

MIT-10M: A Large Scale Parallel Corpus of Multilingual Image Translation

Bo Li^{1,3}, Shaolin Zhu²*, Lijie Wen¹*

¹School of Software, Tsinghua University, Beijing, China

²College of Intelligence and Computing, Tianjin University, Tianjin, China

³Baidu Inc., Beijing, China

libo15@baidu.com, zhushaolin@tju.edu.cn, wenlj@tsinghua.edu.cn

Abstract

Image Translation (IT) holds immense potential across diverse domains, enabling the translation of textual content within images into various languages. However, existing datasets often suffer from limitations in scale, diversity, and quality, hindering the development and evaluation of IT models. To address this issue, we introduce MIT-10M¹, a large-scale parallel corpus of multilingual image translation with over 10M image-text pairs derived from real-world data, which has undergone extensive data cleaning and multilingual translation validation. It contains 840K images in three sizes, 28 categories, tasks with three levels of difficulty and 14 languages image-text pairs, which is a considerable improvement on existing datasets. We conduct extensive experiments to evaluate and train models on MIT-10M. The experimental results clearly indicate that our dataset has higher adaptability when it comes to evaluating the performance of the models in tackling challenging and complex image translation tasks in the real world. Moreover, the performance of the model fine-tuned with MIT-10M has tripled compared to the baseline model, further confirming its superiority.

1 Introduction

Image Translation (IT), the task of translating embedded text within an image from a source language to a target language (Watanabe et al., 1998; Yang et al., 2002; Lan et al., 2023), holds significant promise for various applications. Its utility spans domains such as scene text translation, document image translation (Liang et al., 2024), and photo translation, enhancing accessibility and cross-lingual communication. This is in contrast to traditional neural machine translation (Chen et al.,

*The corresponding author.

¹Our datasets and code are publicly available at: <https://huggingface.co/datasets/liboaccn/MIT-10M>

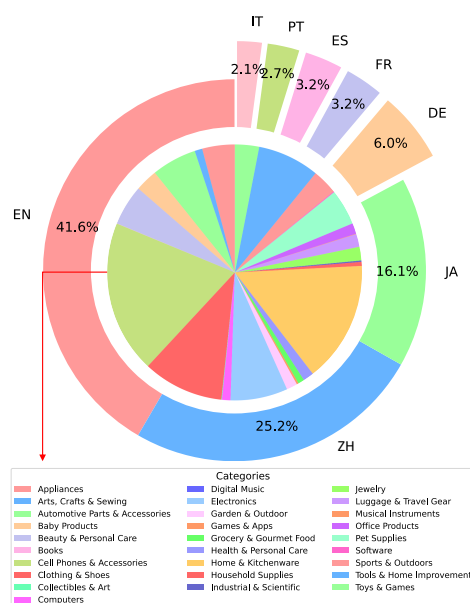


Figure 1: Categories and languages of MIT-10M. It includes 28 categories and 14 languages image-text pairs (8 languages images).

2022; Zhu et al., 2024), which is based solely on textual information.

Traditional IT approaches often employed cascade methods (Zhang et al., 2019; Zhao et al., 2020; Shekar et al., 2021; Hinami et al., 2021; Affi and Way, 2016). The emergence of end-to-end IT models (Zhu et al., 2023; Liang et al., 2024; Niu et al., 2024; Ma et al., 2023a) offered a more streamlined approach, utilizing a unified model to directly translate image text. Multimodal Large Language Models (MLLMs) (Bai et al., 2023; Chen et al., 2024b; Hong et al., 2024; Liu et al., 2024; Yao et al., 2024; Lu et al., 2024; Li et al., 2023) have further fueled progress.

However, IT model development and evaluation lies in the scarcity of high-quality, large-scale datasets (Ma et al., 2022, 2023b; Zhu et al., 2023; Ma et al., 2023a; Liang et al., 2024). Models trained on limited real-world data often struggle



Figure 2: Examples of MIT-10M dataset. Each image contains the original text and the corresponding language. Additionally, we annotate the image category, the token length of the text and the number of bounding boxes and used them for the difficulty level. In addition to the original text, 13 languages translations were annotated.

to generalize to complex scenarios, while those trained on synthetic datasets (Chen et al., 2021; Su et al., 2021; Niu et al., 2024) may not adequately capture the nuances of real-world image characteristics. Furthermore, existing datasets frequently lack fine-grained splits and diversity in language representation and task difficulty, hindering comprehensive model assessment.

To address these issues, we introduce MIT-10M, a large-scale multilingual image translation corpus. MIT-10M is the largest real-world image translation dataset to date. As shown in Figure 1, it includes 840K images in 8 languages and 10M image-text pairs across 14 languages, and split into a train set a test set. Figure 2 shows some examples from the MIT-10M dataset. Each image contains the original text and the corresponding language. Additionally, we annotate the image category, the token length of the text and the number of bounding boxes and used them for the difficulty level. In addition to the original text, 13 further translations were annotated. Detailed examples can be found in Appendix B. The MIT-10M dataset construction process consists of three main stages: data collection and pre-processing, OCR annotation and cleaning, and multilingual translation and validation. We compared MIT-10M with several existing popular image translation datasets in terms of data scale, task difficulty level, diversity, and image quality. The comparison results demonstrate that MIT-10M significantly outperforms other datasets in these aspects. We conduct extensive experiments to evaluate the multilingual translation capabilities of 7 end-to-end IT models using the MIT-10M test set. The results show that using the MIT-10M test

set is beneficial when evaluating challenging tasks of translating multi-line text into images with complex backgrounds. Furthermore, we fine-tune the Qwen2-VL (Bai et al., 2023) model using the MIT-10M training set, BLEU, chrF++, and METEOR scores have increased by 230%, 88%, and 130% respectively. The results demonstrate significant performance improvements in multilingual image translation tasks, further validating the advantages of MIT-10M dataset.

The main contributions are outlined below:

- We present MIT-10M, a dataset comprising 10M image-text pairs and 840K high-resolution real-world images, representing the largest publicly available real-world high-quality dataset specifically designed for multilingual image translation tasks.
- We propose a multi-dimensional evaluation framework for multilingual image-text translation datasets, considering aspects such as data scale, task difficulty level, diversity, and image quality. Our comparative analysis demonstrates the superiority of MIT-10M across these dimensions.
- We conduct comprehensive experiments to validate the effectiveness of MIT-10M for training and evaluating multilingual image translation models.

2 Related Work

IT focuses on translating text embedded within images from a source language into a target language. This section reviews existing approaches and datasets relevant to IT.

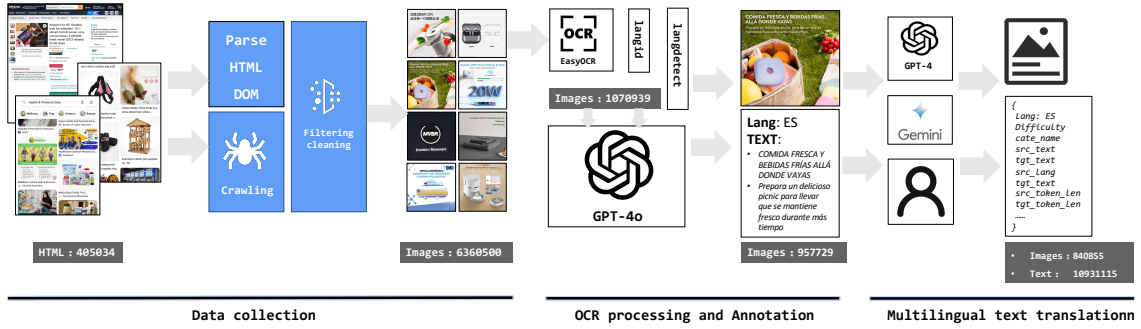


Figure 3: Overview of MIT-10M dataset construction pipeline.

2.1 IT Models

Two primary paradigms dominate the field of image translation:

Cascade Methods. These approaches decompose the task into sequential steps, typically employing Optical Character Recognition (OCR) followed by NMT (Zhao et al., 2020; Shekar et al., 2021; Hinami et al., 2021; Zhong et al., 2024; Chen et al., 2024a; Zhu et al., 2024). While conceptually straightforward, cascade methods suffer from error propagation between stages, increased latency due to separate model processing, and redundant parameterization (Ma et al., 2023b, 2022; Lan et al., 2023; Zhu et al., 2023).

End-to-End Methods. These models aim to directly translate image text using a unified architecture (Mansimov et al., 2020; Jain et al., 2021; Zhu et al., 2023; Liang et al., 2024; Niu et al., 2024; Ma et al., 2023a), typically comprising a visual encoder for extracting image features and a text decoder for generating the target translation. The advent of Multimodal Large Language Models (MLLMs), such as (Bai et al., 2023; Chen et al., 2024b; Hong et al., 2024; Liu et al., 2024; Yao et al., 2024; Lu et al., 2024; Li et al., 2023), has significantly advanced end-to-end IT, enabling more effective cross-modal fusion and improved translation accuracy. However, the success of end-to-end models heavily relies on the availability of large-scale, high-quality parallel image-translation data, which remains a critical challenge in the field (Ma et al., 2022, 2023b; Zhu et al., 2023; Ma et al., 2023a; Liang et al., 2024).

2.2 IT Datasets

Due to the scarcity of real-world image translation data, several studies have utilized synthetic datasets created by rendering text onto background

images. These include datasets focused on single-line text (Mansimov et al., 2020), multi-line text (Jain et al., 2021), Chinese-English translation (Ma et al., 2022), and English-German translation (Niu et al., 2024). While valuable for initial model training, synthetic datasets often oversimplify real-world complexities and may not adequately reflect the diversity and challenges encountered in practical applications. Manually curated real-world datasets are crucial for advancing IT research, but their development is resource-intensive, resulting in limited availability. Existing datasets include OCRMT30K (Lan et al., 2024), a Chinese-English dataset based on OCR annotations; DoTA (Liang et al., 2024), focused on document image translation into Markdown format; and ECOIT (Zhu et al., 2023), targeting the e-commerce domain. Despite these efforts, real-world datasets remain limited in scale, language coverage, and diversity of image characteristics, hindering comprehensive model evaluation and generalization assessment. Our work addresses this gap by introducing **MIT-10M**, a novel large-scale, multilingual, and real-world image translation dataset designed to facilitate the development and evaluation of robust and generalizable IT models.

3 MIT-10M Construction

In this section, the construction pipeline of MIT-10M is described in detail, highlighting the crucial filtering steps to ensure data quality. Figure 3 illustrates the MIT-10M construction pipeline.

3.1 Data collection

Web Crawling. High-quality websites with extensive language coverage were selected for data collection, including google.com, baidu.com, amazon.com, jd.com, and amazon.jp.co. These websites offer a rich source of high-resolution images

Dataset	Source	Languages	Images	Image-Text
E2E_TIT_With_MT (Ma et al., 2022)	synthetic	3	3,000,000	3,000,000
ECOIT (Zhu et al., 2023)	realistic	2	479,490	479,490
OCRMT30K (Lan et al., 2023)	realistic	2	30,000	30,000
DoTA (Liang et al., 2024)	realistic	2	126,345	126,000
IIMT (Lan et al., 2024)	synthetic	1	89,033	89,033
MIT-10M (Ours)	realistic	14	840,855	10,931,115

Table 1: Comparison of MIT-10M with other popular image translation datasets in terms of statistical data.

and support multiple languages. Crawling was conducted across eight languages: English, French, Chinese, Japanese, Portuguese, Italian, German, and Spanish. To ensure data uniqueness, SHA256 hashing was employed to identify and remove duplicate pages. Documents lacking images or containing an excessive number of images (over 50) were excluded. This process resulted in 405K unique HTML files, occupying 440 GB.

Image collection. We use the BeautifulSoup (Leonard, 2004) library to parse the HTML documents into a tree structure and extract the image tags. To ensure the quality of the data, we filter the images according to their resolution and only keep the images with a resolution of more than 800x800 pixels that are in the main content area of the HTML document to exclude logos and advertisements. Duplicate images are then removed based on their MD5 hash values.

Filtering and cleaning data. To ensure that the content is appropriate, we apply a NSFW detection tool (Laborde, 2024) to all images. If NSFW content is detected, we discard all images from the corresponding HTML document. The final set contains 6.3M images occupying 900 GB.

3.2 Data annotation

Initial text recognition. To quickly determine whether images contain text, we first apply easy-OCR² for text recognition and remove images where no text is recognized or the recognized text is meaningless, which accounts for about 50% of the images. This results in 3M images being retained. Next, we use the tools langid³ and langdetect⁴ to identify the language of the recognized text and perform cross-validation. Only images for which both tools consistently identify the same language

are retained. The result is 1 million images.

Precise text recognition. We then use GPT-4o (OpenAI, 2024) for detailed text recognition to extract precise text information from the images and discard those images where the text cannot be recognized.

Filtering sensitive content and text cleaning.

To reduce the risk of sharing personal data, we remove images whose OCR-recognized text contains sensitive information such as email addresses or IP addresses. We also exclude images where the recognized text contains NSFW characters to exclude potentially inappropriate content. To further improve the quality of the dataset, we filter out images with advertising or meaningless content (e.g. excessive numbers or punctuation) and remove images with excessively long text (over 450 tokens or 60 words). After the cleanup, 957K images remain, taking up 150 GB.

3.3 Multilingual text translation

Machine translation. To translate the text into other languages, we first use GPT-4 (Achiam et al., 2023) to translate the OCR-recognized text. We create a highly optimized translation prompt that translates the image text from 8 languages into 13 languages, with the additional languages being Korean, Thai, Arabic, Turkish, Hindi and Russian.

Validation of the translation. Although GPT-4 performs well in multilingual translation, we cross-validate with Google Translate (via the Google Gemini 1.5 Pro API) (Gemini, 2024). We use the tool spacy to convert the results of GPT-4 and Google Translate into word vectors and calculate their semantic similarity. For text pairs with a similarity score below 0.8, we filter out those with significant differences and keep translations with a semantic similarity above 0.8. The final parallel corpus contains 840K images.

²<https://github.com/JaidedAI/EasyOCR>

³<https://github.com/saffsd/langid.py>

⁴<https://github.com/fedeloopez77/langdetect>

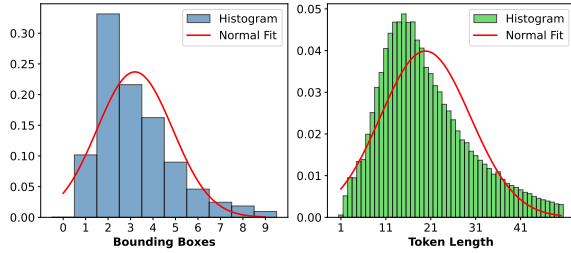


Figure 4: Distribution of the number of bounding boxes in the images and the token lengths in English text.

	@Train	@Test
Easy	3,154,034 (28.9%)	3,120 (30.0%)
Medium	3,156,621 (28.9%)	3,120 (30.0%)
Hard	4,610,060 (42.2%)	4,160 (40.0%)

Table 2: The number and proportion of corpora of different difficulty levels of MIT-10M.

Human evaluation. We divide the translated image-text pairs and their corresponding text translations into 10,000 batches. We randomly select 10 batches and perform a human evaluation, which results in a translation accuracy of 99.4%.

4 Analysis MIT-10M

This section presents a comprehensive analysis of the MIT-10M dataset. First, we compare MIT-10M with existing popular image translation datasets (see Table 1). Then, we analyze MIT-10M in detail in terms of data scale, difficulty, diversity, and image quality.

4.1 Data Scale Comparison

Table 1 shows a comparison of MIT-10M with other popular image translation datasets in terms of statistical data. MIT-10M clearly outperforms the other datasets in terms of the number of languages, images and image-text pairs.

MIT-10M includes 14 languages and is therefore better suited for cross-language multimodal image translation research, especially for scenarios requiring the processing of multiple languages. In addition, MIT-10M is derived from real-world scenarios and contains 840K images and 10M image-text pairs. This is a significantly larger volume than data sets such as DoTA and ECOIT, which also consist of real images.

4.2 Difficulty Levels

We visualized the distribution of the number of bounding boxes in the images and the correspond-

Cate ID	Category Name	Images
10000	Appliances	34711
11000	Digital Music	44
12000	Electronics	60884
13000	Garden & Outdoor	11559
14000	Games & Apps	1940
15000	Grocery & Gourmet Food	6654
16000	Health & Personal Care	11447
17000	Home & Kitchenware	129192
18000	Household Supplies	4298
19000	Industrial & Scientific	1486
20000	Arts, Crafts & Sewing	8077
21000	Luggage & Travel Gear	13643
22000	Musical Instruments	451
23000	Office Products	11462
24000	Pet Supplies	38158
25000	Software	977
26000	Sports & Outdoors	26742
27000	Tools & Home Improvement	65771
28000	Toys & Games	25776
29000	Computers	9039
30000	Automotive Parts & Accessories	47869
31000	Jewelry	14538
40000	Baby Products	23931
50000	Beauty & Personal Care	42964
60000	Books	211
70000	Cell Phones & Accessories	162428
80000	Clothing & Shoes	85998
90000	Collectibles & Art	605

Table 3: Distribution of image category.

ing lengths of the English tokens in the MIT-10M dataset, as shown in Figure 4. Most of the images have a bounding box count between 1 and 3, with the proportion of images with a bounding box count of 2 being the highest. In addition, the length of text tokens in the images is mainly in the range of 10 to 25 tokens, which is consistent with the distributional characteristics of text length in natural language.

Based on the above analysis, we categorize the MIT-10M dataset into three difficulty levels: easy, medium, and hard. Table 2 shows the number and proportion of corpora of different difficulty levels of the MIT-10M training and testing sets.

Easy (number of bounding boxes ≤ 2 and token length ≤ 16) contain fewer bounding boxes and shorter texts, resulting in a relatively easy translation task.

Hard (number of bounding boxes > 5 or length of tokens > 25) contain more bounding boxes or longer texts, which places higher demands on the model’s attention mechanism.

Medium (other cases) have a wider spread in terms of the number of bounding boxes and length of tokens, indicating more realistic and complex



Figure 5: Example of images with different resolutions.

image translation tasks. With this category, the generalization ability of multimodal translation models can be better tested.

4.3 Diversity

As shown in Figure 1, MIT-10M is divided into 28 categories based on image content, ranging from daily objects (e.g., Home & Kitchenware, Health & Personal Care) to specialized equipment (e.g., Industrial & Scientific, Jewelry) and digital goods (e.g., Games & Apps, Digital Music). The detailed categories and data can be found in Table 3. And categories such as "Cell Phones & Accessories" (162K images) and "Clothing & Shoes" (86K images) have a significant number of images. The products in these categories are frequently used in daily life, which means they are more common. This diversity provides a wealth of scenarios for image translation models and enables a comprehensive evaluation of their performance in different image types and text contexts. Consequently, training with such diverse data helps in developing models with superior generalization abilities. The English images dominate the dataset (about 49%), while Chinese is the second most common language and accounts for about 60% of the English data. This diversity of language and image data helps to improve the generalization ability of the model in different cultural and linguistic environments.

4.4 Image resolution

When translating images to text, the quality and resolution of the images have a direct impact on how the model extracts semantic features from the visual information, which in turn affects translation accuracy.

Unlike previous works, which often only include images with a single resolution (e.g. the ECOIT dataset with an image resolution of only 64x600), MIT-10M selects images with a resolution of more

than 1000x1000 pixels (width) when creating the dataset. It also provides images in three different sizes: the original size, a "large" version with a width of 768 pixels and a "small" version with a width of 500 pixels (Figure 5). These different image sizes allow the model to learn visual features under different conditions, which increases its robustness.

5 Experiments

In this section, we empirically demonstrate the efficacy of the MIT-10M dataset both as a pretraining dataset as well as an evaluation set for image-text translate task.

5.1 Experiment Details

Setup. The operating system which we use is CentOS Linux release 7.5, and the programming language is Python 3.9.12. Our experiments are conducted on NVIDIA TESLA A100-40G GPU, the CUDA version is 12.2, and the deep learning framework is torch with version 2.1.0, torchvision with version 0.16.0 and Transformers with 4.44.2. During the inference stage, we set the following hyperparameters: *temperature* to 0.2.

Models. In this study, we perform a comparative analysis of our dataset using both cascaded and end-to-end models for image translation. The cascaded model first applies EasyOCR to extract text from images and then translates the extracted text using the NLLB (NLLB Team et al., 2022) model. This choice of established components makes our baseline representative of typical cascaded methods and facilitates reproducibility. As for the end-to-end model, we compare with mainstream MLLM. The detailed introduction is as follows:

- **LLaVA-NeXT** (Liu et al., 2024) is a large-scale multimodal model that can process a variety of data types, including text, images and video. It has been trained on large multimodal datasets and shows strong performance in understanding multiple images and videos. In addition, LLaVA-NeXT has multi-language support, enabling it to understand multilingual text in images, including most European languages, Japanese, Korean, Arabic, Vietnamese and more.
- **Qwen2-VL** (Bai et al., 2023) can process images with different resolutions and aspect ratios and supports the understanding of multi-

Model	Size	BLEU				chrF++				METEOR			
		Base	Large	Small	Avg	Base	Large	Small	Avg	Base	Large	Small	Avg
EasyOCR_NLLB	-	6.4	6.3	6.2	6.3	19.7	19.1	17.9	18.9	17.6	16.4	14.9	16.3
DeepSeek-VL	7B	5.9	5.7	5.5	5.7	15.5	15.1	14.8	15.1	10.9	10.5	10.2	10.5
LLaVA-NeXT	7B	8.5	8.4	8.3	8.4	20.5	20.4	20.2	20.4	13.8	13.6	13.4	13.6
Qwen2-VL	7B	14.7	14.6	13.6	14.3	29.2	29.1	28.8	29.0	21.6	21.4	21.0	21.3
CogVLM2-LLaMA3	19B	8.6	8.4	8.0	8.3	19.9	19.5	19.2	19.5	14.5	14.4	14.2	14.4
MiniCPM-Llama3-V	8B	11.0	9.8	9.6	10.1	22.9	21.4	20.7	21.7	17.1	16.2	15.6	16.3
InternVL2	8B	14.3	14.1	13.4	13.9	27.4	26.6	26.5	26.8	23.8	23.3	22.8	23.3
Avg		9.9	9.6	9.2	9.6	22.1	21.6	21.1	21.6	17.0	16.6	16.0	16.5

Table 4: Comparison of the results of the image translation. It shows the results of comparison and evaluation of BLEU, chrF++, and METEOR using cascaded and end-to-end IT models on the MIT-10M test set.

lingual texts. Qwen2-VL has achieved world-leading performance in several visual comprehension benchmarks and has open-source 2B and 7B scale models.

- **CogVLM2-LLaMA3** (Hong et al., 2024) builds on LLaMA3 and shows exceptional performance on multimodal tasks involving images and text, especially when processing long texts and high-resolution images. It has a visual expert module that focuses on complex visual tasks and achieves deep integration of vision and speech through a sophisticated fusion strategy.
- **InternVL2** (Chen et al., 2024b) boosts the model’s performance through visual encoder improvements, dynamic strategies for high resolutions and high-quality bilingual data sets. It performs admirably on tasks such as OCR, multimodal assessment, mathematical reasoning and dialog with multiple interlocutors.
- **DeepSeek-VL** (Lu et al., 2024) is an open-source multimodal model designed to improve performance in real-world scenarios. It accepts high-resolution images as input and has general multimodal comprehension capabilities that process logic diagrams, web pages, formula recognition, scientific literature and natural images.
- **MiniCPM-Llama3-V** (Yao et al., 2024) is designed for consumer devices and focuses on providing advanced AI capabilities on resource-constrained devices such as cell phones. It can process various types of data, including text and images, and is capable of describing images, answering text questions

with one or more images, writing and debugging code, conducting dialogs with multiple images, conducting dialogs to understand videos, formatting JSON, and performing OCR in high resolution.

In order to make the model output the translated content stably, after several attempts, we use the prompt “*Translate the text in the image from {src_lang} into {tgt_lang}. Please output the translation directly without any explanation or other words.*” to obtain the inference result of the model.

Metrics. We use the *BLEU* score (Papineni et al., 2002) to assess the similarity between predicted translations and reference translations. This method calculates n-gram precision and applies a brevity penalty for shorter translations. We compute the *chrF++* score (Popović, 2017), which operates on both character and word-level n-grams. It effectively handles morphologically complex languages and is highly sensitive to spelling errors and minor mistakes. We employ the *METEOR* metric (Banerjee and Lavie, 2005) to evaluate translation quality. This metric incorporates stemming, synonym matching, and word reordering to improve its correlation with human evaluation.

5.2 Evaluate IT models with MIT-10M

In this paper, we systematically evaluate the performance of 7 state-of-the-art models on the MIT-10M, focusing on the comparison of cascaded models and multimodal end-to-end models on the image translation task. Table 4 shows the performance of all models on the MIT-10M. Among all models, EasyOCR_NLLB is the only cascaded model, while the others (e.g. LLaVA-NeXT, Qwen2-VL) are multimodal end-to-end models.

The Table 4 shows that DeepSeek-VL model perform poorly in IT task, with BLEU, chrF++,

	Easy	Medium	Hard
EasyOCR_NLLB	19.2	18.9	10.0
DeepSeek-VL	10.9	10.6	10.0
LLaVA-NeXT	15.6	14.3	10.8
Qwen2-VL	23.3	21.5	18.8
CogVLM2-LLaMA3	15.6	14.6	12.5
MiniCPM-LLaMA3-V	17.7	15.0	15.8
InternVL2	25.9	23.2	19.9
Avg	18.3	16.9	14.0

Table 5: METEOR with different levels of difficulty.

and METEOR of 5.7, 15.1, 10.5, well below average. Although the BLEU, chrF++ and METEOR values of EasyOCR_NLLB (6.3, 18.9 and 16.3, respectively) are close to the average, they lag behind the end-to-end models. In contrast, end-to-end models such as InternVL2 and Qwen2-VL perform better on IT task, with scores of 13.9, 16.8, 23.3 and 14.3, 29.0, 21.3 respectively, far outperforming those of the other models. In particular, Qwen2-VL achieved the best performance. The detail comparison of the BLEU of the individual models in multiple languages can be found in Appendix C.

We attribute this to the ability of the end-to-end models to generate text outputs directly from the image inputs and effectively utilize the correlations between images and text for deeper feature fusion and contextual understanding, resulting in superior performance in the fluency, consistency, and accuracy of the generated text. In contrast, cascaded models suffer from error propagation from the OCR phase, which limits their final translation quality. It is noteworthy that all models show relatively low evaluation results compared to their performance on other datasets. This highlights the realism of the dataset, which includes a wider range of challenges, including text with different fonts, colors and backgrounds, as well as instances of text occlusion and blurring. Consequently, the MIT-10M provides a more realistic and sophisticated benchmark for assessing the generalization ability of models.

5.3 Effect of Resolution and Difficulty

To further investigate the properties of the MIT-10M, we conducted finer analyzes at different image resolutions and task difficulties.

Table 4 illustrates the performance differences of the models in the image resolution tasks. A comparison of the average score of all models in different

sizes shows that almost all models perform better with larger images. The BLEU for the base, large and small resolutions are 9.9, 9.6 and 9.2 respectively, while the chrF++ and METEOR are 22.1, 21.6, 21.1 and 17.0, 16.6, 16.0 respectively. This emphasizes the importance of image resolution for the model’s understanding of the image content. High-resolution images provide more detailed information that helps the model to recognize and translate the text in the image more accurately.

Furthermore, MIT-10M contains tasks with three levels of difficulty: *Easy*, *Medium* and *Difficult* (see §Section 4.2). Table 5 illustrates the differences in model performance for the different tasks. When analyzing the performance of the models on the different difficulty levels, we observe a clear trend that the performance decreases with increasing difficulty. The average METEOR drops from 18.3 to 14.0, and this is particularly evident for Qwen2-VL, whose score decreases from 23.3 to 18.8.

This shows that the carefully designed difficulty distribution in the MIT-10M can effectively differentiate the performance of the models and provide the models with different difficulty levels.

As for the evaluation metrics, given the abundance of short sentences in MIT-10M, we recommend focusing on chrF++ and METEOR in addition to BLEU, as they are more suitable for evaluating the quality of short sentence translations.

Overall, our experimental results show that MIT-10M is a comprehensive, high quality and realistic dataset for image to text translation. MIT-10M represents a challenging benchmark for image translation research and serves as a valuable resource to advance the development of robust and versatile multimodal models.

5.4 Fine-tune with MIT-10M

To evaluate the effectiveness of our proposed MIT-10M for training image-to-text translation models, we perform fine-tuning experiments using the Qwen2-VL model as the baseline. During training, we use the following important hyperparameter settings: `per_device_train_batch_size`, `gradient_accumulation_steps`, `learning_rate`, `num_train_epochs`, `lr_scheduler_type` and `warmup_ratio` are set to 4, 32, 1.0e-4, 1, cosine and 0.1, respectively. Additionally, we utilize mixed precision training to accelerate the training process and reduce memory consumption.

We fine-tune the model with different subsets

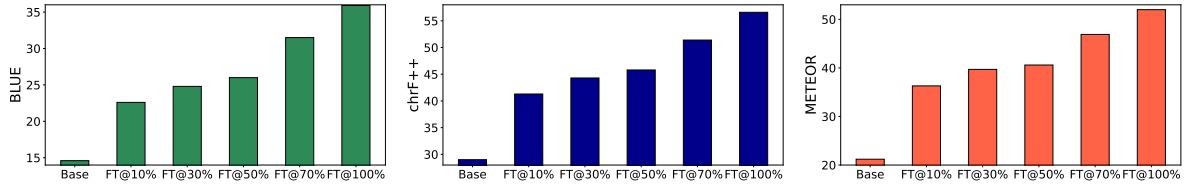


Figure 6: The performance comparison of the 5 model in the metrics: BLEU, chrF++, and METEOR. The 5 models are the base version and the version fine-tuned with 10%, 30%, 50%, 70%, and 100% MIT-10M train set.

Dataset	DE-EN	EN-DE	EN-FR	FR-EN
IIMT (Lan et al., 2024)	14.5	15.0	16.4	21.5
MIT-10M (ours)	16.4	16.2	19.3	25.3

Table 6: BLEU of fine-tuning on subsets of IIMT (170K) and MIT-10M (150K).

of the MIT-10M to investigate the effects of data size on model performance. Specifically, we take 10%, 30%, 50%, 70%, and 100% of the training set from the MIT-10M and compare their performance to analyze the effects of data size. As shown in Figure 6, the results demonstrate that increasing the size of the training data leads to significant improvements in all three metrics. For example, the BLEU increases from 14.6 for the base model (base@10%) to 35.9. Similarly, the chrF++ score increases continuously from 29.0 to 56.6 and the METEOR score improves from 21.2 to 52.0. It strongly emphasize the crucial role of data size in improving the performance of the image-to-text translation model: with richer training data, the model can capture finer correlations between images and text, improving translation accuracy and fluency. It is noteworthy that even with only 10% of the training data, the performance of the model is still significantly better than that of the base model, indicating the high quality of the MIT-10M. This result emphasizes that our dataset can effectively improve the performance of the model even with relatively small amounts of data.

5.5 Comparison with Existing IT Datasets

To further demonstrate the advantages of MIT-10M over existing real-world datasets, we perform fine-tuning experiments with comparable subsets. We compare the fine-tuning performance of MIT-10M with the IIMT dataset. We select subsets of both datasets to control the data size. We use 150K image-text pairs from MIT-10M and the slightly larger 170K pairs from IIMT. This choice ensures that any performance differences are not simply at-

tributable to the volume of training data, but rather to inherent dataset qualities. We focus on four language pairs: DE-EN, EN-DE, EN-FR, and FR-EN. The BLEU after fine-tuning can be found in Table 6. Even with less data, MIT-10M consistently outperforms IIMT in all four language pairs. The performance gains are particularly noticeable in the FR-EN pair, where MIT-10M shows a nearly 4 BLEU improvement. The results strongly support that MIT-10M is a valuable resource for multilingual image translation, enabling the training of more robust and generalizable models. This comparative analysis provides compelling evidence that the quality and diversity of MIT-10M contribute significantly to improved fine-tuning results, even when controlling for data size.

6 Conclusion

In this paper, we propose MIT-10M, a large-scale, high-quality dataset designed for multilingual image translation. MIT-10M contains over 10M image-text pairs in 14 languages and 840K high-resolution real-world images. We described the dataset creation process in detail and conducted a comprehensive analysis of the dataset across multiple dimensions. The results of experiments with different end-to-end IT models and evaluation metrics show that MIT-10M significantly improves multilingual translation performance and has strong generalization capabilities, especially for complex tasks. In the future, we will focus on further improving translation accuracy, extending support for more languages and processing even more diverse and complex images.

Acknowledgments

The present research was supported by the National Key Research and Development Program of China (No. 2024YFB3309702) and the National Natural Science Foundation of China Youth Fund (Grant No.62306210).

Limitations

In this work, we introduce MIT-10M, a novel, large-scale multilingual image translation parallel corpus that significantly advances research in the field of cross-lingual image translation. However, our approach is not without limitations. First, it is a major challenge to achieve a balanced representation of language and domain within the dataset. Furthermore, despite careful annotation and translation, the inherent complexity of multilingual data may lead to inaccuracies that could affect the reliability of the dataset. Furthermore, while the dataset has been extensively cleaned and filtered to address ethical concerns, including the removal of privacy and sensitive content issues, unforeseen possibilities remain. We acknowledge these limitations transparently to promote ethical research and encourage the community to make improvements.

References

- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Haithem Afli and Andy Way. 2016. [Integrating optical character recognition and machine translation of historical documents](#). In *Proceedings of the Workshop on Language Technology Resources and Tools for Digital Humanities (LT4DH)*, pages 109–116, Osaka, Japan. The COLING 2016 Organizing Committee.
- Jinze Bai, Shuai Bai, Shusheng Yang, Shijie Wang, Sinan Tan, Peng Wang, Junyang Lin, Chang Zhou, and Jingren Zhou. 2023. Qwen-vl: A frontier large vision-language model with versatile abilities. *arXiv preprint arXiv:2308.12966*.
- Satanjeev Banerjee and Alon Lavie. 2005. [METEOR: An automatic metric for MT evaluation with improved correlation with human judgments](#). In *Proceedings of the ACL Workshop on Intrinsic and Extrinsic Evaluation Measures for Machine Translation and/or Summarization*, pages 65–72, Ann Arbor, Michigan. Association for Computational Linguistics.
- Andong Chen, Lianzhang Lou, Kehai Chen, Xuefeng Bai, Yang Xiang, Muyun Yang, Tiejun Zhao, and Min Zhang. 2024a. [DUAL-REFLECT: Enhancing large language models for reflective translation through dual learning feedback mechanisms](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 693–704, Bangkok, Thailand. Association for Computational Linguistics.
- Kehai Chen, Masao Utiyama, Eiichiro Sumita, Rui Wang, and Min Zhang. 2022. [Synchronous refinement for neural machine translation](#). In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 2986–2996, Dublin, Ireland. Association for Computational Linguistics.
- Zhe Chen, Jiannan Wu, Wenhai Wang, Weijie Su, Guo Chen, Sen Xing, Muyan Zhong, Qinglong Zhang, Xizhou Zhu, Lewei Lu, et al. 2024b. Internvl: Scaling up vision foundation models and aligning for generic visual-linguistic tasks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 24185–24198.
- Zhuo Chen, Fei Yin, Xu-Yao Zhang, Qing Yang, and Chena-Lin Liu. 2021. [Cross-lingual text image recognition via multi-task sequence to sequence learning](#). In *2020 25th International Conference on Pattern Recognition (ICPR)*, pages 3122–3129.
- Team Gemini. 2024. [Gemini 1.5: Unlocking multimodal understanding across millions of tokens of context](#). *Preprint*, arXiv:2403.05530.
- Ryota Hinami, Shonosuke Ishiwatari, Kazuhiko Yasuda, and Yusuke Matsui. 2021. Towards fully automated manga translation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 12998–13008.
- Wenyi Hong, Weihai Wang, Ming Ding, Wenmeng Yu, Qingsong Lv, Yan Wang, Yean Cheng, Shiyu Huang, Junhui Ji, Zhao Xue, et al. 2024. Cogvlm2: Visual language models for image and video understanding. *arXiv preprint arXiv:2408.16500*.
- Puneet Jain, Orhan Firat, Qi Ge, and Sihang Liang. 2021. Image translation network. *Github.com*.
- Gant Laborde. 2024. [Deep nn for nsfw detection](#).
- Zhibin Lan, Liqiang Niu, Fandong Meng, Jie Zhou, Min Zhang, and Jinsong Su. 2024. [Translatotron-V\(ision\): An end-to-end model for in-image machine translation](#). In *Findings of the Association for Computational Linguistics ACL 2024*, pages 5472–5485, Bangkok, Thailand and virtual meeting. Association for Computational Linguistics.
- Zhibin Lan, Jiawei Yu, Xiang Li, Wen Zhang, Jian Luan, Bin Wang, Degen Huang, and Jinsong Su. 2023. Exploring better text image translation with multimodal codebook. *arXiv preprint arXiv:2305.17415*.
- Richardson Leonard. 2004. [Beautiful soup: a library designed for screen-scraping html and xml](#). *crummy.com*.
- Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. 2023. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. In *International conference on machine learning*, pages 19730–19742. PMLR.

- Yupu Liang, Yaping Zhang, Cong Ma, Zhiyang Zhang, Yang Zhao, Lu Xiang, Chengqing Zong, and Yu Zhou. 2024. [Document image machine translation with dynamic multi-pre-trained models assembling](#). In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 7084–7095, Mexico City, Mexico. Association for Computational Linguistics.
- Haotian Liu, Chunyuan Li, Yuheng Li, Bo Li, Yuanhan Zhang, Sheng Shen, and Yong Jae Lee. 2024. [Llava-next: Improved reasoning, ocr, and world knowledge](#).
- Haoyu Lu, Wen Liu, Bo Zhang, Bingxuan Wang, Kai Dong, Bo Liu, Jingxiang Sun, Tongzheng Ren, Zhuoshu Li, Yaofeng Sun, et al. 2024. [Deepseek-vl: towards real-world vision-language understanding](#). *arXiv preprint arXiv:2403.05525*.
- Cong Ma, Yaping Zhang, Mei Tu, Xu Han, Linghui Wu, Yang Zhao, and Yu Zhou. 2022. Improving end-to-end text image translation from the auxiliary text translation task. *2022 26th International Conference on Pattern Recognition (ICPR)*, pages 1664–1670.
- Cong Ma, Yaping Zhang, Mei Tu, Yang Zhao, Yu Zhou, and Chengqing Zong. 2023a. [E2timt: Efficient and effective modal adapter for text image machine translation](#). In *International Conference on Document Analysis and Recognition*, pages 70–88. Springer.
- Cong Ma, Yaping Zhang, Mei Tu, Yang Zhao, Yu Zhou, and Chengqing Zong. 2023b. [Multi-teacher knowledge distillation for end-to-end text image machine translation](#). In *International Conference on Document Analysis and Recognition*, pages 484–501. Springer.
- Elman Mansimov, Mitchell Stern, Mia Chen, Orhan Firat, Jakob Uszkoreit, and Puneet Jain. 2020. [Towards end-to-end in-image neural machine translation](#). In *Proceedings of the First International Workshop on Natural Language Processing Beyond Text*, pages 70–74, Online. Association for Computational Linguistics.
- Liqiang Niu, Fandong Meng, and Jie Zhou. 2024. [UMTIT: Unifying recognition, translation, and generation for multimodal text image translation](#). In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 16953–16972, Torino, Italia. ELRA and ICCL.
- NLLB Team, Marta R. Costa-jussà, James Cross, Onur Çelebi, Maha Elbayad, Kenneth Heafield, Kevin Heffernan, Elahe Kalbassi, Janice Lam, Daniel Licht, Jean Maillard, Anna Sun, Skyler Wang, Guillaume Wenzek, Al Youngblood, Bapi Akula, Loic Barrault, Gabriel Mejia-Gonzalez, Prangthip Hansanti, John Hoffman, Semaire Jarrett, Kaushik Ram Sadagopan, Dirk Rowe, Shannon Spruit, Chau Tran, Pierre Andrews, Necip Fazil Ayan, Shruti Bhosale, Sergey Edunov, Angela Fan, Cynthia Gao, Vedanuj Goswami, Francisco Guzmán, Philipp Koehn, Alexandre Mourachko, Christophe Ropers, Safiyyah Saleem, Holger Schwenk, and Jeff Wang. 2022. No language left behind: Scaling human-centered machine translation. *Github.com*.
- OpenAI. 2024. [Hello gpt-4o](#). *openai.com*.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. [Bleu: a method for automatic evaluation of machine translation](#). In *Proceedings of the 40th Annual Meeting on Association for Computational Linguistics, ACL '02*, page 311–318, USA. Association for Computational Linguistics.
- Maja Popović. 2017. [chrF++: words helping character n-grams](#). In *Proceedings of the Second Conference on Machine Translation*, pages 612–618, Copenhagen, Denmark. Association for Computational Linguistics.
- K Chandra Shekar, Maria Anisha Cross, and Vignesh Vasudevan. 2021. Optical character recognition and neural machine translation using deep learning techniques. In *Innovations in Computer Science and Engineering: Proceedings of 8th ICICSE*, pages 277–283. Springer.
- Tonghua Su, Shuchen Liu, and Shengjie Zhou. 2021. [Rtnet: An end-to-end method for handwritten text image translation](#). In *Document Analysis and Recognition-ICDAR 2021: 16th International Conference, Lausanne, Switzerland, September 5–10, 2021, Proceedings, Part II 16*, pages 99–113. Springer.
- Yasuhiko Watanabe, Yoshihiro Okada, Yeun-Bae Kim, and Tetsuya Takeda. 1998. Translation camera. In *Proceedings. Fourteenth International Conference on Pattern Recognition (Cat. No. 98EX170)*, volume 1, pages 613–617. IEEE.
- Jie Yang, Xilin Chen, Jing Zhang, Ying Zhang, and Alex Waibel. 2002. Automatic detection and translation of text from natural scenes. In *2002 IEEE International conference on acoustics, speech, and signal processing*, volume 2, pages II–2101. IEEE.
- Yuan Yao, Tianyu Yu, Ao Zhang, Chongyi Wang, Junbo Cui, Hongji Zhu, Tianchi Cai, Haoyu Li, Weilin Zhao, Zhihui He, et al. 2024. [Minicpm-v: A gpt-4v level mllm on your phone](#). *arXiv preprint arXiv:2408.01800*.
- Yaping Zhang, Shuai Nie, Wenju Liu, Xing Xu, Dongxiang Zhang, and Heng Tao Shen. 2019. Sequence-to-sequence domain adaptation network for robust text image recognition. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2740–2749.
- Yang Zhao, Lu Xiang, Junnan Zhu, Jiajun Zhang, Yu Zhou, and Chengqing Zong. 2020. [Knowledge graph enhanced neural machine translation via multi-task learning on sub-entity granularity](#). In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 4495–4505, Barcelona,

Spain (Online). International Committee on Computational Linguistics.

Meizhi Zhong, Lemao Liu, Kehai Chen, Mingming Yang, and Min Zhang. 2024. [Context consistency between training and inference in simultaneous machine translation](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 13465–13476, Bangkok, Thailand. Association for Computational Linguistics.

Shaolin Zhu, Shangjie Li, Yikun Lei, and Deyi Xiong. 2023. [PEIT: Bridging the modality gap with pre-trained models for end-to-end image translation](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 13433–13447, Toronto, Canada. Association for Computational Linguistics.

Shaolin Zhu, Leiyu Pan, Bo Li, and Deyi Xiong. 2024. [LANDeRMT: Detecting and routing language-aware neurons for selectively finetuning LLMs to machine translation](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 12135–12148, Bangkok, Thailand. Association for Computational Linguistics.

A Prompt

A.1 OCR Prompt

To enhance the precision of text identification within images, we employ the GPT-4o model for Optical Character Recognition (OCR) across the dataset. We meticulously construct a prompt for this task, ensuring that the textual data extracted is both accurate and reliable for subsequent analysis and translation processes. It is important to note that we instruct the GPT-4o model to output both English and Chinese content. This bilingual output capability is used to identify and translate the text, which is essential for the subsequent calibration phase.

```
Please perform text recognition on an image and translate the recognized text into other language. The output should be in json format\:  
'lang' refers to the original language of the recognized text, such as 'es', 'de', etc.  
'text' is the text recognized from the image. If there are multiple lines in the original text, please output them line by line.  
'en' is the result translated into English, corresponding line by line with the original text.  
'zh' is the result translated into Simplified Chinese.
```

```
Translate the following sentence into multiple languages, keeping the original paragraph structure and translating line by line.  
Sentence: ``text``  
Using this JSON schema:  
Lang = {"hi": str, "tr": str, "zh": str, "en": str, "fr": str, "de": str, "ja": str, "it": str, "ko": str, "th": str, "ru": str, "pt": str, "es": str, "ar": str}  
Return a `dict[str, Lang]`  
'en' is English.  
'zh' is Chinese (Simplified).  
'de' is German.  
'fr' is French.  
'ja' is Japanese.  
'it' is Italian.  
'ko' is Korean.  
'th' is Thai.  
'ru' is Russian.  
'pt' is Portuguese.  
'es' is Spanish.  
'hi' is Hindi.  
'tr' is Turkish.  
'ar' is Arabic.
```

A.2 Translation Prompt

We use the GPT-4 model to translate the extracted text into 13 target languages. To improve the quality of the translations, we use carefully crafted prompts. We also call Google Translate with the same prompts for cross-validation. The translations are further refined by filtering out those with ambiguous meanings based on semantic similarity scores to ensure the accuracy and reliability of the translations.

B Dataset Examples

Figure 7 shows the detail fields in the dataset, including 6 labels and 14 languages text.

C Experiments

Figure 8 show the detail comparison of the BLEU of each model in multiple language pairs.



ID: 10729452_EN-XX

SRC_TOEKN_LEN: 28 BOX_CNT: 5

DIFFICULTY: 3 (HARD) SRC_LANG: EN

CATE_ID: 24000 (Pet Supplies)

src_text: SCRATCHING POST\nFUN SKYLIGHT\nCLIMBING PLATFORM\nREINFORCED BASE\nExtra battens further enhance stability

```
{
  "tgt_lang": "ZH",
  "tgt_text": "抓撻柱\n有趣的天窗\n攀爬平台\n加固底座\n额外的横档进一步增强稳定性",
  "tgt_token_len": 48
}
{
  "tgt_lang": "JA",
  "tgt_text": "爪とぎポスト\n楽しいスカイライト\nクライミングプラットフォーム\nフォーム\n強化ベース\n追加のバテンが安定性をさらに高めます",
  "tgt_token_len": 57
}
{
  "tgt_lang": "KO",
  "tgt_text": "스크래칭 포스트\n재미있는 스카이라이트\n클라이밍 플랫폼\n강화된 베이스\n추가 배트가 안정성을 더욱 향상시킵니다",
  "tgt_token_len": 70
}
{
  "tgt_lang": "DE",
  "tgt_text": "KRATZPFEILER\nPASSIGE HIMMELICHT\nKLETTERRPLATTFORM\nVERSTÄRKTES FUNDAMENT\nZusätzliche Leisten erhöhen die Stabilität weiter",
  "tgt_token_len": 48
}
{
  "tgt_lang": "PT",
  "tgt_text": "POSTE ARRANHADOR\nLUZ DIVERTIDA\nPLATAFORMA DE ESCALADA\nBASE REFORÇADA\nBarras extras aumentam ainda mais a estabilidade",
  "tgt_token_len": 39
}
{
  "tgt_lang": "ES",
  "tgt_text": "POSTE RASCADOR\nDIVERTIDO VENTANAL\nPLATAFORMA DE ESCALADA\nBASE REFORZADA\nListones adicionales mejoran aún más la estabilidad",
  "tgt_token_len": 40
}
{
  "tgt_lang": "FR",
  "tgt_text": "POSTE À GRATTAGE\nLUMIÈRE AMUSANTE\nPLATEFORME D'ESCALADE\nBASE RENFORCÉE\nDes Lattes supplémentaires renforcent encore la stabilité",
  "tgt_token_len": 50
}
{
  "tgt_lang": "IT",
  "tgt_text": "POSTE PER GRATTARE\nLUCE DIVERTENTE\nPIATTAFORMA PER ARRAMPICARSI\nBASE RINFORZATA\nIstelli aggiuntivi migliorano ulteriormente la stabilità",
  "tgt_token_len": 46
}
{
  "tgt_lang": "RU",
  "tgt_text": "КОШКА-ЦАРЬ\nУДОВОЛЬСТВИЕ СОЛНЦА\nПЛАТФОРМА ДЛЯ ЛЕЗЕНИЯ\nУКРЕПЛЕННАЯ ОСНОВА\nДополнительные планки еще больше увеличиваю т стабильность",
  "tgt_token_len": 97
}
{
  "tgt_lang": "AR",
  "tgt_text": "تزيد القاعدة معززة\nالمنطقة تسلق\nالمنطقة سما، منطقة\nالمنطقة من الاسفل\nالاضافه من الاستقرار",
  "tgt_token_len": 62
}
{
  "tgt_lang": "TH",
  "tgt_text": "เสาข่วน\nแสงสนุก\nแพลตฟอร์มปีน\nฐานเสริม\nตะกั่วเพิ่มความปลอดภัย",
  "tgt_token_len": 76
}
{
  "tgt_lang": "TR",
  "tgt_text": "ÇİZME POSTASI\nEĞLENCELİ ÇATIDAN İŞIK ALMA DELİĞİ\nTIRMANMA PLATFORMU\nPEKİSTİRİLMİŞ TABAN\nEk çatalar kararlılığı daha da artırır",
  "tgt_token_len": 67
}
{
  "tgt_lang": "HI",
  "tgt_text": "खरब पोस्ट\nमजेदार रोमदान\nचढ़ाई प्लेटफॉर्म\nव्यवस्थित आधार\nअतिरिक्त बैटन स्थिरता को और बढ़ाते हैं",
  "tgt_token_len": 108
}
```



ID: 10360277_IT-XX

SRC_TOEKN_LEN: 39 BOX_CNT: 4

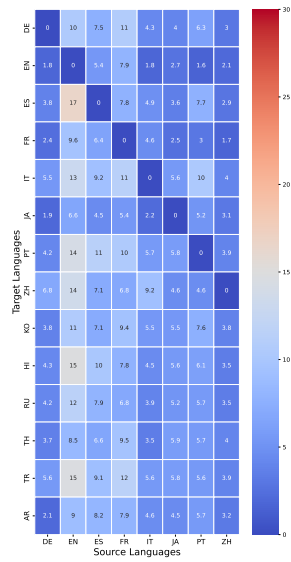
DIFFICULTY: 3 (HARD) SRC_LANG: IT

CATE_ID: 27000 (Tools & Home Improvement)

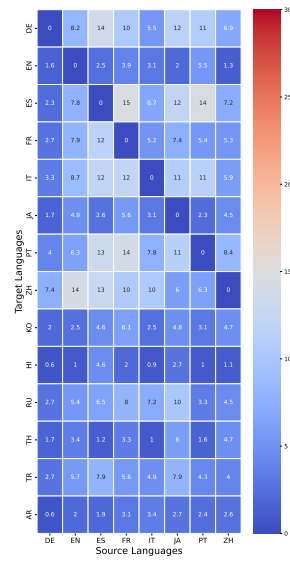
src_text: Soddisfare ogni tipo di esigenza\nAC 4000W/16A\nUSB C 15W/3.0\nUSB A 12W/2.4

```
{
  "tgt_lang": "EN",
  "tgt_text": "Meet every type of need\nAC 4000W/16A\nUSB C 15W/3.0\nUSB A 12W/2.4",
  "tgt_token_len": 34
}
{
  "tgt_lang": "ZH",
  "tgt_text": "满足各种需求\n交流电 4000W/16A\nUSB C 15W/3.0\nUSB A 12W/2.4",
  "tgt_token_len": 40
}
{
  "tgt_lang": "JA",
  "tgt_text": "あらゆるニーズに対応\nAC 4000W/16A\nUSB C 15W/3.0\nUSB A 12W/2.4",
  "tgt_token_len": 42
}
{
  "tgt_lang": "KO",
  "tgt_text": "모든 유형의 요구를 충족하다\nAC 4000W/16A\nUSB C 15W/3.0\nUSB A 12W/2.4",
  "tgt_token_len": 47
}
{
  "tgt_lang": "DE",
  "tgt_text": "Jeden Bedarf erfüllen\nAC 4000W/16A\nUSB C 15W/3.0\nUSB A 12W/2.4",
  "tgt_token_len": 36
}
{
  "tgt_lang": "ES",
  "tgt_text": "Satisfacer todo tipo de necesidades\nAC 4000W/16A\nUSB C 15W/3.0\nUSB A 12W/2.4",
  "tgt_token_len": 37
}
{
  "tgt_lang": "FR",
  "tgt_text": "Répondre à tous Les types de besoins\nAC 4000W/16A\nUSB C 15W/3.0\nUSB A 12W/2.4",
  "tgt_token_len": 39
}
{
  "tgt_lang": "RU",
  "tgt_text": "Удоволемся все типы потребности\nAC 4000W/16A\nUSB C 15W/3.0\nUSB A 12W/2.4",
  "tgt_token_len": 48
}
{
  "tgt_lang": "AR",
  "tgt_text": "الاحتياج كل نوع من\nAC 4000W/16A\nUSB C 15W/3.0\nUSB A 12W/2.4",
  "tgt_token_len": 49
}
{
  "tgt_lang": "TH",
  "tgt_text": "ตอบสนองทุกความต้องการ\nAC 4000W/16A\nUSB C 15W/3.0\nUSB A 12W/2.4",
  "tgt_token_len": 47
}
{
  "tgt_lang": "TR",
  "tgt_text": "Her türlü ihtiyaca cevap verin\nAC 4000W/16A\nUSB C 15W/3.0\nUSB A 12W/2.4",
  "tgt_token_len": 43
}
{
  "tgt_lang": "HI",
  "tgt_text": "हर तरह की जरूरत को पूरा करें\nAC 4000W/16A\nUSB C 15W/3.0\nUSB A 12W/2.4",
  "tgt_token_len": 57
}
```

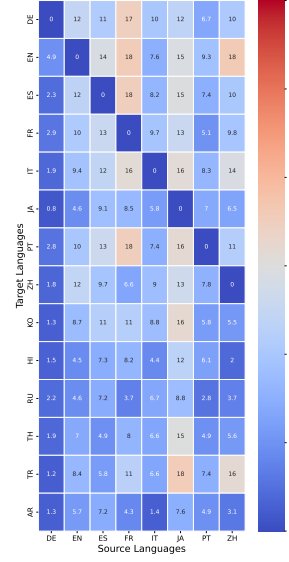
Figure 7: The detail field description for the MIT-10M dataset.



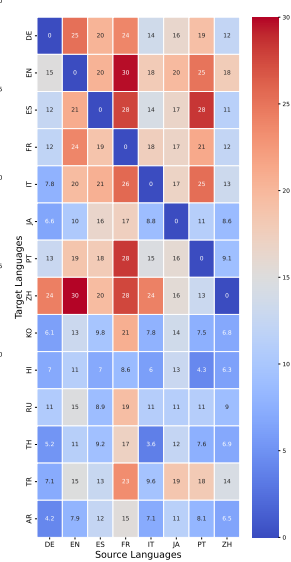
(a) EasyOCR_NLLB



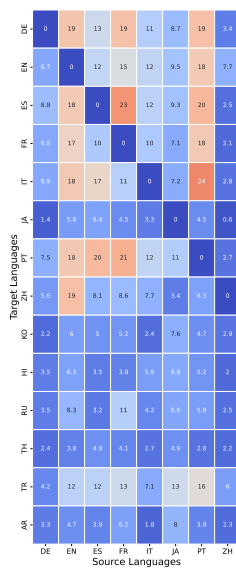
(b) DeepSeek-VL



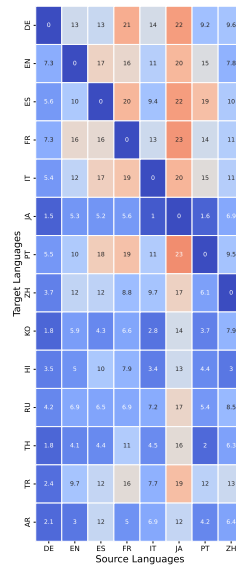
(c) LLaVA-NeXT



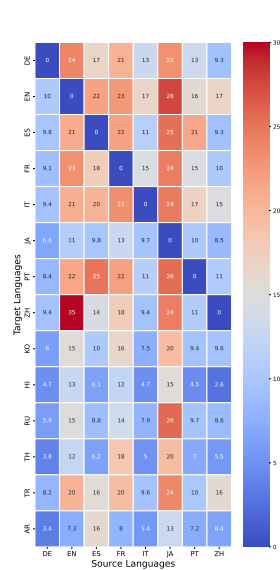
(d) Qwen2-VL



(e) CogVLM2-LLaMA3



(f) MiniCPM-LLaMA3-V



(g) InternVL2

Figure 8: The heatmap of various models in image translation tasks.