +1.72	+22.86 / +0.74	+6.16 / +3.77

curved text layout, and handwritten menu. Existing MLLMs frequently exhibit omissions, mistranslations, and hallucinations in these complex settings. For example, in menus and posters with small, scattered, or cluttered text, Qwen3-VL and InternVL3 fail to capture subtle textual cues, misrepresenting or omitting critical elements, whereas GLoTran consistently preserves textual fidelity and layout coherence.

5.4 Ablation Studies

Ablation in Backbone. We evaluate the effectiveness of GLoTran with different backbone models, including InternVL2.5 4B, Qwen2.5-VL 7B and Qwen3-VL 8B, as summarized in Tab.3. Across all backbones, integrating GLoTran consistently yields substantial improvements in both BLEU and COMET scores. For InternVL2.5 4B, GLoTran achieves notable gains, particularly on low-resource translation tasks such as ko→zh and jp→zh, with BLEU improvements of 132.3% and 5.50%, respectively. Across other backbones, relative gains range from 4.54% to 141.0% compared to Qwen2.5-VL 7B, and up to 22.86% BLEU in the Introduction and 6.16% in the leaflet for Qwen3-VL 8B. These results reveal two key insights: (i) the global-local dual-perception strategy consistently improves translation fidelity across diverse backbones, and (ii) its benefits are particularly pronounced for small-scale or visually intricate text regions, highlighting that coordinated global-local perception yields greater gains than mere parameter scaling.

Components analysis. We employ the same text detector as GLoTran to extract text regions and feed these regions into existing MLLMs to evaluate their translation performance under accurate text localization. As shown in Tab.4, even when provided with localized text regions, the evaluated MLLMs (InternVL3 $\diamond$, Qwen2.5-VL $\diamond$, and Qwen3-VL $\diamond$) still perform unsatisfactorily, with accuracy further declining in complex scenes, e.g., News and Sign. To further isolate the effect of the text detector, we replace it in GLoTran with a parameter-free shape-adaptive cropping (PSC) [22, 56] strategy that partitions each high-resolution image into fixed-size local sub-images. As reported in Tab.4, GLoTran (PSC) exhibits a significant performance decline due to the non-grounded random cropping strategy, a common practice adopted by existing

Table 4: Components analysis results (BLEU/COMET) of GLoTran on the MCiTon dataset.

Method	Size	Document			Scene			Poster	
		Paper $\uparrow$	News $\uparrow$	Novel $\uparrow$	Title $\uparrow$	Sign $\uparrow$	Introduction $\uparrow$	Cover $\uparrow$	Leaflet $\uparrow$
InternVL3	8B	59.8/86.0	48.9/83.9	31.2/79.8	50.5/82.1	41.8/81.6	32.8/72.6	29.5/66.1	35.8/70.9
InternVL3 $\diamond$	8B	60.5/86.0	45.0/83.7	36.8/81.2	48.1/77.9	36.8/75.6	37.4/73.4	27.8/63.7	46.9/72.9
Qwen2.5-VL	7B	60.7/85.8	53.4/85.5	27.4/78.1	47.9/80.1	44.0/82.2	36.2/72.8	30.3/64.4	43.6/72.4
Qwen2.5-VL $\diamond$	7B	63.5/86.5	46.3/84.0	37.6/81.9	42.6/78.8	37.5/75.2	40.3/73.4	32.9/65.8	44.6/72.9
Qwen3-VL	8B	60.9/86.1	54.3/86.2	43.6/83.1	52.3/81.8	44.2/83.0	33.6/75.8	36.2/70.4	48.2/75.1
Qwen3-VL $\diamond$	8B	60.6/85.8	53.1/84.4	42.5/82.5	53.5/82.5	37.4/76.4	38.2/75.1	31.8/67.1	45.4/73.9
GLoTran (PSC)	Qwen3-VL 8B	57.1/83.4	50.8/78.1	36.8/78.4	46.2/80.3	36.2/75.3	33.2/73.3	30.1/61.7	43.8/72.7
GLoTran (w/o Global Image)	Qwen3-VL 8B	60.7/85.9	44.8/83.7	44.1/82.9	41.0/75.6	37.8/77.4	33.1/74.6	34.1/67.6	44.9/73.1
GLoTran (w/o Replay Mechanism)	Qwen3-VL 8B	64.8/85.9	56.0/86.7	45.5/84.8	53.1/81.2	45.0/82.7	40.0/74.5	36.1/69.9	49.5/75.8
GLoTran	Qwen3-VL 8B	66.3/87.9	57.5/88.6	46.8/86.6	54.3/83.2	46.3/84.8	41.3/76.4	37.7/72.0	51.1/77.9

MLLMs for high-resolution inputs, which introduces non-text regions, fragments text, and overlooks small text instances. Moreover, we conduct an ablation by removing the global image and retaining only local slices. As shown in Tab.4, GLoTran (w/o Global Image) exhibits a performance drop compared to GLoTran, onfirming the critical role of global contextual information.

Hyperparameter Sensitivity Analysis. Fig.5(c) shows the impact of varying global image resolutions (224, 448, 896, and 1792) on performance. The results indicate that moderate resolutions (224 or 448) have the best accuracy. Beyond 448, performance on introduction and leaflet scenes noticeably declines, likely due to excessive visual clutter and dispersed attention, which hinders effective global-context grounding. Consequently, we adopt 224 as the default global resolution and complement it with local slices to retain fine-grained textual details. Fig.5(d) presents the effect of the pre-region translation replay window η for GLoTran (Qwen2.5-VL 7B). BLEU peaks at 43.54 when $\eta = 4$, while smaller or larger window sizes lead to performance degradation. COMET scores fluctuate less but also favor moderate settings ($\eta \in [3, 5]$) to achieve coherent translations. An overly small η fails to provide sufficient target-language context, resulting in local ambiguity and inconsistency. Conversely, an excessively large η introduces accumulated noise and translation drift, degrading contextual fidelity and increasing computational overhead. In general, these sensitivity analyzes demonstrate that a balanced configuration of global visual resolution and contextual replay window maximizes the translation benefits of global-local cooperative modeling.

Complexity Analysis. To evaluate the trade-off between efficiency and accuracy, we conduct a complexity analysis of representative MLLMs on two challenging benchmarks, Novel (all images < 2K) and zh $\rightarrow$ jp (130/200 images $\geq$ 2K and 53/200 images $\geq$ 4K), which contain text-rich and visually complex high-resolution images. Tab.5 reports the average number of visual tokens per image (Token^V), First Token Latency (FTL_s), frame processing speed (FPS), and BLEU score. Under a unified resolution of 448 $\times$ 448, GLoTran $\heartsuit$ uses more visual tokens than Qwen3-VL $\heartsuit$ and Qwen2.5-VL $\heartsuit$, but achieves substantially higher BLEU scores, surpassing Qwen3-VL $\heartsuit$ by 38.14 and 7.95 points on Novel and zh $\rightarrow$ jp, respectively. Achieving comparable accuracy with Qwen3-VL $\clubsuit$ re-

Table 5: Comparison of MLLMs in terms of efficiency and accuracy for TIMI. The superscripts ♣, ♥, and ♠ denote input resolutions of 224×224 , 448×448 , and the original size, respectively.

Method	Size	Token ^V ↓	Novel			zh→jp		
			FTL _s ↓	FPS↑	BLEU↑	FTL _s ↓	FPS↑	BLEU↑
Qwen2.5-VL [♥]	7B	6.1K	0.098	0.061	27.36	0.098	3.755	11.85
InternVL3 [♥]	8B	10K	0.033	0.065	31.74	0.149	2.016	15.02
Qwen3-VL [♥]	8B	6.1K	0.099	0.028	9.98	0.093	2.572	18.79
Qwen3-VL [♠]	8B	164K	1.081	0.041	43.56	3.297	0.275	22.06
GLoTran [♥]	8B	8.4K	3.002	0.035	48.12	0.253	2.285	26.74
GLoTran [♣]	8B	4.9K	2.975	0.037	46.84	0.241	2.303	25.53

quires processing full-resolution images, which dramatically increases visual tokens (164K vs. 6.1K), first token latency (3.297 vs. 0.093), and computational cost ($\sim 215\times$ more FLOPs) in Novel. In contrast, our global-local dual perception framework enables GLoTran[♣] to maintain competitive performance even at lower resolutions 224×224 , capturing fine-grained textual details and global contextual layouts while keeping computational overhead manageable, requiring only 4.9K Token^V and about 18GB of memory.

These results indicate that existing MLLMs rely heavily on resolution scaling to improve accuracy, which introduces substantial visual redundancy with limited gains, whereas our global-local dual perception framework captures fine-grained text and global context even at low resolutions, enabling high-quality translation with lower computational cost.

6 Conclusion

In this paper, we propose GLoTran, a global-local dual visual perception framework for MLLMs in TIMT. By integrating global contextual understanding with fine-grained local text focus, GLoTran enables MLLMs to capture holistic semantics and precise localized details, thereby supporting accurate translation of high-resolution text-rich images. To facilitate global-local dual-perception learning for existing MLLMs, we further construct GLoD, a large-scale dataset containing over 510K global-local image-text pairs, systematically curated through a unified data generation pipeline. Extensive experiments demonstrate that GLoTran achieves improvements over state-of-the-art MLLMs, delivering superior translation completeness and reliability.

Acknowledgements

This work is supported by the National Natural Science Foundation of China under Project 62576206.

References

1. Bai, J., Bai, S., Yang, S., Wang, S., Tan, S., Wang, P., Lin, J., Zhou, C., Zhou, J.: Qwen-vl: A frontier large vision-language model with versatile abilities. arXiv preprint arXiv:2308.12966 **1**(2), 3 (2023)
2. Bai, S., Chen, K., Liu, X., Wang, J., Ge, W., Song, S., Dang, K., Wang, P., Wang, S., Tang, J., et al.: Qwen2.5-vl technical report. arXiv preprint arXiv:2502.13923 (2025)
3. Blecher, L., Cucurull, G., Scialom, T., Stojnic, R.: Nougat: Neural optical understanding for academic documents. arXiv preprint arXiv:2308.13418 (2023)
4. Cai, L., Kang, J., Li, S., Ma, W., Xie, B., Qin, Z., Liang, J.: From local details to global context: Advancing vision-language models with attention-based selection. In: Forty-second International Conference on Machine Learning
5. Castilho, S., Knowles, R.: A survey of context in neural machine translation and its evaluation. *Nat. Lang. Process.* **31**(4), 986–1016 (2025)
6. Chen, B., Martí Monsó, D., Du, Y., Simchowitz, M., Tedrake, R., Sitzmann, V.: Diffusion forcing: Next-token prediction meets full-sequence diffusion. *NIPS* **37**, 24081–24125 (2024)
7. Chen, Z., Wu, J., Wang, W., Su, W., Chen, G., Xing, S., Zhong, M., Zhang, Q., Zhu, X., Lu, L., et al.: Internvl: Scaling up vision foundation models and aligning for generic visual-linguistic tasks. In: *CVPR*. pp. 24185–24198 (2024)
8. Cui, C., Sun, T., Lin, M., Gao, T., Zhang, Y., Liu, J., Wang, X., Zhang, Z., Zhou, C., Liu, H., et al.: Paddleocr 3.0 technical report. arXiv preprint arXiv:2507.05595 (2025)
9. Dong, D., Zheng, M., Xu, D., Zhuang, B., Zhang, W., Luo, C., Wang, H., Zhao, Z., Li, J., Li, Y., et al.: Qianfan-vl: Domain-enhanced universal vision-language models. arXiv preprint arXiv:2509.18189 (2025)
10. Dong, H., Kang, Z., Yin, W., Liang, X., Feng, C., Ran, J.: Scalable vision language model training via high quality data curation. arXiv preprint arXiv:2501.05952 (2025)
11. Dong, X., Zhang, P., Zang, Y., Cao, Y., Wang, B., Ouyang, L., Wei, X., Zhang, S., Duan, H., Cao, M., et al.: Internlm-xcomposer2: Mastering free-form text-image composition and comprehension in vision-language large model. arXiv preprint arXiv:2401.16420 (2024)
12. Dosovitskiy, A.: An image is worth 16x16 words: Transformers for image recognition at scale. arXiv preprint arXiv:2010.11929 (2020)
13. Fang, W., Wu, Q., Chen, J., Xue, Y.: guided mllm reasoning: Enhancing mllm with knowledge and visual notes for visual question answering. In: *CVPR*. pp. 19597–19607 (2025)
14. Fang, Z., Lyu, P., Wu, J., Zhang, C., Yu, J., Lu, G., Pei, W.: Recognition-synergistic scene text editing. In: *CVPR*. pp. 13104–13113 (2025)
15. Feng, F., Yang, Y., Cer, D., Arivazhagan, N., Wang, W.: Language-agnostic bert sentence embedding. arXiv preprint arXiv:2007.01852 (2020)
16. Feng, Z., Liang, Y., Cao, S., Su, J., Ren, J., Xu, Z., Hu, Y., Huang, W., Wu, J., Liu, Z.: Mt3: Scaling mllm-based text image machine translation via multi-task reinforcement learning. arXiv preprint arXiv:2505.19714 (2025)
17. Gao, P., He, Z., Wu, H., Wang, H.: Towards boosting many-to-many multilingual machine translation with large language models. arXiv preprint arXiv:2401.05861 (2024)

18. Gao, T., Yao, X., Chen, D.: Simcse: Simple contrastive learning of sentence embeddings. In: EMNLP. pp. 6894–6910 (2021)
19. Gao, Y., Zhao, J., Sun, S., Qiao, X., Song, T., Yang, H.: Multimodal machine translation with text-image in-depth questioning. In: finding of ACL. pp. 9274–9287 (2025)
20. Guo, D., Yang, D., Zhang, H., Song, J., Zhang, R., Xu, R., Zhu, Q., Ma, S., Wang, P., Bi, X., et al.: Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. arXiv preprint arXiv:2501.12948 (2025)
21. Hinami, R., Ishiwatari, S., Yasuda, K., Matsui, Y.: Towards fully automated manga translation. In: AAAI. vol. 35, pp. 12998–13008 (2021)
22. Hu, A., Xu, H., Zhang, L., Ye, J., Yan, M., Zhang, J., Jin, Q., Huang, F., Zhou, J.: mplug-docowl2: High-resolution compressing for ocr-free multi-page document understanding. In: Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers). pp. 5817–5834 (2025)
23. Hurst, A., Lerer, A., Goucher, A.P., Perelman, A., Ramesh, A., Clark, A., Ostrow, A., Welihinda, A., Hayes, A., Radford, A., et al.: GPT-4o system card. arXiv preprint arXiv:2410.21276 (2024)
24. Lan, Z., Niu, L., Meng, F., Zhou, J., Zhang, M., Su, J.: Translatotron-v (ison): An end-to-end model for in-image machine translation. In: finding of ACL. pp. 5472–5485 (2024)
25. Lan, Z., Yu, J., Li, X., Zhang, W., Luan, J., Wang, B., Huang, D., Su, J.: Exploring better text image translation with multimodal codebook. In: ACL. pp. 3479–3491 (2023)
26. Li, B., Zhu, S., Wen, L.: Mit-10m: A large scale parallel corpus of multilingual image translation. In: COLING. pp. 5154–5167 (2025)
27. Li, H., Jie, Z., Lu, W.: Non-autoregressive machine translation as constrained hmm. In: finding of ACL. pp. 12361–12372 (2024)
28. Liang, Y., Zhang, Y., Ma, C., Zhang, Z., Zhao, Y., Xiang, L., Zong, C., Zhou, Y.: Document image machine translation with dynamic multi-pre-trained models assembling. In: NAACL. pp. 7084–7095 (2024)
29. Liang, Y., Zhang, Y., Zhang, Z., Chen, Z., Zhao, Y., Xiang, L., Zong, C., Zhou, Y.: Improving mllm’s document image machine translation via synchronously self-reviewing its ocr proficiency. In: finding of ACL. pp. 23659–23678 (2025)
30. Liang, Y., Zhang, Y., Zhang, Z., Zhao, Y., Xiang, L., Zong, C., Zhou, Y.: Single-to-mix modality alignment with multimodal large language model for document image machine translation. In: ACL. pp. 12391–12408 (2025)
31. Lu, H., Liu, W., Zhang, B., Wang, B., Dong, K., Liu, B., Sun, J., Ren, T., Li, Z., Yang, H., et al.: Deepseek-vl: towards real-world vision-language understanding. arXiv preprint arXiv:2403.05525 (2024)
32. Luan, B., Feng, H., Chen, H., Wang, Y., Zhou, W., Li, H.: Textcot: Zoom in for enhanced multimodal text-rich image understanding. arXiv preprint arXiv:2404.09797 (2024)
33. Ma, C., Han, X., Wu, L., Zhang, Y., Zhao, Y., Zhou, Y., Zong, C.: Modal contrastive learning based end-to-end text image machine translation. *ACM Trans. Audio Speech Lang. Process.* **32**, 2153–2165 (2023)
34. Ma, C., Zhang, Y., Tu, M., Han, X., Wu, L., Zhao, Y., Zhou, Y.: Improving end-to-end text image translation from the auxiliary text translation task. In: ICPR. pp. 1664–1670. IEEE (2022)
35. Ma, C., Zhang, Y., Tu, M., Zhao, Y., Zhou, Y., Zong, C.: Multi-teacher knowledge distillation for end-to-end text image machine translation. In: ICDAR. pp. 484–501. Springer (2023)

36. Mansimov, E., Stern, M., Chen, M.X., Firat, O., Uszkoreit, J., Jain, P.: Towards end-to-end in-image neural machine translation. In: NLPBW. pp. 70–74 (2020)
37. Niu, L., Meng, F., Zhou, J.: Umtit: Unifying recognition, translation, and generation for multimodal text image translation. In: LREC-COLING. pp. 16953–16972 (2024)
38. Papineni, K., Roukos, S., Ward, T., Zhu, W.J.: Bleu: a method for automatic evaluation of machine translation. In: ACL. pp. 311–318 (2002)
39. Park, J., Choi, J.Y., Park, J., Han, B.: Hierarchical visual feature aggregation for ocr-free document understanding. In: NeurIPS. vol. 37, pp. 105972–105996 (2024)
40. Qian, Z., Zhang, P., Yang, B., Fan, K., Ma, Y., Wong, D., Sun, X., Ji, R.: Anytrans: Translate anytext in the image with large scale models. In: finding of EMNLP. pp. 2432–2444 (2024)
41. Rafailov, R., Sharma, A., Mitchell, E., Manning, C.D., Ermon, S., Finn, C.: Direct preference optimization: Your language model is secretly a reward model. In: NeurIPS (2023), <https://arxiv.org/abs/2305.18290>
42. Rei, R., Stewart, C., Farinha, A.C., Lavie, A.: Comet: A neural framework for mt evaluation. In: EMNLP. pp. 2685–2702 (2020)
43. Shen, H., Zhao, K., Zhao, T., Xu, R., Zhang, Z., Zhu, M., Yin, J.: Zoomeye: Enhancing multimodal llms with human-like zooming capabilities through tree-based image exploration. In: Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing. pp. 6613–6629 (2025)
44. Su, T., Liu, S., Zhou, S.: Rtnet: An end-to-end method for handwritten text image translation. In: ICDAR. pp. 99–113. Springer (2021)
45. Tian, Y., Li, X., Liu, Z., Guo, Y., Wang, B.: In-image neural machine translation with segmented pixel sequence-to-sequence model. In: finding of EMNLP. pp. 15046–15057 (2023)
46. Tian, Y., Han, Y., Chen, X., Wang, W., Chawla, N.V.: Beyond answers: Transferring reasoning capabilities to smaller llms using multi-teacher knowledge distillation. In: WSDM. pp. 251–260 (2025)
47. Vu, Y.L., Duong, D.T., Duong, T.B., Nguyen, A.K., Nguyen, T.H., Nguyen, L.T.P., Xing, J., Li, X., Wang, T., Bagci, U., et al.: Describe anything model for visual question answering on text-rich images. arXiv preprint arXiv:2507.12441 (2025)
48. Wang, B., Xu, C., Zhao, X., Ouyang, L., Wu, F., Zhao, Z., Xu, R., Liu, K., Qu, Y., Shang, F., et al.: Mineru: An open-source solution for precise document content extraction. arXiv preprint arXiv:2409.18839 (2024)
49. Wang, P., Bai, S., Tan, S., Wang, S., Fan, Z., Bai, J., Chen, K., Liu, X., Wang, J., Ge, W., et al.: Qwen2-vl: Enhancing vision-language model’s perception of the world at any resolution. arXiv preprint arXiv:2409.12191 (2024)
50. Wang, W., Lv, Q., Yu, W., Hong, W., Qi, J., Wang, Y., Ji, J., Yang, Z., Zhao, L., Song, X., et al.: Cogvlm: visual expert for pretrained language models. In: NeurIPS. pp. 121475–121499 (2024)
51. Wang, W., Ding, L., Zeng, M., Zhou, X., Shen, L., Luo, Y., Yu, W., Tao, D.: Divide, conquer and combine: A training-free framework for high-resolution image perception in multimodal large language models. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 39, pp. 7907–7915 (2025)
52. Wu, H., Guo, M., Ma, Y., Sun, Y., Wang, J., Matusik, W., Long, M.: Flashbias: Fast computation of attention with bias. arXiv preprint arXiv:2505.12044 (2025)
53. Wu, P., Xie, S.: V*: Guided visual search as a core mechanism in multimodal llms. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 13084–13094 (2024)

54. Yang, A., Li, A., Yang, B., Zhang, B., Hui, B., Zheng, B., Yu, B., Gao, C., Huang, C., Lv, C., et al.: Qwen3 technical report. arXiv preprint arXiv:2505.09388 (2025)
55. Yang, W., Wu, J., Wang, C., Zong, C., Zhang, J.: Language imbalance driven rewarding for multilingual self-improving. In: ICLR (2025)
56. Ye, J., Hu, A., Xu, H., Ye, Q., Yan, M., Xu, G., Li, C., Tian, J., Qian, Q., Zhang, J., et al.: Ureader: Universal ocr-free visually-situated language understanding with multimodal large language model. In: finding of EMNLP. pp. 2841–2858 (2023)
57. Ye, J., Xu, H., Liu, H., Hu, A., Yan, M., Qian, Q., Zhang, J., Huang, F., Zhou, J.: mplug-owl3: Towards long image-sequence understanding in multi-modal large language models. In: ICLR
58. Zhang, H., Chen, K., Bai, X., Xiang, Y., Zhang, M.: Paying more attention to source context: Mitigating unfaithful translations from large language model. In: finding of ACL (2024)
59. Zhang, Z., Zhang, Y., Liang, Y., Xiang, L., Zhao, Y., Zhou, Y., Zong, C.: Layoutdit: Layout-aware end-to-end document image translation with multi-step conductive decoder. In: finding of EMNLP. pp. 10043–10053 (2023)
60. Zhang, Z., Zhang, Y., Liang, Y., Xiang, L., Zhao, Y., Zhou, Y., Zong, C.: From chaotic ocr words to coherent document: A fine-to-coarse zoom-out network for complex-layout document image translation. In: COLING. pp. 10877–10890 (2025)
61. Zhang, Z., Zhang, Y., Xiang, L., Zhao, Y., Zhou, Y., Zong, C.: A novel dataset and benchmark analysis on document image translation. In: CCMT. pp. 103–115. Springer (2023)
62. Zhao, H.H., Zhou, P., Gao, D., Bai, Z., Shou, M.Z.: Lova3: Learning to visual question answering, asking and assessment. In: NeurIPS. vol. 37, pp. 115146–115175 (2024)
63. Zhao, Y., Zhang, J., Zhou, Y., Zong, C.: Knowledge graphs enhanced neural machine translation. In: IJCAI. pp. 4039–4045 (2021)
64. Zhu, J., Wang, W., Chen, Z., Liu, Z., Ye, S., Gu, L., Tian, H., Duan, Y., Su, W., Shao, J., et al.: Internvl3: Exploring advanced training and test-time recipes for open-source multimodal models. arXiv preprint arXiv:2504.10479 (2025)
65. Zhu, S., Li, S., Lei, Y., Xiong, D.: Peit: bridging the modality gap with pre-trained models for end-to-end image translation. In: ACL. pp. 13433–13447 (2023)
66. Zuo, F., Chen, K., Zhang, Y., Xue, Z., Zhang, M.: Inimagetrans: Multimodal llm-based text image machine translation. In: finding of ACL. pp. 20256–20277 (2025)