

# LARGE MULTILINGUAL MODELS PIVOT ZERO-SHOT MULTIMODAL LEARNING ACROSS LANGUAGES

Jinyi Hu<sup>1</sup> Yuan Yao<sup>1\*</sup> Chongyi Wang<sup>5</sup> Shan Wang<sup>5</sup> Yinxu Pan<sup>6</sup>  
 Qianyu Chen<sup>1</sup> Tianyu Yu<sup>1</sup> Hanghao Wu<sup>6</sup> Yue Zhao<sup>2</sup> Haoye Zhang<sup>1</sup>  
 Xu Han<sup>1,3</sup> Yankai Lin<sup>4</sup> Jiao Xue<sup>5</sup> Dahai Li<sup>5</sup> Zhiyuan Liu<sup>1\*</sup> Maosong Sun<sup>1\*</sup>

<sup>1</sup> Tsinghua University <sup>2</sup> Beijing University of Posts and Telecommunications

<sup>3</sup> Shanghai Artificial Intelligence Laboratory <sup>4</sup> Renmin University of China

<sup>5</sup> Zhihu Inc. <sup>6</sup> ModelBest Inc.

hu-jy21@mails.tsinghua.edu.cn

## ABSTRACT

Recently there has been a significant surge in multimodal learning in terms of both image-to-text and text-to-image generation. However, the success is typically limited to English, leaving other languages largely behind. Building a competitive counterpart in other languages is highly challenging due to the low-resource nature of non-English multimodal data (i.e., lack of large-scale, high-quality image-text data). In this work, we propose MPM, an effective training paradigm for training large multimodal models in non-English languages. MPM demonstrates that **M**ultilingual language models can **P**ivot zero-shot **M**ultimodal learning across languages. Specifically, based on a strong multilingual large language model, multimodal models pretrained on English-only image-text data can well generalize to other languages in a (quasi)-zero-shot manner, even surpassing models trained on image-text data in native languages. Taking Chinese as a practice of MPM, we build large multimodal models VISCPM in image-to-text and text-to-image generation, which achieve state-of-the-art (open-source) performance in Chinese. To facilitate future research, we open-source codes and model weights at <https://github.com/OpenBMB/VisCPM>.

## 1 INTRODUCTION

With the rapid advancement of powerful models such as GPT-4 (OpenAI, 2023) and Stable Diffusion (Rombach et al., 2022) in their multimodal capabilities, large multimodal models have emerged as the latest frontier in pursuing achieving Artificial General Intelligence (AGI). Generally, the multimodal generative capabilities across images and text can be divided into two categories: (i) In the field of image-to-text generation, prominent multimodal large language models like GPT-4 (OpenAI, 2023), LLaVA (Liu et al., 2023a) and InstructBLIP (Dai et al., 2023) exhibit remarkable multimodal conversational and reasoning abilities based on images; (ii) In the field of text-to-image generation, models such as Imagen (Saharia et al., 2022) and Stable Diffusion (Rombach et al., 2022) excel in generating highly realistic and relevant images based on text prompts. These models possess exceptional capabilities in processing images and text, profoundly reshaping the landscape of multimodal AI in both academia and industry.

However, the success of large multimodal models has mainly been achieved within the English community, while the multimodal capabilities in other non-English languages significantly trail behind. Bridging this gap is challenging due to the extensive image-text pair data requirements for training multimodal models. For instance, the pretraining of BLIP-2 (Li et al., 2023a) involves more than 100M high-quality image-text pairs, while Stable Diffusion (Rombach et al., 2022) utilizes more than 2B pairs. As a result of the paucity of such multimodal data resources in non-English languages, the progress of multimodal research in these languages remains hindered.

To address this challenge, we propose MPM, an effective training paradigm for large multimodal models in non-English languages. MPM utilizes the **M**ultilingual language model to **P**ivot **M**ultimodal

\*Corresponding authors.

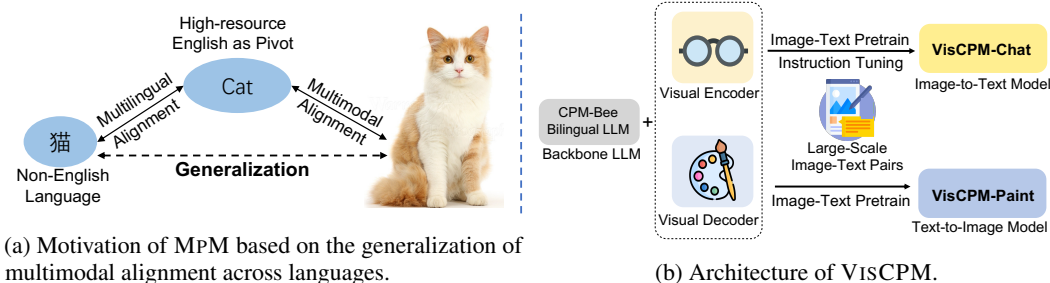


Figure 1: Overview of the motivation and architecture of MPM and VisCPM.

learning across languages and considers English, which contains substantial multimodal data resources, as a pivot between the visual signals and non-English languages which commonly lacking in multimodal data. MPM draws inspiration from the *Bilingual Dual-coding Theory* (Paivio, 2014; Clark & Paivio, 1991) that argue that *visual semantics are largely language-agnostic*. Intuitively, as portrayed in Fig. 1a, multilingual learners can effectively align the visual semantics with newly acquired language based on established multimodal and multilingual alignment. Simulating the human learning process, MPM also divides the non-English multimodal learning into two consecutive stages: multilingual alignment and multimodal alignment. The former focuses on building a multilingual model, while the latter culminates in a multimodal model spanning multiple languages.

Specifically, for multilingual alignment, MPM harnesses a pretrained multilingual large language model (LLM) as the backbone language model, which can provide aligned representations for different languages. Next, for the multimodal alignment, MPM trains the visual modules based on the multilingual model exclusively on English image-text pairs to align English and visual semantics. Similar to how humans learn, using the multilingual model as a pivot point, the resultant multimodal model naturally acquires zero-shot multimodal capability in other non-English languages.

Taking Chinese as a practical instance for MPM, we develop Chinese large multimodal models named VisCPM based on English-Chinese bilingual large language model CPM-Bee (Zhang et al., 2021). Notably, pretraining exclusively on English image-text pairs, the zero-shot performance of VisCPM in Chinese still surpasses that of existing Chinese multimodal models trained on image-text pairs in native Chinese. The promising performance of MPM in Chinese sheds light on its potential application in broader languages. Following the same training process, we further extend MPM to develop a multilingual multimodal conversation model supporting six languages based on the LLaMA (Touvron et al., 2023), including English, German, French, Spanish, Italian, and Portuguese.

In summary, the contributions of this paper are as follows: (i) We propose MPM, an effective training paradigm specifically designed for non-English languages lacking multimodal resources. Researchers worldwide can utilize MPM to rapidly adapt advanced multimodal learning methods from English to their respective languages. (ii) We develop a series of Chinese large multimodal models VisCPM as a practical application of MPM, which achieves state-of-the-art performance among open-source Chinese multimodal models. (iii) We open-source the model weights of VisCPM and provide experimental details, serving as a valuable reference for other researchers. (iv) We validate the generalization capability of VisCPM in diverse languages and develop the multilingual multimodal conversation model across six languages.

## 2 RELATED WORK

**Image-to-text Models.** Traditional image-to-text generation models mainly focus on the task of image caption and visual question answering (Yuan et al., 2021; Yu et al., 2022a; Wang et al., 2022b). Recently, the mainstream of image-to-text has turned to multimodal LLM, focusing on rendering LLM capable of multimodal interaction with users. These models connect the visual module and LLM with perceivers, such as BLIP-2 (Li et al., 2023a) and InstructBLIP (Dai et al., 2023) or linear projectors such as LLaVA (Liu et al., 2023a) and PandaGPT (Su et al., 2023). VPGTrans (Zhang et al., 2023) explores the transferability of visual modules across LLM. Meanwhile, many efforts have been dedicated to building multimodal instruction following datasets. LLaVA (Liu et al., 2023a)

and MiniGPT-4 (Zhu et al., 2023) build visual content-related dialog by transferring image captions into conversation data using GPT-4. InstructBLIP (Dai et al., 2023) and M<sup>3</sup>IT (Li et al., 2023b) incorporate downstream vision-language datasets to construct instruction data. LLaVA-RLHF (Sun et al., 2023) and RLHF-v (Yu et al., 2023b) propose multimodal RLHF for trustworthy behavior.

**Text-to-image Models.** In the early stages of text-to-image model development, generative adversarial networks (Zhu et al., 2019; Li et al., 2019) and auto-regressive generation (Esser et al., 2021) are popularly chosen architectures for text-to-image synthesis models (Li et al., 2019). More recently, large-scale diffusion-based text-to-image models such as DALLÉ-2 (Ramesh et al., 2022), Imagen (Saharia et al., 2022), and Stable Diffusion (Rombach et al., 2022) have taken center stage, demonstrating exceptional generative capabilities.

**Multilingual Multimodal Models.** The extension of multimodal models to include more languages has become a key research focus over the past few years. Researchers have made efforts to extend the powerful image-text model CLIP (Radford et al., 2021) to handle more languages using techniques of knowledge distillation (Carlsson et al., 2022; Hu et al., 2023) or contrastive learning (Bianchi et al., 2021; Chen et al., 2023). Other studies have aimed to create a universal framework for multilingual vision-language pertaining and simultaneously achieve multilingual and multimodal alignment (Ni et al., 2021; Zhou et al., 2021; Zeng et al., 2023). In the era of LLMs, PaLI (Chen et al., 2022) develops a 17B multilingual language-image model based on 10B image-text pairs spanning 100 languages. Ying-VLM (Li et al., 2023b) shows that instruction tuning in English can generalize to other languages. MultiFusion (Bellagente et al., 2023) discover that the multilingual language model can help cross-lingual transfer in text-to-image generation. In comparison, this work provides a more systematical formulation for training multilingual multimodal models and demonstrates that the zero-shot transfer performance of these models can surpass that of models trained on native-language multimodal data.

### 3 MPM TRAINING PARADIGM

In this section, we first present the formulation of multilingual multimodal learning and provide an overview of the training procedure of MPM. Following this, we detail the specific training procedures of MPM for both image-to-text and text-to-image generation.

#### 3.1 PROBLEM FORMULATION AND OVERVIEW

Multimodal learning can be formulated as modeling the relationship between images, denoted as  $x$ , and text, denoted as  $y$ , in a target language  $l_t$ . In this context, the image-to-text generation, which can be roughly summarized as generating description for input images, is to learn the conditional distribution  $p_\theta(y^{l_t}|x)$  parameterized by  $\theta$ ; the text-to-image generation, which is to synthesize relevant images given input text prompts, is to learn  $p_\phi(x|y^{l_t})$  parameterized by  $\phi$ .

In the vanilla setting, these conditional distributions are typically trained using image-text pairs  $\mathcal{D}_t = \{(x_i, y_i^{l_t})\}_{i=1}^N$  in the target language  $l_t$  (Radford et al., 2021; Yang et al., 2022). However, high-quality image-text pairs are extremely scarce in most languages. To mitigate the dependency on native image-text pairs, we introduce the pivot language  $l_p$ , which contains abundant multimodal pairs  $\mathcal{D}_p = \{(x_i, y_i^{l_p})\}_{i=1}^M$ , where  $M \gg N$ . Imitating the human learning mechanism that can naturally align visual concepts with various learned languages, MPM aims to transfer visual concepts learned in the pivot language to the target language.

MPM divides the multimodal learning process in target language  $l_t$  into two consecutive stages: **multilingual alignment** and **multimodal alignment**. For the multilingual alignment, MPM aims to establish the cross-lingual alignment for  $l_t$  and  $l_p$ . This is achieved by directly leveraging a pretrained multilingual LLM, denoted as  $f_\sigma$ , which can provide close hidden representations for text pair  $y^{l_t}$  and  $y^{l_p}$  with similar semantics, i.e.,  $f_\sigma(y^{l_t}) \approx f_\sigma(y^{l_p})$ . For the multimodal alignment, MPM utilize the sufficient multimodal resource  $\mathcal{D}_p$  in the pivot language and optimize the image-to-text objective  $p_\theta(y^{l_p}|x)$  and text-to-image objective  $p_\phi(x|y^{l_p})$ . In the following sections, we introduce the training process of **multimodal alignment** stage. It’s worth noting that MPM is agnostic to the specific model architecture and training method, which enables us to flexibly utilize existing highly effective model architectures and training techniques in each task.

### 3.2 IMAGE-TO-TEXT GENERATION

In image-to-text generation, we incorporate an image encoder module  $h_\xi$  parameterized by  $\xi$  to provide visual feature  $\mathbf{z} = h_\xi(x)$ . These visual features  $\mathbf{z}$  are then concatenated with the text embedding as input into the multilingual LLM. Following recent work to train multimodal conversation models (Zhu et al., 2023; Liu et al., 2023a), MPM’s training process for image-to-text generation consists of two sub-stages: Multimodal Pretraining and Instruction Tuning.

**Multimodal Pretraining.** In this sub-stage, we pretrain the visual module to align it with LLM on a large scale of image-text pairs using the language modeling objective:

$$\mathcal{L}_1(p_\theta, \mathcal{D}_p) = - \sum_{i=1}^M \log p_\theta(y_i^l | h_\xi(x_i)). \quad (1)$$

Here, we fix the parameters of LLM ( $\theta = \{\xi\}$ ) to prevent the powerful capabilities of LLM from being influenced by short texts in the image-text pairs.

**Instruction Tuning.** To enhance models’ capabilities in following human instructions, we conduct instruction tuning on elaborately curated multimodal instruction tuning datasets built by blending the existing multimodal instruction tuning datasets in the pivot language and their translated version in the target language. We denote this multilingual instruction tuning datasets as  $\mathcal{D}_i = \{x_k, y_{q,k}^l, y_{a,k}^l\}_{k=1}^S$ , where  $y_q^l$  is the instructions and  $y_a^l$  is the response in certain language  $l$ . Both the visual module and multilingual LLM are fine-tuned, i.e.,  $\theta = \{\xi, \sigma\}$ , by maximizing the probability of the response:

$$\mathcal{L}_2(p_\theta, \mathcal{D}_i) = - \sum_{k=1}^S \log p_\theta(y_{a,k}^l | h_\xi(x_k), f_\sigma(y_{q,k}^l)). \quad (2)$$

Interestingly, we find a *quasi-zero-shot* transfer capability of multilingual multimodal models in this scenario. If excluding the translated variant in the target language and solely performing instruction tuning using the pivot language, when given an image  $x$  and a question or an instruction  $y_q^{l_t}$  in the target language, the resultant model responds accurately though mostly in the pivot language. This can be attributed to the close resemblance between the hidden representation of instructions in two languages provided by the multilingual LLM, i.e.,  $f_\sigma(y_q^{l_p}) \approx f_\sigma(y_q^{l_t})$ . Consequently, we have  $p_\theta(y_a^{l_p} | h_\xi(x), f_\sigma(y_q^{l_p})) \approx p_\theta(y_a^{l_p} | h_\xi(x), f_\sigma(y_q^{l_t}))$ . Since both the pretraining and instruction tuning stages employ text components solely in the pivot language, the LLM can understand the question in the target language but cannot calibrate the response in the same language.

To stimulate the model to respond in the target language, MPM incorporates a small number of translated pairs in the target language during instruction tuning. In this way, MPM simultaneously improves the model’s instruction-following capability and calibrates the response language, ultimately realizing a multimodal chatbot in the target language.

### 3.3 TEXT-TO-IMAGE GENERATION

In the text-to-image generation, we adopt a similar architecture with Stable Diffusion (Rombach et al., 2022). It incorporates a denoising network  $g_\delta$  with a UNet architecture (Ronneberger et al., 2015) parameterized by  $\delta$  as an image decoder to generate images given the input prompt. The LLM  $f_\sigma$  and image decoder  $g_\delta$  are interconnected with cross-attention mechanism (Vaswani et al., 2017).

Diffusion models (Ho et al., 2020; Song et al., 2020) involve learning an iterative process of denoising a Gaussian noise into data distribution. The denoise network is optimized to remove the noise of noised image  $x_\tau$ , conditioned on the hidden states of text input provided by the LLM. The training objective is defined as follows:

$$\mathcal{L}(p_\phi, \mathcal{D}_p) = \mathbb{E}_{x, y^{l_p}, \varepsilon, \tau} [\|g_\delta(x_\tau, f_\sigma(y^{l_p}), \tau) - \varepsilon\|_2^2], \quad (3)$$

In this stage,  $\phi = \{\delta\}$ , i.e., the image decoder is trained to align with frozen LLM. In this way, when input with the unseen prompt in the target language  $y^{l_t}$ , the multilingual LLM  $f_\sigma$  can inherently provide a representation  $f_\sigma(y^{l_t})$  close to the seen representation  $f_\sigma(y^{l_p})$  of the pivot language prompt with similar semantics. Therefore, the capability of text-to-image in the target language can be seamlessly transferred from the pivot language in a zero-shot fashion, illustrated as  $g_\delta(x_\tau, f_\sigma(y^{l_t}), \tau) \approx g_\delta(x_\tau, f_\sigma(y^{l_p}), \tau)$ .

## 4 VISCPM

As a practice of MPM, we develop a series of large-scale Chinese multimodal models called VISCPM. We use Chinese as the target language and English as the pivot language. The Chinese-English bilingual language model CPM-Bee (Zhang et al., 2021) serves as the backbone multilingual LLM. We have two variations of the model: VISCPM-Chat for image-to-text multimodal conversation and VISCPM-Paint for text-to-image synthesis. In the following sections, we begin by providing an overview of existing multimodal datasets in Chinese and then introduce the training procedure of VISCPM-Chat and VISCPM-Paint.

### 4.1 ARE CHINESE MULTIMODAL DATASETS ENOUGH TO TRAIN A MULTIMODAL MODEL?

The largest publicly available multimodal dataset of native Chinese is Wukong (Gu et al., 2022), consisting of 100M image-text pair pairs. However, by visualizing the CLIP score computed with Chinese-CLIP (Yang et al., 2022), as shown in Fig. 2, and manually inspecting, as introduced in Appendix G, we discover that only a minor fraction of image-text pairs in Wukong possess semantically matched content. The poor quality of the dataset escalates existing data resource shortfalls. A straightforward method to enhance the count of Chinese image-text pairs is to translate the English image-text pairs into Chinese, which has been utilized in previous work (Wang et al., 2022a; Qiu et al., 2022). However, translation requires an external machine translation model, and translating a large-scale dataset used in pretraining consumes substantial computing resources. Also, as discussed in Sec.




Figure 2: The histogram of CLIP score of 100M Chinese image-text pairs from Wukong dataset. We set 0.18 as a moderate filtering threshold.

5.3, we practically find that incorporating translated image-text only has marginal improvement on the performance when using powerful multimodal LLM as backbone language model, which already possesses strong cross-lingual generalization capability. Based on this analysis, we argue that effectively utilizing the existing English data to achieve knowledge transfer in multimodal alignment is the key to developing a powerful Chinese large multimodal model.

### 4.2 VISCPM-CHAT

VISCPM-Chat is a Chinese-English bilingual multimodal chatbot capable of responding to users’ instructions based on the input image. VISCPM-Chat utilizes the Muffin architecture (Yu et al., 2023a) as the image encoder. Specifically, Muffin directly leverages a pretrained vision-language model BEiT-3 (Wang et al., 2023) as an inherent bridge module between vision and language. In the *multimodal pretraining stage*, the visual module is trained on 100M image-text pairs to align with the frozen LLM for 180K steps. In the *instruction tuning sub-stage*, we utilize bilingual versions of LLaVA 150K (Liu et al., 2023a) and UniMM-Chat (Yu et al., 2023a), and the English part of M<sup>3</sup>IT (Li et al., 2023b) to fine-tune the image encoder and LLM for 80K steps. The details of these datasets are presented in Appendix B.2. Due to the *quasi-zero-shot* phenomenon in Chinese introduced in Sec. 3.2, we incorporate certain Chinese data by translating LLaVA 150K and UniMM-Chat into Chinese using machine translation.<sup>1</sup>

To demonstrate the effect of Chinese image-text pairs, we also train an additional version of VISCPM-Chat, which adds additional Chinese image-text pairs during pretraining, including 20M native Chinese image-text pairs filtered from 100M Wukong (Gu et al., 2022) and 20M Zero-Chinese (Xie et al., 2022) dataset based on a CLIP score threshold greater than 0.18, and 136M image-text pairs translated from Laion-COCO dataset. We name this model as VISCPM-Chat+.

<sup>1</sup>We employ CPM-Bee for machine translation. CPM-Bee can obtain 41.6 BLEU score on WMT 2017 test set (Bojar et al., 2017). Except for special instructions, all translation processes in this work are carried out using the CPM-Bee model.

Table 1: Experimental results on LLaVA Test Set accessed by GPT-4. Con: Conversation, DD: Detailed Description, CR: Complex Reasoning, AVG: the average score of three tasks. The best/second best results are marked in **bold** and underlined, respectively.

Model		LLM Backbone	English				Chinese			
			Con	DD	CR	AVG	Con	DD	CR	AVG
English Model	MiniGPT-4 (Zhu et al., 2023)	Vicuna-13B	65.0	67.3	76.6	69.7	-	-	-	-
	InstructBLIP (Dai et al., 2023)	Vicuna-13B	81.9	68.0	91.2	80.5	-	-	-	-
	LLaVA (Liu et al., 2023a)	Vicuna-13B	<b>89.5</b>	<b>70.4</b>	<u>96.2</u>	<b>85.6</b>	-	-	-	-
En-Zh Bilingual Model	mPLUG-Owl (Ye et al., 2023)	BLOOMZ-7B	64.6	47.7	80.1	64.2	76.3	61.2	77.8	72.0
	VisualGLM (Du et al., 2022)	ChatGLM-6B	62.4	63.0	80.6	68.7	76.6	87.8	83.6	82.7
	Ziya-Visual (Wang et al., 2022a)	Ziya-LLaMA-13B	<u>82.7</u>	<u>69.9</u>	92.1	81.7	85.0	74.7	82.4	80.8
	Qwen-VL-Chat (Bai et al., 2023)	Qwen-7B	82.4	<u>76.9</u>	91.9	<u>83.8</u>	82.3	<b>93.4</b>	89.5	88.2
	ViSCPM-Chat	CPM-Bee-10B	81.4	69.2	93.1	81.4	<u>90.0</u>	87.4	<u>95.0</u>	<u>90.9</u>
	ViSCPM-Chat+	CPM-Bee-10B	80.1	67.1	<b>97.1</b>	81.5	<b>91.3</b>	<u>90.7</u>	<b>95.4</b>	<b>92.5</b>

### 4.3 ViSCPM-PAINT

ViSCPM-Paint is a text-to-image synthesis model that can accept prompts in both Chinese and English. ViSCPM-Paint employs the UNet in Stable Diffusion (Rombach et al., 2022) as the image decoder. To maintain the generative capability of UNet, the training process only involves the cross-attention layer of UNet and the linear transblock between the multilingual LLM and UNet. We optimize these parameters using an extensive dataset of English image-text pairs Laion-2B (Schuhmann et al., 2022) for 300K steps.

Similar to ViSCPM-Chat+, we train an additional version of ViSCPM-Paint, which is fine-tuned on Chinese image-text pairs. The component of these pairs are identical to that ViSCPM-Chat+ uses. We name this model as ViSCPM-Paint+.

## 5 EXPERIMENTS

### 5.1 EVALUATION OF ViSCPM-CHAT

#### 5.1.1 EVALUATION SETTING

**Baselines.** We compare ViSCPM-Chat with existing multimodal conversation models, which include the English-only models: MiniGPT-4 (Zhu et al., 2023), InstructBLIP (Dai et al., 2023), and LLaVA (Liu et al., 2023a), as well as Chinese-English bilingual models: mPLUG-Owl (Ye et al., 2023), VisualGLM (Du et al., 2022), Ziya-Visual (Wang et al., 2022a), and Qwen-VL-Chat (Bai et al., 2023). All of these Chinese-English models have performed large-scale training on both Chinese and English multimodal datasets. More details of baselines are presented in Appendix C.

**Evaluation Benchmark.** We assess the multimodal conversational abilities of the ViSCPM-Chat in both English and Chinese on the LLaVA Test Set (Liu et al., 2023a) and UniMM-Bench (Yu et al., 2023a). We use CPM-Bee to translate them into Chinese and manually check the translation results. Specifically, LLaVA comprehensively evaluates the model’s multimodal conversation capability from “conversation”, “detailed description”, and “complex reasoning”. UniMM-Bench, drawn from commonly used visual question-answering datasets—including OKVQA (Marino et al., 2019), AOKVQA (Schwenk et al., 2022), GQA (Hudson & Manning, 2019), and VQAv2 (Antol et al., 2015)—is designed particularly for evaluating multimodal conversation models on the abilities associated with reasoning, commonsense, and world knowledge. Considering that multimodal conversation models commonly generate responses with complete sentences to answer the question, traditional string-matching-based metrics, such as VQAScore (Antol et al., 2015), are not suitable in this context. Therefore, for the LLaVA Test Set and UniMM-Bench, we follow Liu et al. (2023a); Yu et al. (2023a) and employ GPT-4<sup>2</sup> to rate the model-generated responses and the reference answers. Further details of the LLaVA Test Set and UniMM-Bench are presented in Appendix D.3.

<sup>2</sup>Specifically, we use the GPT-4-0314 version to evaluate the responses.

Table 2: Evaluation of bilingual multimodal models on UniMM-Bench in Chinese. See results in English at Table 6.

Model	OKVQA	AOKVQA	GQA	VQAv2	AVG
mPLUG-Owl	52.8	55.8	60.8	56.7	56.5
VisualGLM	53.7	57.5	64.8	62.5	59.6
Ziya-Visual	<u>59.2</u>	58.1	61.9	59.1	59.6
Qwen-VL-Chat	58.5	<u>58.5</u>	<b>70.6</b>	<b>72.0</b>	64.9
VisCPM-Chat	<b>62.8</b>	<b>64.9</b>	<u>68.3</u>	<u>71.8</u>	<b>67.0</b>

Table 3: Results of extension to more languages based on Vicuna.

Lang	Con	DD	CR	AVG
<i>en</i>	87.6	76.8	101.0	88.6
<i>de</i>	90.8	80.8	93.6	88.7
<i>fr</i>	87.6	81.9	94.6	88.2
<i>es</i>	87.1	79.6	102.3	89.8
<i>it</i>	84.5	79.5	93.3	85.9
<i>pt</i>	81.3	81.6	92.4	85.1

### 5.1.2 EXPERIMENTAL RESULTS

**Quantitative Results.** The evaluation results on the LLaVA Test Set and UniMM-Bench are presented in Table 1 and Table 2. In **Chinese**, VisCPM-Chat achieves remarkable results and outperforms all baseline models on both benchmarks. It is worth noting that, unlike VisualGLM and Qwen-VL-Chat, which leverage a substantial number of Chinese image-text pairs during pretraining, VisCPM-Chat does not incorporate any Chinese multimodal data in its pretraining process. Despite this, VisCPM-Chat still presents superior performance on conversational abilities, commonsense, and world knowledge. Such results strongly demonstrate the effectiveness of MPM in transferring visual knowledge from English to Chinese. In **English**, the performance of VisCPM-Chat exceeds that of mPLUG-Owl, VisualGLM, InstructBLIP, and MiniGPT-4, and is roughly on par with Ziya-Visual and Qwen-VL-Chat while remaining comparable to the strong baseline LLaVA.

**Case Study.** In addition to acquiring broad world knowledge and commonsense, MPM training also implicitly stimulated VisCPM-Chat’s competency in transferring knowledge unique to Chinese culture. For example, as shown in Fig. 3, For instance, as illustrated in Fig. 3, VisCPM-Chat is capable of relating the Chinese classical poems present in Su Shi’s "Water Melody" according to the full moon and pavilion. Although this image-text pair deeply related to Chinese culture is not included in the training data, VisCPM-Chat successfully navigates such a non-trivial knowledge transfer through MPM training. For a more detailed analysis of cases involving VisCPM-Chat, please refer to Appendix B.3.

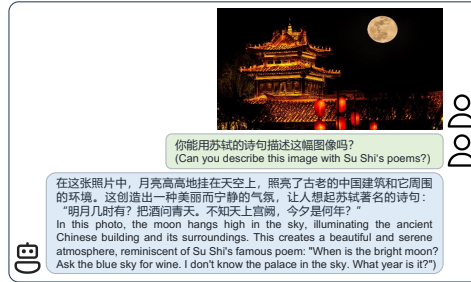


Figure 3: Multimodal conversation cases of VisCPM-Chat in Chinese.

### 5.1.3 GENERALIZATION TO MORE LANGUAGES

The remarkable results of VisCPM-Chat on Chinese of MPM illuminate the potential broader application to a more diverse set of languages. Specifically, we leverage the multilingual LLM LLaMA (Touvron et al., 2023) as the LLM backbone and consider German, French, Spanish, Portuguese, and Italian as target languages. Following the same training procedure as in VisCPM-Chat, we develop a multilingual multimodal chatbot proficiently supporting six languages. We begin by pretraining the visual encoder with LLaMA to achieve visual feature alignment in English image-text pairs. Next, we employ M2M-100 (Fan et al., 2021) to translate the instruction training set of LLaVA into five target languages. The original English instruction training set and the five translated sets are merged and shuffled, then used for fine-tuning the visual module and LLM.

Table 3 presents the evaluation results for English and five target languages on LLaVA Testset. See the results on IGLUE benchmark (Bugliarello et al., 2022) in Appendix D.5. Notably, for the relatively popular languages, such as German, French, and Spanish, the average scores exceed 88. Additionally, for Italian and Portuguese, the average scores are above 85. These results are highly encouraging as they demonstrate the chatbot’s coherent and accurate responses to visual-related questions in all six languages, even though the five targeted languages are simply blended during instruction tuning. The results in these languages validate the generalization and robustness of MPM in building powerful multimodal models in diverse linguistic contexts.



Table 4: Zero-shot FID on MSCOCO dataset.

Model	FID↓	
	En	Ch
GLIDE (Nichol et al., 2022)	12.2	-
Make-A-Scene (Gafni et al., 2022)	11.8	-
DALL-E-2 (Ramesh et al., 2022)	10.4	-
UniDiffuser (Bao et al., 2023)	9.7	-
CogView2 (Ding et al., 2022)	-	24.0
Stable Diffusion (Rombach et al., 2022)	<b>8.6</b>	-
AltDiffusion (Chen et al., 2023)	17.2	16.1
TaiyiDiffusion (Wang et al., 2022a)	-	15.6
ViSCPM-Paint	<u>9.5</u>	<u>10.9</u>
ViSCPM-Paint+	9.9	<b>9.6</b>




Figure 4: Results of human evaluation for text-to-images on Chinese Drawbench.

## 5.2 EVALUATION OF ViSCPM-PAINT

### 5.2.1 EVALUATION SETTING

We compare ViSCPM-Paint with several strong text-to-image model, which including English-only models: GLIDE (Nichol et al., 2022), Make-A-Scene (Gafni et al., 2022), DALL-E-2 (Ramesh et al., 2022), UniDiffuser (Bao et al., 2023), and Chinese or Chinese-English bilingual text-to-image models: CogView2 (Ding et al., 2022), AltDiffusion (Chen et al., 2023) and TaiyiDiffusion (Wang et al., 2022a). We mainly compare ViSCPM-Paint with AltDiffusion (Chen et al., 2023) and TaiyiDiffusion (Wang et al., 2022a). More details of baselines are presented in Appendix C.

### 5.2.2 AUTOMATIC EVALUATION

For the text-to-image tasks, we assess the zero-shot Frechet Inception Distance (FID) (Heusel et al., 2017) and the CLIP score (Ramesh et al., 2021) using the MSCOCO validation set (Lin et al., 2014). We sample 30K prompts from MSCOCO, and for the Chinese evaluations, the text prompts are translated from the original English captions. We maintain the same sampling configuration for ViSCPM-Paint, AltDiffusion, and TaiyiDiffusion and grid search the optimal FID across eight separate classifier guidance scales.

We present the zero-shot FID on MSCOCO validation in Table 4. In **Chinese**, ViSCPM-Paint achieves the best FID performance. By solely training on English image-text pairs, ViSCPM-Paint displays a significant advantage over AltDiffusion (Chen et al., 2023) and TaiyiDiffusion (Wang et al., 2022a). In **English**, the performance of ViSCPM-Paint is comparable to existing powerful text-to-image models, such as Stable Diffusion (Rombach et al., 2022) and UniDiffuser (Bao et al., 2023). See more analysis of fidelity and alignment trade-offs in Appendix E.2.

### 5.2.3 HUMAN EVALUATION

Following previous work (Yu et al., 2022b; Chang et al., 2023), we perform a human evaluation to have a more comprehensive understanding of model performance. Due to the lack of human evaluation benchmarks in Chinese, we create a Chinese text-to-image human evaluation benchmark, named *Chinese Drawbench*. *Chinese Drawbench* consists of 174 prompts and evaluates text-to-image models' proficiency across different aspects. We conduct a human evaluation involving ViSCPM-Paint, AltDiffusion (Chen et al., 2023), and TaiyiDiffusion (Wang et al., 2022a) on *Chinese Drawbench*. We ask five independent annotators to judge the best image generated by three models for each prompt. Their evaluation criteria included Overall, Alignment, and Fidelity. More details of human evaluation and *Chinese Drawbench* are provided in Appendix E.3.

Figure 4 presents the human evaluation results. The figure shows the preference shares of 5 evaluators for each generated image with different marginals over consensus in three aspects. Notably, ViSCPM-Paint receives the strongest preference across all three evaluation aspects. In each aspect, more than 40 percent of images generated by ViSCPM-Paint are the favored choice. Impressively, in the more objective metrics of Overall and Alignment, ViSCPM-Paint earns more than 20 percent 5/5



Table 5: Performance in Chinese with different combinations of dataset languages in each training stage. IT means instruction tuning.

(a) Image-to-text: Evaluation results on LLaVA Test Set.					(b) Text-to-image: zero-shot FID on MSCOCO.			
Dataset Language		Con	DD	CR	AVG	Dataset Language		FID↓
Pretrain	IT					Pretrain	Finetune	
Native Chinese	Chinese	82.4	81.8	91.7	85.5	Native Chinese	-	15.1
English	Chinese	84.0	<b>83.6</b>	91.2	86.3	English	-	10.9
English	Bilingual	85.4	81.4	<b>96.6</b>	88.0	English	Native Chinese w/o Filter	12.8
English+ Native Chinese	Bilingual	89.1	82.3	91.4	87.8	English	Native Chinese	10.7
English+ Native Chinese+ Translated Chinese	Bilingual	<b>90.3</b>	81.4	92.1	<b>88.2</b>	English	Native Chinese + Translated Chinese	<b>9.6</b>

preference. These results strongly demonstrate the superior quality of VISCPM-Paint’s generated images compared to the two baseline models.

### 5.3 ABLATION STUDY

We conduct the ablation study of dataset languages to investigate the impact of varied dataset language combinations on multimodal models’ performance in image-to-text and text-to-image tasks. For better efficiency, we only use LLaVA 150K for instruction tuning in the image-to-text task. The detailed configurations of each experiment are reported in Appendix F. Based on the results presented in Table 5a and Table 5b, We have the following observations: (i) *Relying solely on a native Chinese dataset is insufficient for achieving good performance in both image-to-text and text-to-image tasks.* Models trained exclusively on native Chinese datasets yield worse scores, with the image-to-text obtaining an 85.5 average score and the text-to-image model attaining a 15.1 FID. (ii) *English data plays a crucial role in improving the chat capability of the model in Chinese during the instruction tuning stage.* When the image-to-text model is pretrained on a large dataset of English but then fine-tuned using the monolingual Chinese instruction tuning dataset, its average performance experiences a decline from 88.0 to 86.3 compared to the model utilizing the bilingual instruction tuning dataset. (iii) *The filtering process applied to the native dataset is essential for the Chinese performance.* In the text-to-image task, after pretraining on English data and then fine-tuning with unfiltered native Chinese multimodal pairs, the FID worsens from 10.9 in the zero-shot scenario to 12.8. (iv) *Incorporating the native Chinese dataset and translated Chinese dataset yields a marginal improvement in the model performance.* Specifically, adding native Chinese data into the pretraining stage of the image-to-text model or fine-tuning stage of the text-to-image model only results in a 0.2 variation of metrics, while further mixing translated datasets improves the FID from 10.7 to 9.6. The improvement from VISCPM-Chat to VISCPM-Chat+, shown in Table 1, also confirms the same results. We present a more systematic exploration of the influence on the response language accuracy and content quality with different percentages of the Chinese instruction tuning datasets in Appendix D.4.

## 6 CONCLUSION

In this work, we introduce MPM, an innovative training paradigm designed for effectively training large multimodal models in non-English languages. By utilizing a multilingual LLM as a pivotal intermediary between vision signals and target languages, MPM facilitates the efficient transfer of multimodal alignment knowledge across different languages. Based on MPM, we develop a series of open-sourced Chinese large multimodal models VISCPM, which show remarkable capability in Chinese image-to-text and text-to-image tasks. Our experimental results demonstrate that by solely relying on English multimodal data, VISCPM can achieve the SOTA performance among Chinese open-sourced multimodal models. We further scale the scope of language by constructing a versatile multimodal chatbot that supports six distinct languages. We believe that the effectiveness of MPM will contribute to the development of large multimodal models worldwide, thereby fostering the development of sophisticated multimodal models across various languages and cultures.

## ACKNOWLEDGEMENT

This work is supported by the National Key R&D Program of China (No.2022ZD0160501).

## REFERENCES

- Stanislaw Antol, Aishwarya Agrawal, Jiasen Lu, Margaret Mitchell, Dhruv Batra, C. Lawrence Zitnick, and Devi Parikh. VQA: Visual Question Answering. In *Proceedings of ICCV*, 2015.
- Jinze Bai, Shuai Bai, Shusheng Yang, Shijie Wang, Sinan Tan, Peng Wang, Junyang Lin, Chang Zhou, and Jingren Zhou. Qwen-vl: A frontier large vision-language model with versatile abilities. *arXiv preprint arXiv:2308.12966*, 2023.
- Fan Bao, Shen Nie, Kaiwen Xue, Chongxuan Li, Shi Pu, Yaole Wang, Gang Yue, Yue Cao, Hang Su, and Jun Zhu. One transformer fits all distributions in multi-modal diffusion at scale. In Andreas Krause, Emma Brunskill, Kyunghyun Cho, Barbara Engelhardt, Sivan Sabato, and Jonathan Scarlett (eds.), *Proceedings of ICML*, volume 202, pp. 1692–1717. PMLR, 23–29 Jul 2023.
- Marco Bellagente, Manuel Brack, Hannah Teufel, Felix Friedrich, Björn Deiseroth, Constantin Eichenberg, Andrew Dai, Robert Baldock, Souradeep Nanda, Koen Oostermeijer, et al. Multifusion: Fusing pre-trained models for multi-lingual, multi-modal image generation. *arXiv preprint arXiv:2305.15296*, 2023.
- Federico Bianchi, Giuseppe Attanasio, Raphael Pisoni, Silvia Terragni, Gabriele Sarti, and Sri Lakshmi. Contrastive language-image pre-training for the italian language. *arXiv preprint arXiv:2108.08688*, 2021.
- Ondřej Bojar, Rajen Chatterjee, Christian Federmann, Yvette Graham, Barry Haddow, Shujian Huang, Matthias Huck, Philipp Koehn, Qun Liu, Varvara Logacheva, et al. Findings of the 2017 conference on machine translation (wmt17). Association for Computational Linguistics, 2017.
- Emanuele Bugliarello, Fangyu Liu, Jonas Pfeiffer, Siva Reddy, Desmond Elliott, Edoardo Maria Ponti, and Ivan Vulić. IGLUE: A benchmark for transfer learning across modalities, tasks, and languages. In Kamalika Chaudhuri, Stefanie Jegelka, Le Song, Csaba Szepesvari, Gang Niu, and Sivan Sabato (eds.), *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pp. 2370–2392. PMLR, 17–23 Jul 2022.
- Fredrik Carlsson, Philipp Eisen, Faton Rekathati, and Magnus Sahlgren. Cross-lingual and multilingual clip. In *Proceedings of LREC*, pp. 6848–6854, 2022.
- Huiwen Chang, Han Zhang, Jarred Barber, AJ Maschinot, Jose Lezama, Lu Jiang, Ming-Hsuan Yang, Kevin Murphy, William T Freeman, Michael Rubinstein, et al. Muse: Text-to-image generation via masked generative transformers. *arXiv preprint arXiv:2301.00704*, 2023.
- Soravit Changpinyo, Piyush Sharma, Nan Ding, and Radu Soricut. Conceptual 12M: Pushing web-scale image-text pre-training to recognize long-tail visual concepts. In *Proceedings of CVPR*, pp. 3558–3568, June 2021.
- Xi Chen, Xiao Wang, Soravit Changpinyo, AJ Piergiovanni, Piotr Padlewski, Daniel Salz, Sebastian Goodman, Adam Grycner, Basil Mustafa, Lucas Beyer, et al. Pali: A jointly-scaled multilingual language-image model. In *Proceedings of ICLR*, 2022.
- Zhongzhi Chen, Guang Liu, Bo-Wen Zhang, Qinghong Yang, and Ledell Wu. AltCLIP: Altering the language encoder in CLIP for extended language capabilities. In *Findings of ACL*, pp. 8666–8682, Toronto, Canada, July 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.findings-acl.552.
- Schuhmann Christoph, Köpf Andreas, Vencu Richard, Coombes Theo, and Beaumont Romain. Laion COCO: 600M synthetic captions from Laion2B-En, 2022. URL <https://laion.ai/blog/laion-coco/>.

- James M Clark and Allan Paivio. Dual coding theory and education. *Educational psychology review*, 3:149–210, 1991.
- Wenliang Dai, Junnan Li, Dongxu Li, Anthony Meng Huat Tiong, Junqi Zhao, Weisheng Wang, Boyang Li, Pascale Fung, and Steven Hoi. InstructBLIP: Towards general-purpose vision-language models with instruction tuning, 2023.
- Abhishek Das, Satwik Kottur, Khushi Gupta, Avi Singh, Deshraj Yadav, José M.F. Moura, Devi Parikh, and Dhruv Batra. Visual Dialog. In *Proceedings of CVPR*, 2017.
- Ming Ding, Wendi Zheng, Wenyi Hong, and Jie Tang. Cogview2: Faster and better text-to-image generation via hierarchical transformers. *NeurIPS*, 35:16890–16902, 2022.
- Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. In *Proceedings of ICLR*, 2020.
- Zhengxiao Du, Yujie Qian, Xiao Liu, Ming Ding, Jiezhong Qiu, Zhilin Yang, and Jie Tang. Glm: General language model pretraining with autoregressive blank infilling. In *Proceedings of ACL*, pp. 320–335, 2022.
- Patrick Esser, Robin Rombach, and Bjorn Ommer. Taming transformers for high-resolution image synthesis. In *Proceedings of CVPR*, pp. 12873–12883, 2021.
- Angela Fan, Shruti Bhosale, Holger Schwenk, Zhiyi Ma, Ahmed El-Kishky, Siddharth Goyal, Mandeep Baines, Onur Celebi, Guillaume Wenzek, Vishrav Chaudhary, et al. Beyond english-centric multilingual machine translation. *JMLR*, 22(1):4839–4886, 2021.
- Oran Gafni, Adam Polyak, Oron Ashual, Shelly Sheynin, Devi Parikh, and Yaniv Taigman. Make-a-Scene: Scene-based text-to-image generation with human priors. In *Proceedings of ECCV*, pp. 89–106. Springer, 2022.
- Jiaxi Gu, Xiaojun Meng, Guansong Lu, Lu Hou, Minzhe Niu, Hang Xu, Xiaodan Liang, Wei Zhang, Xin Jiang, and Chunjing Xu. Wukong: 100 million large-scale chinese cross-modal pre-training dataset and a foundation framework, 2022.
- Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. GANs trained by a two time-scale update rule converge to a local nash equilibrium. *NeurIPS*, 30, 2017.
- Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *NeurIPS*, 33: 6840–6851, 2020.
- Jinyi Hu, Xu Han, Xiaoyuan Yi, Yutong Chen, Wenhao Li, Zhiyuan Liu, and Maosong Sun. Efficient cross-lingual transfer for chinese stable diffusion with images as pivots. *arXiv preprint arXiv:2305.11540*, 2023.
- Drew A Hudson and Christopher D Manning. Gqa: A new dataset for real-world visual reasoning and compositional question answering. In *Proceedings of CVPR*, pp. 6700–6709, 2019.
- Gabriel Ilharco, Mitchell Wortsman, Ross Wightman, Cade Gordon, Nicholas Carlini, Rohan Taori, Achal Dave, Vaishaal Shankar, Hongseok Namkoong, John Miller, Hannaneh Hajishirzi, Ali Farhadi, and Ludwig Schmidt. Openclip, July 2021. If you use this software, please cite it as below.
- Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C. Berg, Wan-Yen Lo, Piotr Dollár, and Ross Girshick. Segment anything. *arXiv:2304.02643*, 2023.
- Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, David A Shamma, et al. Visual Genome: Connecting language and vision using crowdsourced dense image annotations. *IJCV*, 123:32–73, 2017.

- Bowen Li, Xiaojuan Qi, Thomas Lukasiewicz, and Philip Torr. Controllable text-to-image generation. *NeurIPS*, 32, 2019.
- Junnan Li, Dongxu Li, Caiming Xiong, and Steven Hoi. BLIP: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In Kamalika Chaudhuri, Stefanie Jegelka, Le Song, Csaba Szepesvari, Gang Niu, and Sivan Sabato (eds.), *Proceedings of ICML*, volume 162, pp. 12888–12900. PMLR, 17–23 Jul 2022.
- Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. BLIP-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. In Andreas Krause, Emma Brunskill, Kyunghyun Cho, Barbara Engelhardt, Sivan Sabato, and Jonathan Scarlett (eds.), *Proceedings of ICML*, volume 202, pp. 19730–19742. PMLR, 23–29 Jul 2023a.
- Lei Li, Yuwei Yin, Shicheng Li, Liang Chen, Peiyi Wang, Shuhuai Ren, Mukai Li, Yazheng Yang, Jingjing Xu, Xu Sun, et al. M<sup>3</sup> IT: A large-scale dataset towards multi-modal multilingual instruction tuning. *arXiv preprint arXiv:2306.04387*, 2023b.
- Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft COCO: Common objects in context. In *Proceedings of ECCV*, pp. 740–755. Springer, 2014.
- Fangyu Liu, Emanuele Bugliarello, Edoardo Maria Ponti, Siva Reddy, Nigel Collier, and Desmond Elliott. Visually grounded reasoning across languages and cultures. In Marie-Francine Moens, Xuanjing Huang, Lucia Specia, and Scott Wen-tau Yih (eds.), *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pp. 10467–10485, Online and Punta Cana, Dominican Republic, November 2021. Association for Computational Linguistics.
- Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. *arXiv preprint arXiv:2304.08485*, 2023a.
- Shanyuan Liu, Dawei Leng, and Yuhui Yin. Bridge diffusion model: bridge non-english language-native text-to-image diffusion model with english communities. *arXiv preprint arXiv:2309.00952*, 2023b.
- Shilong Liu, Zhaoyang Zeng, Tianhe Ren, Feng Li, Hao Zhang, Jie Yang, Chunyuan Li, Jianwei Yang, Hang Su, Jun Zhu, et al. Grounding DINO: Marrying dino with grounded pre-training for open-set object detection. *arXiv preprint arXiv:2303.05499*, 2023c.
- Kenneth Marino, Mohammad Rastegari, Ali Farhadi, and Roozbeh Mottaghi. OK-VQA: A visual question answering benchmark requiring external knowledge. In *Proceedings of CVPR*, 2019.
- Niklas Muennighoff, Thomas Wang, Lintang Sutawika, Adam Roberts, Stella Biderman, Teven Le Scao, M Saiful Bari, Sheng Shen, Zheng-Xin Yong, Hailey Schoelkopf, Xiangru Tang, Dragomir Radev, Alham Fikri Aji, Khalid Almubarak, Samuel Albanie, Zaid Alyafeai, Albert Webson, Edward Raff, and Colin Raffel. Crosslingual generalization through multitask finetuning, 2023.
- Minheng Ni, Haoyang Huang, Lin Su, Edward Cui, Taroon Bharti, Lijuan Wang, Dongdong Zhang, and Nan Duan. M3P: Learning universal representations via multitask multilingual multimodal pre-training. In *Proceedings of CVPR*, pp. 3977–3986, 2021.
- Alexander Quinn Nichol, Prafulla Dhariwal, Aditya Ramesh, Pranav Shyam, Pamela Mishkin, Bob McGrew, Ilya Sutskever, and Mark Chen. GLIDE: Towards photorealistic image generation and editing with text-guided diffusion models. In Kamalika Chaudhuri, Stefanie Jegelka, Le Song, Csaba Szepesvari, Gang Niu, and Sivan Sabato (eds.), *Proceedings of ICML*, volume 162, pp. 16784–16804. PMLR, 17–23 Jul 2022.
- OpenAI. GPT-4 technical report. *arXiv preprint arXiv:2303.08774*, 2023.
- Allan Paivio. *Bilingual Dual Coding Theory and Memory*, pp. 41–62. Springer New York, New York, NY, 2014. ISBN 978-1-4614-9218-4. doi: P10.1007/978-1-4614-9218-4\_3.

- Jonas Pfeiffer, Gregor Geigle, Aishwarya Kamath, Jan-Martin Steitz, Stefan Roth, Ivan Vulić, and Iryna Gurevych. xGQA: Cross-lingual visual question answering. In Smaranda Muresan, Preslav Nakov, and Aline Villavicencio (eds.), *Findings of the Association for Computational Linguistics: ACL 2022*, pp. 2497–2511, Dublin, Ireland, May 2022. Association for Computational Linguistics.
- Chen Qiu, Dan Oneață, Emanuele Bugliarello, Stella Frank, and Desmond Elliott. Multilingual multimodal learning with machine translated text. In Yoav Goldberg, Zornitsa Kozareva, and Yue Zhang (eds.), *Findings of the Association for Computational Linguistics: EMNLP 2022*, pp. 4178–4193, Abu Dhabi, United Arab Emirates, December 2022. Association for Computational Linguistics.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *ICML*, pp. 8748–8763. PMLR, 2021.
- Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. Zero-shot text-to-image generation. In Marina Meila and Tong Zhang (eds.), *Proceedings of ICML*, volume 139, pp. 8821–8831. PMLR, 18–24 Jul 2021.
- Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*, 2022.
- Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of CVPR*, pp. 10684–10695, 2022.
- Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-Net: Convolutional networks for biomedical image segmentation. In *MICCAI*, pp. 234–241. Springer, 2015.
- Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily L Denton, Kamyar Ghasemipour, Raphael Gontijo Lopes, Burcu Karagol Ayan, Tim Salimans, et al. Photorealistic text-to-image diffusion models with deep language understanding. *NeurIPS*, 35:36479–36494, 2022.
- Christoph Schuhmann, Richard Vencu, Romain Beaumont, Robert Kaczmarczyk, Clayton Mullis, Aarush Katta, Theo Coombes, Jenia Jitsev, and Aran Komatsuzaki. Laion-400M: Open dataset of clip-filtered 400 million image-text pairs. *arXiv preprint arXiv:2111.02114*, 2021.
- Christoph Schuhmann, Romain Beaumont, Richard Vencu, Cade Gordon, Ross Wightman, Mehdi Cherti, Theo Coombes, Aarush Katta, Clayton Mullis, Mitchell Wortsman, et al. Laion-5B: An open large-scale dataset for training next generation image-text models. *NeurIPS*, 35:25278–25294, 2022.
- Dustin Schwenk, Apoorv Khandelwal, Christopher Clark, Kenneth Marino, and Roozbeh Mottaghi. A-OKVQA: A benchmark for visual question answering using world knowledge. *arXiv*, 2022.
- Piyush Sharma, Nan Ding, Sebastian Goodman, and Radu Soricut. Conceptual Captions: A cleaned, hypernymed, image alt-text dataset for automatic image captioning. In *Proceedings of ACL*, pp. 2556–2565, 2018.
- Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. In *Proceedings of ICLR*, 2020.
- Yixuan Su, Tian Lan, Huayang Li, Jialu Xu, Yan Wang, and Deng Cai. PandaGPT: One model to instruction-follow them all. *arXiv preprint arXiv:2305.16355*, 2023.
- Zhiqing Sun, Sheng Shen, Shengcao Cao, Haotian Liu, Chunyuan Li, Yikang Shen, Chuang Gan, Liang-Yan Gui, Yu-Xiong Wang, Yiming Yang, et al. Aligning large multimodal models with factually augmented rlhf. *arXiv preprint arXiv:2309.14525*, 2023.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. LLaMA: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*, 2023.

- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *NeurIPS*, 30, 2017.
- Junjie Wang, Yuxiang Zhang, Lin Zhang, Ping Yang, Xinyu Gao, Ziwei Wu, Xiaoqun Dong, Junqing He, Jianheng Zhuo, Qi Yang, Yongfeng Huang, Xiayu Li, Yanghan Wu, Junyu Lu, Xinyu Zhu, Weifeng Chen, Ting Han, Kunhao Pan, Rui Wang, Hao Wang, Xiaojun Wu, Zhongshen Zeng, Chongpei Chen, Ruyi Gan, and Jiaying Zhang. Fengshenbang 1.0: Being the foundation of chinese cognitive intelligence. *CoRR*, abs/2209.02970, 2022a.
- Peng Wang, An Yang, Rui Men, Junyang Lin, Shuai Bai, Zhikang Li, Jianxin Ma, Chang Zhou, Jingren Zhou, and Hongxia Yang. OFA: Unifying architectures, tasks, and modalities through a simple sequence-to-sequence learning framework. In Kamalika Chaudhuri, Stefanie Jegelka, Le Song, Csaba Szepesvari, Gang Niu, and Sivan Sabato (eds.), *Proceedings of ICML*, volume 162, pp. 23318–23340. PMLR, 17–23 Jul 2022b.
- Wenhui Wang, Hangbo Bao, Li Dong, Johan Bjorck, Zhiliang Peng, Qiang Liu, Kriti Aggarwal, Owais Khan Mohammed, Saksham Singhal, Subhojit Som, et al. Image as a foreign language: Beit pretraining for vision and vision-language tasks. In *Proceedings of CVPR*, pp. 19175–19186, 2023.
- Chunyu Xie, Heng Cai, Jianfei Song, Jincheng Li, Fanjing Kong, Xiaoyu Wu, Henrique Morimitsu, Lin Yao, Dexin Wang, Xiangzheng Zhang, Dawei Leng, Xiangyang Ji, and Yafeng Deng. Zero and R2D2: A large-scale chinese cross-modal benchmark and a vision-language framework, 2022.
- An Yang, Junshu Pan, Junyang Lin, Rui Men, Yichang Zhang, Jingren Zhou, and Chang Zhou. Chinese CLIP: Contrastive vision-language pretraining in chinese. *arXiv preprint arXiv:2211.01335*, 2022.
- Yukang Yang, Dongnan Gui, Yuhui Yuan, Haisong Ding, Han Hu, and Kai Chen. Glyphcontrol: Glyph conditional control for visual text generation. 2023.
- Qinghao Ye, Haiyang Xu, Guohai Xu, Jiabo Ye, Ming Yan, Yiyang Zhou, Junyang Wang, Anwen Hu, Pengcheng Shi, Yaya Shi, Chaoya Jiang, Chenliang Li, Yuanhong Xu, Hehong Chen, Junfeng Tian, Qian Qi, Ji Zhang, and Fei Huang. mPLUG-Owl: Modularization empowers large language models with multimodality, 2023.
- Jiahui Yu, Zirui Wang, Vijay Vasudevan, Legg Yeung, Mojtaba Seyedhosseini, and Yonghui Wu. CoCa: Contrastive captioners are image-text foundation models. *arXiv preprint arXiv:2205.01917*, 2022a.
- Jiahui Yu, Yuanzhong Xu, Jing Yu Koh, Thang Luong, Gunjan Baid, Zirui Wang, Vijay Vasudevan, Alexander Ku, Yinfei Yang, Burcu Karagol Ayan, Ben Hutchinson, Wei Han, Zarana Parekh, Xin Li, Han Zhang, Jason Baldridge, and Yonghui Wu. Scaling autoregressive models for content-rich text-to-image generation. *TMLR*, 2022b. ISSN 2835-8856. Featured Certification.
- Tianyu Yu, Jinyi Hu, Yuan Yao, Haoye Zhang, Yue Zhao, Chongyi Wang, Shan Wang, Yinxv Pan, Jiao Xue, Dahai Li, et al. Reformulating vision-language foundation models and datasets towards universal multimodal assistants. *arXiv preprint arXiv:2310.00653*, 2023a.
- Tianyu Yu, Yuan Yao, Haoye Zhang, Taiwan He, Yifeng Han, Ganqu Cui, Jinyi Hu, Zhiyuan Liu, Hai-Tao Zheng, Maosong Sun, et al. RLhf-v: Towards trustworthy mllms via behavior alignment from fine-grained correctional human feedback. *arXiv preprint arXiv:2312.00849*, 2023b.
- Lu Yuan, Dongdong Chen, Yi-Ling Chen, Noel Codella, Xiyang Dai, Jianfeng Gao, Houdong Hu, Xuedong Huang, Boxin Li, Chunyuan Li, et al. Florence: A new foundation model for computer vision. *arXiv preprint arXiv:2111.11432*, 2021.
- Yan Zeng, Wangchunshu Zhou, Ao Luo, Ziming Cheng, and Xinsong Zhang. Cross-view language modeling: Towards unified cross-lingual cross-modal pre-training. In Anna Rogers, Jordan Boyd-Graber, and Naoaki Okazaki (eds.), *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 5731–5746, Toronto, Canada, July 2023. Association for Computational Linguistics.

Ao Zhang, Hao Fei, Yuan Yao, Wei Ji, Li Li, Zhiyuan Liu, and Tat-Seng Chua. Transfer visual prompt generator across llms. *arXiv preprint arXiv:2305.01278*, 2023.

Zhengyan Zhang, Xu Han, Hao Zhou, Pei Ke, Yuxian Gu, Deming Ye, Yujia Qin, Yusheng Su, Haozhe Ji, Jian Guan, et al. CPM: A large-scale generative chinese pre-trained language model. *AI Open*, 2:93–99, 2021.

Mingyang Zhou, Luwei Zhou, Shuohang Wang, Yu Cheng, Linjie Li, Zhou Yu, and Jingjing Liu. UC2: Universal cross-lingual cross-modal vision-and-language pre-training. In *Proceedings of CVPR*, pp. 4155–4165, 2021.

Deyao Zhu, Jun Chen, Xiaoqian Shen, Xiang Li, and Mohamed Elhoseiny. Minigt-4: Enhancing vision-language understanding with advanced large language models. *arXiv preprint arXiv:2304.10592*, 2023.

Minfeng Zhu, Pingbo Pan, Wei Chen, and Yi Yang. DM-GAN: Dynamic memory generative adversarial networks for text-to-image synthesis. In *Proceedings of CVPR*, pp. 5802–5810, 2019.



## APPENDIX

### A CONTRIBUTIONS

The authors’ contributions can be outlined as follows:

In the preparation of the project, Jinyi Hu and Yuan Yao design the model architecture. Xu Han, Yankai Lin, Jiao Xue, Dahai Li, Zhiyuan Liu, and Maosong Sun offer invaluable guidance in refining the model’s architecture. Shan Wang, Chongyi Wang, and Jinyi Hu take charge of collecting and processing the extensive multimodal dataset required for pretraining. Additionally, Hanghao Wu, Yue Zhao, Haoye Zhang, and Yuan Yao collaborate on constructing the instruction tuning data. Jinyi Hu, Chongyi Wang, Tianyu Yu, Qianyu Chen, and Shan Wang jointly implement the training codebase.

In model training, Chongyi Wang, Tianyu Yu, and Yinxu Pan babysit the VISCPM-Chat training; Jinyi Hu, Shan Wang and Qianyu Chen take care of the VISCPM-Paint training.

In model evaluation, Jinyi Hu and Yuan Yao design the evaluation framework. Chongyi Wang and Tianyu Yu evaluate VISCPM-Chat; Jinyi Hu and Shan Wang execute the automatic evaluation of VISCPM-Paint and organize the human evaluation of VISCPM-Paint.

In paper writing, Jinyi Hu and Yuan Yao write the main paper; Yankai Lin, Zhiyuan Liu, and Maosong Sun provide suggestions to polish the writing.

For public usability, Jinyi Hu, Yinxu Pan, Chongyi Wang, Shan Wang, and Yuan Yao promote the open-source of VISCPM; Yinxu Pan develops the online demo and API of VISCPM; Chongyi Wang and Yinxu Pan implement the low-resource inference of VISCPM.

Throughout the project, Xu Han, Yankai Lin, Jiao Xue, Dahai Li, Zhiyuan Liu, and Maosong Sun provide invaluable technical guidance and advice.

### B DATASET

#### B.1 PRETRAINING DATASET

**COCO** (Lin et al., 2014): The COCO dataset is a meticulously compiled image caption dataset that encompasses everyday scenes with common objects. It includes a training set of 118,287 images, each complemented by five unique captions. For each pass through the dataset, we randomly select one caption, resulting in 591,435 image-text pairs, five times the total number of images.

**Visual Genome** (Krishna et al., 2017) Visual Genome stands as a meticulously labeled image dataset, enriched with detailed annotations for objects. Our employed training set features approximately 100K images, each accompanied by an average of 8 unique captions.

**CC3M** (Sharma et al., 2018) Also known as Conceptual Captions, the CC3M dataset contains approximately 3.3M image-text pairs. CC3M is collected from the web and carefully processed to achieve the cleanliness and informativeness of captions. Due to some broken image links, the final successfully downloaded dataset contains approximately 2.8M pairs.

**CC12M** (Changpinyo et al., 2021) The CC12M dataset is an extension of the CC3M (Sharma et al., 2018), hosting nearly 12M image-text pairs. A more relaxed collection pipeline was used for its creation. After downloading, our dataset comprises approximately 6M pairs.

**Laion2B** (Schuhmann et al., 2022) Laion-2B is a massive dataset populated with image data sourced from publicly accessible areas of the internet. Our successful download yielded a vast collection of around 1.3B images.

**Laion-COCO** (Christoph et al., 2022) Laion-COCO is a subset of the Laion-2B dataset which includes 600M image entries. These images have been captioned using the BLIP (Li et al., 2022) to generate high-quality descriptions mimicking the MS COCO style.

**Wukong** (Gu et al., 2022) Wukong is a 100M image-text pair in Chinese, where the images are filtered according to the image size, and the text is filtered according to its language, length, and frequency.

Table 6: Evaluation of bilingual multimodal chat models on UniMM-Bench in English.

Model	OKVQA	AOKVQA	GQA	VQAv2	AVG
mPLUG-Owl	<u>66.7</u>	62.5	63.0	66.0	64.6
VisualGLM	57.6	62.8	58.1	63.9	60.6
Ziya-Visual	66.1	68.7	67.4	67.3	67.4
Qwen-VL-Chat	<b>71.4</b>	<b>77.9</b>	<u>68.6</u>	<b>77.9</b>	<b>74.0</b>
ViSCPM-Chat	65.4	<u>75.5</u>	<b>71.5</b>	<u>76.6</u>	<b>72.3</b>

Table 7: Summary of multimodal chat models’ architecture, parameters, and training data. Here, the size of mPlug-Owl’s training data and the proportion of Chinese and English in the Ziya-Visual’s training data are reported in their model cards.

Model	Visual Module	LLM	Training Data
mPlug-Owl	ViT-L/14 (0.3B)	BLOOMZ-7B	-
VisualGLM	Q-Former (1.6B)	ChatGLM-6B	English: 300M; Chinese 30M
Ziya-Visual	Q-Former (1.1B)	Ziya-LLaMA-13B-v1	20M
Qwen-VL-Chat	ViT-bigG (1.9B)	Qwen-7B	English: 1.1B; Chinese: 300M
ViSCPM-Chat	Muffin (0.7B)	CPM-Bee-10B	English: 140M; Chinese: 1M

**Zero** (Xie et al., 2022) Sourced from a search engine, the Zero dataset comprises 20M images and corresponding textual descriptions, selected from a pool of 5B image-text pairs based on user click-through rate.

## B.2 INSTRUCTION TUNING DATASET

**LLaVA-Instruct-150K** (Liu et al., 2023a) LLaVA-Instruct-150K is a set of multimodal instruction-following data generated by GPT4. In its creation, image captions and corresponding bounding boxes are harnessed to encode images into textual sequences for the text-only GPT4 model. LLaVA-Instruct-150K divides the dataset into three types: Conversation, Detailed Description, and Complex Reasoning. For each type, manually designed examples are incorporated into prompts for in-context learning of GPT4.

**M<sup>3</sup>IT** (Li et al., 2023b) The M<sup>3</sup>IT dataset is a collection of large-scale multimodal instruction tuning datasets curated by leveraging downstream datasets covering diverse vision-language tasks, including captioning, reasoning, and visual question-answering. The datasets have been reformulated into a unified image-text schema. Certain constructed instances, derived from key datasets, are translated into a total of 100 languages. It is important to note that we do not proceed with additional translations for non-Chinese portions of the dataset. It’s worth noting that we do not proceed with additional translations for non-Chinese portions of the dataset.

**UniMM-Chat** (Yu et al., 2023a) UniMM-Chat is a bespoke knowledge-intensive multimodal conversational dataset. UniMM-Chat utilizes images from the MSCOCO (Lin et al., 2014) and associated labeled datasets as the foundation for data generation. The related datasets include VQAv2 (Antol et al., 2015), Visual Dialogue (Das et al., 2017), OKVQA (Marino et al., 2019), and AOKVQA (Schwenk et al., 2022). The initial step involves merging the annotation information across different datasets for the same image, including pairs of questions and answers, dialogues, rationale information, and image captions. The varied text annotations provide a diversified view of the image content, enabling the comprehensive interpretation of the image by ChatGPT.

Table 8: Summary of text-to-image models’ architecture and training data.

Model	Visual Module	Text Encoder	Training Data
AltDiffusion	UNet	AltCLIP	-
TaiyiDiffusion	UNet	Taiyi-CLIP-RoBERTa	English: 0; Chinese: 20M
ViSCPM-Paint	UNet	CPM-Bee	English: 1.3B; Chinese: 0

Table 9: Description and examples of different categories in Chinese Drawbench

Category	Description	Example
Relation	Ability to generate objects with specific interaction relationship between them	“挂在衣架上的一个帽子” (A hat hanging on a hanger.)
Compositional	Ability to multiply kinds of objects	“一个盘子上放有黄色的梨和紫色的苹果” (There are yellow pears and purple apples on a plate.)
Attribute	Ability to accurately generate objects with given attributes, such as size, color, and action.	“一件蓝色的天鹅绒晚礼服” (A blue velvet evening gown.)
Counterintuitive	Ability to generate unusual scene against common sense	“穿着宇航服的李白开宇宙飞船” (Li Bai flies a spaceship in a space suit.)
Rare Word	Ability to understand rare expression	窗台上的一盆剑兰 (A gladiolus on the windowsill.)
Count	Ability to generate given numbers of objects	三块涂有黄油的面包 (Three pieces of bread with butter.)
Long Input	Ability to understand long and complex input	平静的湖面上，船夫划着船桨在湖面上划过，泛起涟漪，一只鸟飞过水面，抓起来一只鱼 (On the calm lake, the boatman rowed his paddle across the lake, rippling, a bird flew over the water and caught a fish.)
Chinese Culture	Ability to generate scene related to Chinese Culture	海上生明月，天涯共此时 (As the bright moon shines over the sea, From far away you share this moment with me.)

Table 10: Details of image-text datasets used in ViSCPM-Chat’s pretraining, where “bil” means the mixed version of English and translated Chinese.

Datasets	Size	Weight	Epoch
ViSCPM-Chat			
COCO + Visual Genome	626K	12.50%	27.60
CC3M + CC12M	8.4M	25.00%	16.40
Laion-COCO	390M	62.50%	0.35
ViSCPM-Chat+			
COCO + VG	626K	12.50%	15.33
CC12M	5.6M	25.00%	3.43
Zero + Wukong	20M	12.50%	0.48
CC3M-bil	5.6M	12.50%	1.72
Laion-COCO-bil	780M	37.50%	0.04

### B.3 CASE STUDY

Fig. 7 and Fig. 8 vividly depict ViSCPM-Chat’s capabilities, embodying wide-ranging global knowledge assimilation. As illustrated in the two case studies in Fig. 7, ViSCPM-Chat can identify the Mona Lisa painting adapted in a surreal style and recognize a stained map of New York City, further interpreting its real-world meaning. Demonstrating a well-balanced capacity for Chinese-English multi-modal conversation and robust text recognition skills, as exemplified in Fig. 8, ViSCPM-Chat can engage in fluent multimodal conversations on varying topics in English and effectively identify the words “Starbucks”, “Avengers: Endgame”, and its release date, “April 26th”, from visual inputs.

## C DETAILS OF BASELINES

The details of existing Chinese-English bilingual multimodal chat models are summarized in Table 7 and introduced as follows:

- mPLUG-Owl (Ye et al., 2023): mPLUG-Owl consists of a vision encoder, a vision abstractor module, and BLOOMZ (Muennighoff et al., 2023) as language model backbone. It is trained on LAION-400M, COYO, CC3M, and MSCOCO.

Table 11: Performance of MPM on IGLUE benchmark. We compare MPM with mUNITER, xUNITER (Liu et al., 2021), UC<sup>2</sup> (Zhou et al., 2021) and M<sup>3</sup>P (Ni et al., 2021).

Model	xNLI				xGQA						
	ARB	SPA	FRA	RUS	BEN	DEU	IND	KOR	POR	RUS	CMN
mUNITER	46.7	57.0	59.4	51.7	3.1	24.0	9.4	4.2	13.7	8.5	7.0
xUNITER	52.0	58.9	63.3	59.7	10.8	34.8	33.7	12.1	22.1	18.8	19.6
UC <sup>2</sup>	<b>56.2</b>	57.5	69.5	64.9	<b>20.0</b>	42.9	28.7	21.4	30.4	31.0	31.2
M <sup>3</sup> P	55.2	58.9	56.4	62.5	18.6	33.4	32.5	25.1	31.4	27.5	28.7
MPM	55.8	<b>75.9</b>	<b>78.1</b>	<b>75.0</b>	6.6	<b>51.0</b>	<b>37.1</b>	<b>36.9</b>	<b>49.8</b>	<b>47.9</b>	<b>46.1</b>

- VisualGLM (Du et al., 2022): VisualGLM employs Q-Former as image encoder and ChatGLM-6B (Du et al., 2022) as language model backbone. VisualGLM-6B’s pretraining incorporates 30M high-quality Chinese image-text pairs and 300M filtered English image-text pairs. In the fine-tuning phase, VisualGLM is trained on long VQA data.
- Ziya-Visual (Wang et al., 2022a): Ziya-Visual leverage Q-Former (Li et al., 2023a) as image encoder and Ziya-LLaMA-13B-v1 as language model backbone. They utilize 20M Chinese image-text pair, which is built by cleaning high-quality data from open source data, translating English datasets, and extracting coarse-grained information from captions using BLIP (Li et al., 2022) and Grounded SAM (Kirillov et al., 2023; Liu et al., 2023c).
- Qwen-VL-Chat (Bai et al., 2023): Qwen-VL-Chat utilizes a large ViT (Dosovitskiy et al., 2020) with 1.9B parameters initialized from Openclip’s bigG (Ilharco et al., 2021) as image encoder and Qwen-7B as language model backbone. The image encoder and language model are connected by a one-layer cross-attention module augmented with positional encodings. Qwen-VL-Chat p

The details of existing Chinese-English bilingual text-to-image models are summarized in Table 8 and introduced as follows:

- AltDiffusion (Chen et al., 2023): AltDiffusion is a Chinese-English bilingual text-to-image model based on Stable Diffusion and bilingual vision-language encoder AltClip (Chen et al., 2023). The training data are collected from Laion (Schuhmann et al., 2021).
- TaiyiDiffusion (Wang et al., 2022a): TaiyiDiffusion is a Chinese text-to-image model that adapts a Chinese text encoder into Stable Diffusion. The visual part is frozen during training. The training datasets include 20M filtered Chinese image-text pairs.

## D DETAILS OF VISCPM-CHAT

### D.1 MULTIMODAL PRETRAINING

In the visual feature alignment pretraining, we train the visual encoder on a mix of image-text pair datasets, including CC3M (Sharma et al., 2018), CC12M (Changpinyo et al., 2021), COCO (Lin et al., 2014), Visual Genome (Krishna et al., 2017), and Laion-COCO (Christoph et al., 2022). We train VISCPM-Chat and VISCPM-Chat+ for 180K and 480K steps, respectively, with a batch size of 768 and a learning rate of 1e-5. The distribution of each dataset used in pretraining for VISCPM-Chat and VISCPM-Chat+ are shown in Table 10. Only the parameters of the visual module are optimized in this stage.

### D.2 INSTRUCTION TUNING

During the instruction tuning phase, the optimization process involves the visual module and the LLM. We train VISCPM-Chat on the mixed dataset from LLaVA-Instruct-150K, UniMM-Chat, and M<sup>3</sup>IT, where LLaVA-Instruct-150K and UniMM-Chat are Chinese-English bilingual versions. The instruction tuning lasts for 80K steps with a batch size of 64. The training configuration of VISCPM-Chat+ is identical to VISCPM-Chat in this stage. In the inference phase, VISCPM-Chat uses beam search decoding with a beam size of 3.




Figure 5: The trend of response language accuracy and content quality when using different percentages of translated Chinese data during instruction tuning.

### D.3 EVALUATION BENCHMARK

**LLaVA Test Set** (Liu et al., 2023a) LLaVA Test Set consists of 90 instances, each containing an image, a question, and an answer, which comprehensively evaluates the model’s performance in conversation, detailed description, and complex reasoning. Following LLaVA (Liu et al., 2023a), we use GPT-4<sup>3</sup> to rate the model-generated answers and reference answers in a range of 1-10. We norm the average score of each instance to 1-100.

**UniMM-Bench** (Yu et al., 2023a) UniMM-Bench consists of 400 test instances, uniformly sampled from 4 commonly used VQA benchmarks, including OKVQA (Marino et al., 2019), AOKVQA (Schwenk et al., 2022), GQA (Hudson & Manning, 2019), and VQAv2 (Antol et al., 2015), whose annotations has undergone meticulous check. In the traditional evaluation of visual question answering, the metrics are to compute whether the model-generated answers are exactly matched to the reference answers. However, considering that LLMs typically generate responses with one complete sentence to answer the question, UniMM-Bench is built to comprehensively evaluate multimodal chat models’ capabilities in the context of visual question-answering, including reasoning, commonsense, grounding, and world knowledge (Schwenk et al., 2022). Given the cost of using GPT-4 for evaluation, 400 test instances are sampled from reliable VQA benchmarks.

### D.4 THE TREND OF USING DIFFERENT NUMBER OF CHINESE TRANSLATED DATA

To investigate the influence of varying amounts of Chinese data on the model’s ability to respond in the target language during the instruction tuning, surpassing the quasi-zero-shot transfer threshold, we systematically varied the percentages of translated Chinese data. Our analysis, depicted in Fig 5, elucidates the trends in both response language accuracy and content quality as a function of the Chinese data percentage utilized during instruction tuning. Remarkably, introducing a mere 0.1% of translated Chinese data suffices to achieve over 80% accuracy in response language accuracy, signifying the model’s high sensitivity to even minimal language-specific data inputs. Elevating this percentage to 1% results in a scenario where all responses are rendered in Chinese. Regarding content quality, our findings suggest that a range between 25% to 50% of translated Chinese data represents an optimal balance, yielding high-quality content responses. We leave the exploration of determining a more precise optimal proportion of translated Chinese data for maximizing both response content quality in future work.

### D.5 ADDITIONAL RESULTS ON IGLUE

We evaluate the effectiveness of MPM in multilingual image understanding on two benchmarks of IGLUE: xVNL (Bugliarello et al., 2022) and xGQA (Pfeiffer et al., 2022). We choose our model

<sup>3</sup>We use GPT-4 0314 version to evaluate the answers.




Figure 6: The curves of FID  $\downarrow$  vs CLIP score  $\uparrow$  with different classifier guidance scales.

pre-trained exclusively on English image-text pairs with LLaMA as the backbone language model, fine-tune it with IGLUE’s English training set, and conduct a zero-shot evaluation following the official setting of IGLUE.

As Table 11 shows, the performance of MPM demonstrates a notable improvement over baseline models. Specifically, when using Chinese-English bilingual LLM CPM-Bee as the backbone language model, MPM attains a 56.05 score on xGQA for CMN (Chinese Mandarin). Based on the above results, we can observe that for languages where multilingual LLMs demonstrate proficiency, such as Spanish, French, Russian, and German, the multilingual capability can be effectively generalized to achieve strong multimodal capabilities in the target language.

## E DETAILS OF VISCPM-PAINT

### E.1 TRAINING AND SAMPLING

VISCPM-Paint consists of a UNet, initialized from Stable Diffusion, and a bilingual LLM CPM-Bee (Zhang et al., 2021). We pretrain VISCPM-Paint on English image-text pairs dataset Laion-2B (Schuhmann et al., 2022) for 300k steps with batch size 4096. The 300K step consists of 200K-step training on  $256 \times 256$  resolutions and 100K training  $512 \times 512$  resolutions. We only optimize the cross-attention layer of UNet and the linear transblock between LLM and UNet. Linear transblock consists of a linear layer and a layer-norm layer, where the linear layer converts the dimension of LLM’s hidden states into UNet’s dimension. After pretraining on Laion-2B, we continue to train VISCPM-Paint+ on 20M filtered native Chinese image-text pairs and 136M translated image-text pairs in Laion-COCO (Christoph et al., 2022). We sample the generated image with DDPM scheduler (Ho et al., 2020) with 50 steps. For each prompt, we run three models to generate four images and select the one with the highest CLIP score evaluated by Chinese-CLIP (Yang et al., 2022).

### E.2 TRADE OFF OF FIDELITY AND ALIGNMENT

To visualize the trade-off between fidelity and alignment, Fig. 6 plots the curves of the FID against the CLIP score under different classifier guidance scales. The results indicate VISCPM-Paint and VISCPM-Paint+ deliver good overall performance regarding the balance between image quality and semantic alignment. While the CLIP score of AltDiffusion slightly outperforms that of VISCPM-Paint and VISCPM-Paint+, AltDiffusion’s FID falls significantly short. The sub-optimal quality of generated images will affect the practical use, as shown in the human evaluation below.

### E.3 HUMAN EVALUATION

To better understand the model performance on text-to-image generation in Chinese, we create a comprehensive set of prompts in Chinese, named *Chinese Drawbench*. *Chinese Drawbench* contains 174 prompts and consists of 8 categories, including Relation, Compositional, Attribute, Counterintuitive, Rare Word, Count, Long Input, and Chinese Culture. The descriptions and examples of each category are shown in Table 9. During the construction of *Chinese Drawbench*, we asked 7 annotators to create 25 prompts per person conditioned on the given categories. To diversify the prompts, we also restrain the field of 20 prompts ranging from 9 common classes, such as animal, plants, food, space, etc, and leave five prompts with no restrain. After this, we double-check the prompts and polish them minor.

We invite 5 independent human evaluators to assess the performance of VISCPM-Paint, TaiyiDiffusion (Wang et al., 2022a), and AltDiffusion on *Chinese Drawbench*. For each instance, three random shuffled images and the input prompt are shown to the human evaluators. We ask them to select the best one among three images from 3 aspects: Overall, Alignment, and Fidelity. Alignment measures the consistency between the generated image and the input prompt. Fidelity measures the clarity, aesthetics appeal, and object realism of the image. Overall measures the comprehensive quality of generated images. Figure 10 provides detailed results of each category of *Chinese DrawBench*.

## F DETAIL OF ABLATION STUDY

We introduce the detailed experimental configuration in the ablation study.

For image-to-text tasks shown in Table 5a, the native Chinese dataset consists of 20M image-text pairs filtered from Wukong (Gu et al., 2022) and Zero (Xie et al., 2022). For the “Pretrain” column in Table 5a, “Native Chinese” stands for the 20M native Chinese image-text pairs filtered from Wukong (Gu et al., 2022) and Zero (Xie et al., 2022). “English” corresponds to the dataset used in VISCPM-Chat, i.e., COCO, Visual Genome, CC3M, CC12M, and Laion-COCO. “Translated Chinese” consists of 390M Chinese image-text pairs whose captions are translated from English in Laion-COCO. In instruction tuning, i.e., supervised fine-tuning (SFT), we only utilize LLaVA and its translated Chinese version to reduce the computation time in the ablation study. Specifically, “Chinese” in the “SFT” column of Table 5a means translated Chinese LLaVA instruction tuning dataset. “Bilingual” means the mixed version of English and translated Chinese LLaVA.


For text-to-image tasks shown in Table 5b, “English” stands for Laion-2B. The other setting is the same as the image-to-text tasks introduced above.

## G DETAIL OF CLIP SCORE

We use Chinese-CLIP (Yang et al., 2022) to compute the CLIP score in Chinese image-text pairs. Chinese CLIP models have been used in several Chinese text-to-image models for data filtering and text encoding (Wang et al., 2022a; Liu et al., 2023b; Yang et al., 2023). We referred to Taiyidiffusion’s criteria for data volume and filtering standards (including CLIP score) in training Chinese text-to-image models. We manually check the caption in the Wukong dataset. We observe that the overall quality of text content in Wukong is unsatisfactory. Here are some findings about low CLIP scores:

- Broken Images: Some image URLs are broken or redirected to the same default image for downloading.
- Unrelated captions: Some content of captions are totally unrelated to image content.
- Incomplete sentences: Some captions are truncated or missing punctuation.





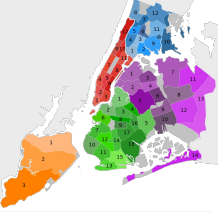
Do Do

你怎么看这幅画?  
(Please appreciate this painting)

这幅画描绘了一只猫的肖像，它有长长的胡须和大眼睛。这只猫直勾勾地盯着观众看，给人一种好奇或顽皮的感觉。在背景中可以看到一个绿色的墙，为画面增添了一抹色彩。整幅作品具有艺术性与趣味性，捕捉到了猫咪独特而富有表现力的特征。  
(This painting depicts a portrait of a cat with long whiskers and large eyes. The cat is staring straight at the viewer, giving it a sense of curiosity or playfulness. A green wall can be seen in the background, adding a splash of color to the picture. The whole composition is artistic and playful, capturing the unique and expressive characteristics of cats.)

这只猫的形象和哪副世界名画中的人物形象有关? 为什么?  
(The image of this cat is related to the characters in which world famous painting? Why?)

这只猫的形象与达芬奇名画《蒙娜丽莎》中女人的形象相似，因为它被描绘成一个长着大眼睛、有胡须和头发的女人。这幅画以其神秘而富有表现力的微笑著称，而这只猫似乎也在做同样的表情——睁大了眼睛直勾勾地盯着观众。这种相似性可能是由于两者都具有独特的面部特征或表现出一种类似于达芬奇艺术作品中所捕捉到的情感和氛围。  
(The image of the cat is similar to the image of the woman in Leonardo da Vinci's famous painting "Mona Lisa", as it is depicted as a woman with big eyes, beard and hair. The painting is known for its enigmatic and expressive wechat, and the cat seems to be doing the same—staring wide-eyed and intently at the viewer. This similarity could be due to both having distinctive facial features or displaying an emotion and atmosphere similar to those captured in da Vinci's artwork.)



Do Do


地图中画的是哪个城市?  
(Which city is drawn on the map?)

这是地图显示的是纽约市。  
(This is a map of New York City.)

图中不同的色块是什么含义?  
(What do the different color blocks in the picture mean?)

图中不同的颜色块代表了纽约市各个区域，每个区域都有自己独特的颜色。这些色彩编码系统有助于人们快速、容易地识别和区分城市内不同的街区或社区。通过使用这种编码方式，居民和游客可以很容易地找到他们想要去的地方并更好地理解这座城市的布局与组织结构。  
(The different colored blocks in the map represent the various boroughs of New York City, each with its own unique color. These color-coding systems help people quickly and easily identify and distinguish between different neighborhoods or neighborhoods within a city. By using this coding method, residents and tourists can easily find where they want to go and better understand the layout and organization of the city.)

Figure 7: Multimodal conversation cases of VISCPM-Chat in Chinese. VISCPM-Chat can recognize artwork like the Mona Lisa, albeit adapted in a surrealistic style, and a detailed, stained map of New York City.




Which movie is this poster from and which characters appear in the poster?

This poster is for the movie "Avengers: Endgame", which is set to release on April 26th. The characters featured in the poster are Thor, Hulk, Iron Man, Captain America, Black Widow, and Hawkeye.

Who about the villain?

The villain in the movie is Thanos, a powerful and dangerous character from the Avengers Movie.



What is this place in the picture?

The place in the picture is a Starbucks coffee shop, with a large green and white sign on its side that reads "Starbucks".


What can I buy here?

In the image, there is a Starbucks coffee shop sign hanging on the side of a building. This suggests that one can buy beverages, such as coffee or tea, at this location. The presence of the Starbucks logo also indicates that they offer various products and services related coffee, including coffee-related merchandise like mugs, T-shirts, and other accessories.


Figure 8: Multimodal conversation cases of VISCPM-Chat in English. VISCPM-Chat can conduct fluent multimodal conversations on different topics in English and recognize the words “Starbucks”, “Avengers: Endgame” and its release date “April 26th” in the image.




Figure 9: Generated images of VISCPM-Paint.



(a) Overall



(b) Alignment



(c) Fidelity

Figure 10: Human preference rate on Chinese Drawbench in different kinds of aspects.