

TEXTBIND: Multi-turn Interleaved Multimodal Instruction-following in the Wild

Huayang Li^{♥♣*} Siheng Li^{♥♣*} Deng Cai^{♥*} Longyue Wang[♥]

Lemao Liu[♥] Taro Watanabe[♠] Yujiu Yang[♣] Shuming Shi[♥]

[♥]Tencent AI Lab [♠]Nara Institute of Science and Technology [♣]Tsinghua University

Abstract

Large language models with instruction-following abilities have revolutionized the field of artificial intelligence. These models show exceptional generalizability to tackle various real-world tasks through their natural language interfaces. However, their performance heavily relies on high-quality exemplar data, which is often difficult to obtain. This challenge is further exacerbated when it comes to multimodal instruction following. We introduce TEXTBIND, an almost annotation-free framework for empowering LLMs with multi-turn interleaved multimodal instruction-following capabilities. Our approach requires only image-caption pairs and generates multi-turn multimodal instruction-response conversations from a language model. To accommodate interleaved image-text inputs and outputs, we devise MIM, a language model-centric architecture that seamlessly integrates image encoder and decoder models. Extensive quantitative and qualitative experiments demonstrate that MIM trained on TEXTBIND achieves remarkable generation capability in multimodal conversations compared to recent baselines.

1 Introduction

Artificial intelligence (AI) has experienced a significant paradigm shift with the rise of large language models (LLMs). These models are capable of processing a wide range of natural language processing (NLP) applications through natural language interactions with users (OpenAI, 2022, 2023).

Recently, a number of efforts have been made to augment LLMs with visual perception and understanding abilities. Prior work uses template-based instruction-following datasets for training (Xu et al., 2023b; Dai et al., 2023; Li et al., 2023c).

These datasets comprise a variety of classic computer vision (CV) tasks, e.g., object detection, with each task being converted into an instructional format using a handful of human-written natural language instructions. However, classic CV tasks often represent manageable and focused abstractions or simplifications of real-world tasks (Marr, 2010), they generally fall short in representing the true variety and complexity of real-world tasks and capturing the lexical diversity of human language. For example, most of them are single-turn inquiries about a single input image, albeit a small fraction supports multi-turn textual interactions or multiple image inputs. Consequently, the instruction-following capabilities of models trained on these datasets remain limited in open-world scenarios (Xu et al., 2023a). This is reminiscent of the early development of instruction tuning in NLP, where public NLP tasks were initially adopted but eventually superseded by high-quality, diverse open-world instruction data (Ouyang et al., 2022). Nevertheless, collecting such data for multimodal models can be extremely costly.

In this paper, we address the above challenge by introducing TEXTBIND, an almost annotation-free framework for augmenting LLMs with multi-turn interleaved multimodal instruction-following capabilities. The main idea is to represent images through their textual descriptions, e.g., captions, and utilize an LLM to generate multi-turn instructions and responses. To ensure the coherence and meaningfulness of the constructed multi-turn conversations, we propose a series of strategies such as topic-aware image sampling and human-in-the-loop refinement of in-context demonstrations. TEXTBIND can harvest large-scale datasets given the abundance of public image-caption pairs. TEXTBIND provides examples of processing and generating arbitrarily interleaved image-text content. To accommodate interleaved image-text inputs and outputs, we devise MIM, a multimodal

*Equal Contribution. Work done during HL and SL's internships at Tencent AI Lab. Correspondence to DC (thisisjcykcd@gmail.com).

model that emphasizes the reasoning abilities of LLMs and seamlessly integrates image encoder and decoder models.

For a comprehensive evaluation of open-world multimodal instruction following, we also construct TEXTBINDEVAL, an evaluation set of TEXTBIND validated and refined by human annotators. To verify the effectiveness of TEXTBIND, we compare MIM models trained on TEXTBIND and various existing datasets such as MiniGPT-4 (Zhu et al., 2023) and LLaVA (Liu et al., 2023b). The results demonstrate that TEXTBIND substantially outperforms these datasets across multiple evaluation metrics. Additional evaluations on conventional vision-language, including MME (Fu et al., 2023), MMBench (Liu et al., 2023c), and MM-Vet (Yu et al., 2023), further confirm that TEXTBIND provides complementary benefits to existing multimodal datasets. For image generation quality, we compare the proposed MIM with GILL (Koh et al., 2023a) and SDXL (Podell et al., 2023), showcasing improved performance (e.g., 0.673 vs. 0.608 CLIP scores). The advantage is particularly pronounced in multi-turn scenarios. Moreover, a holistic human evaluation further substantiates the superior performance of MIM trained on TEXTBIND in contrast to baseline models in the context of open-world multi-turn interleaved multimodal instruction-following. Most interestingly, rather than requiring users to provide explicit commands and direct descriptions of the desired image, our model can spontaneously generate images in proper conversation contexts. Our qualitative experiments show that MIM trained on TEXTBIND can perform a wide range of tasks, including composing engaging stories inspired by a set of images, identifying commonalities and differences across multiple images, explaining concepts with vivid images, generating long coherent stories with illustrations, etc. Examples can be found in Appx. A. We hope TEXTBIND serves as an initial step towards building AGI that can interact with humans flexibly in different modalities and broad real-world scenarios.

2 Related Work

Multimodal Datasets Existing multimodal datasets can be broadly classified into two categories: (1) Conventional datasets for specific vision-language tasks such as image captioning (Chen et al., 2015; Agrawal et al., 2019; Young et al., 2014) and visually-grounded question answering (Hudson and Manning, 2019; Marino

et al., 2019; Singh et al., 2019; Lu et al., 2022; Zhou et al., 2018; Goyal et al., 2017; Gurari et al., 2018). (2) Recent datasets for general instruction following. For instance, MultiInstruct (Xu et al., 2023b), InstructBLIP (Dai et al., 2023), and M3IT (Li et al., 2023c) convert existing vision-language datasets into a unified instructional format with handcrafted templates. This approach is reminiscent of the early explorations on instruction tuning in NLP (Wei et al., 2022; Sanh et al., 2022), where existing NLP tasks were phrased as instructions. However, it has been reported that such instruction-tuned multimodal models still generalize poorly to open-world scenarios (Xu et al., 2023a). This finding also aligns with the observations in NLP (Ouyang et al., 2022), where template-based instruction tuning is less effective than instruction tuning data collected from real-world scenarios due to its restricted diversity. There are also some attempts to convert the output of existing vision-language models into natural language answers for constructing instruction-tuning data (Liu et al., 2023b; Zhu et al., 2023; Chen et al., 2023a).

Compared to existing instruction-tuning data, the examples in TEXTBIND (1) generally exhibit greater task and lexicon diversity; (2) typically involve multiple images scattered throughout a multi-turn conversation; (3) support multimodal output (image generation).

Multimodal Models To augment existing LLMs with visual abilities, one straightforward approach is to employ off-the-shelf vision models as external tools. That is, the LLM calls expert vision models through their language interfaces for completing specific visual tasks when needed (Wu et al., 2023a; Shen et al., 2023; Chen et al., 2023b; Zou et al., 2022; Yang et al., 2023; Surís et al., 2023). However, these approaches may suffer from cross-modal information loss and lack of generality.

Recently, end-to-end multimodal language models have garnered significant interest. Flamingo (Alayrac et al., 2022) and OpenFlamingo (Alayrac et al., 2022) are among the pioneering work to extend LLMs to vision-language pretraining. Different from training from scratch, subsequent research efforts have focused on integrating pre-trained vision and language models. BLIP-2 (Li et al., 2023b) proposes Qformer to align the feature spaces of vision models and language models. To date, various network architectures and training

strategies have been proposed (Zhu et al., 2023; Liu et al., 2023b; Ye et al., 2023; Li et al., 2023a; Zhang et al., 2023; Du et al., 2022; Chen et al., 2023a; Dai et al., 2023; Liu et al., 2023a). However, these models are limited to the use of visual content as input. Our work is inspired by recent work on LLM-empowered image retrieval or generation (Koh et al., 2023b,a) and the pioneer work of (Sun et al., 2022) for chitchat in the context of single photo sharing. Contrary to prior work, we aim to present the first instruction-following model capable of processing and generating arbitrarily interleaved image-text inputs and outputs.

Numerous contemporary studies also exist in this field (Team, 2023; Yao et al., 2023; Dong et al., 2023; Zheng et al., 2023; Ge et al., 2023). A unique characteristic of our work lies in our emphasis on the aspect of data creation, whereas other studies primarily concentrate on architecture design or training algorithms. Pan et al. (2023) and Chen et al. (2023c) focus on image editing, while Wu et al. (2023b) and Moon et al. (2023) place emphasis on incorporating additional modalities, such as audio and video.

Evaluation Conventional vision datasets designed for specific tasks and scenarios may suffer from data contamination issues for evaluating LLMs. Recently, efforts have been made to provide systematic evaluations with a broader coverage of diverse visual abilities. MME (Fu et al., 2023) is an evaluation dataset containing visually-grounded Yes/No questions. OwlEval (Ye et al., 2023) is a benchmark comprising 82 questions based on 50 images and relies on human feedback evaluation. The test size is limited, and the results may suffer from subjective bias. In response to these challenges, MMbench (Liu et al., 2023c) and MM-Vet (Yu et al., 2023) are two recent benchmarks aiming to offer more comprehensive evaluations by incorporating the use of ChatGPT/GPT4 for answer verification. LVLM Arena (Xu et al., 2023a), an online evaluation framework that ranks different models using human judgment, is also introduced. However, the above benchmarks primarily focus on question answering based on a single image at the beginning of a conversation.

3 TEXTBIND

In this work, we seek to enhance the multi-turn instruction-following capabilities of a language model in the context of arbitrarily interleaved im-

ages and text. Constructing such datasets poses significant challenges: 1) it demands inventive thinking for devising high-quality visually-grounded instructions and their responses; 2) it requires specialized expertise to craft appropriate images. To tackle these issues, we introduce TEXTBIND, a method that predominantly resorts to existing *text-only language models* to produce the desired data.

3.1 Definition of Data

The goal of TEXTBIND is to construct a collection of multi-turn conversation such as $[\mathbf{x}_u^1, \mathbf{x}_a^1, \dots, \mathbf{x}_u^T, \mathbf{x}_a^T]$, where T is the number of turns, \mathbf{x}_u^i denotes the i -th instruction from the user, and \mathbf{x}_a^i represents the i -th response from the assistant. The conversation is also accompanied by an image set $\{\mathbf{m}_1, \dots, \mathbf{m}_n\}$, where n is the number of unique images in this conversation. Each instruction \mathbf{x}_u^i or response \mathbf{x}_a^i is a sequence of tokens in $\mathcal{V}_{\text{lang}} \cup \mathcal{V}_{\text{img}}$, where $\mathcal{V}_{\text{lang}}$ is the ordinary vocabulary of a language model and \mathcal{V}_{img} contains n distinct pointers to the images $\mathbf{m}_1, \dots, \mathbf{m}_n$ respectively. It is worth noting that every image can appear at any point within the conversation.

3.2 Automatic Data Generation

TEXTBIND consists of a three-step pipeline: 1) topic-aware image sampling for ensuring the coherence of each conversation and the diversity across conversations; 2) LLM-empowered multi-turn instruction-response generation to create natural and practical conversations; 3) post-processing and filtering to eliminate low-quality data. An overview of the pipeline is shown in Fig. 1.

Topic-Aware Image Sampling The initial step of TEXTBIND entails assembling groups of images that will serve as the foundation for generating multi-turn conversations. In order to facilitate coherent, meaningful, and practical conversations, the images within each group should exhibit meaningful interconnections. Furthermore, to guarantee a comprehensive representation of real-world scenarios, the topics of images across different conversations should demonstrate a wide range of diversity.

Following the above inspirations, we employ unsupervised clustering algorithms to group the images in our dataset into clusters and execute a two-step image sampling process for each conversation. Concretely, we use the image encoder of the CLIP model (Radford et al., 2021) to obtain vector representations of images. Then, we execute the k -means algorithm to classify all images

into K clusters (topics). Examples of such clusters are given in Fig. 1. For each conversation, we randomly sample a cluster from the available K clusters, then sample $n \in \{2, 3, 4\}$ images from the chosen cluster. We want to highlight that the clustered images are semantically relevant, rather than visually similar.

Generation of Multi-turn Conversations After selecting a list of images, we proceed to leverage a text-only LLM, such as GPT-4, to simulate a conversation between a user and an assistant based on the chosen images. The core idea is to let LLMs receive and process the textual descriptions of the images as if they see the actual images. Given the abundance of publicly available image-caption pairs, we propose representing an image with an XML-like string `<imgX> DESCRIPTION </imgX>`, where DESCRIPTION serves as a placeholder for the image caption, `<imgX>` and `</imgX>` mark the caption boundaries, and X denotes the image index in the input image list. After generating the conversation, we replace the XML-like strings in the conversation with the original images. Importantly, to ensure that a caption faithfully describes its corresponding image, we employ the CLIP model (Radford et al., 2021) to filter out image-caption pairs with matching scores below a high threshold.

The detailed prompt can be found in Appx. B, and examples of generated conversations before mapping the textual descriptions back to visual images are shown in Appx. C. In the prompt, we also provide in-context examples to improve the generation quality. We collect the in-context examples through a human-in-the-loop refinement process, which is elaborated in §3.3.

Post-processing and Low-quality Filtering To ensure data quality, we filter out conversations where there is a pair of input and output image descriptions with an edit distance higher than 0.1. We also exclude conversations containing image descriptions not present in the provided image list and conversations containing formatting errors such as co-reference errors and invalid image tags.

3.3 Human-in-the-loop Refinement

In-context learning has been demonstrated to be crucial for enhancing the generation quality of LLMs (Brown et al., 2020; Wang et al., 2023). Therefore, we also construct a seed set of high-quality in-context examples \mathcal{S} . The seed set \mathcal{S} begins as an empty set and is iteratively updated

Statistics	
# of conversations	25,629
Avg. # turns in conversations	3.36
Avg. # images	
in conversations	2.46
in instructions	0.94
in responses	1.52
Avg. # words	
in conversations	285.90
in instructions	78.66
in responses	207.24

Table 1: Statistics of the dataset by applying TEXTBIND to GPT-4.

with human feedback. In each iteration, we follow the steps detailed below:

1. We employ the latest \mathcal{S} and the template in Appx. B, and generate 100 new conversations using TEXTBIND (§3).
2. We manually analyze the generated conversations. Each conversation is assigned a quality label (“Excellent”, “Satisfactory”, or “Poor”). Besides, we label the visual abilities required for each conversation. The detailed annotation guideline for quality labels and visual abilities is outlined in Tab. 8 (Appx. E).
3. We add the generated conversations with “Excellent” or “Satisfactory” labels to \mathcal{S} .

To ensure diversity across different conversations, we randomly sample three in-context examples from the seed set for each generation. We further require that at least one in-context example is labeled “Excellent” and the three examples must cover all four visual abilities. After three iterations, we fix the seed set and employ it to generate the remaining data. The percentage of “Poor” data annotated by humans declines from 30% to 9%.

4 TEXTBIND Data from GPT4

We apply TEXTBIND to GPT4 and the Conceptual Captions (CC3M) dataset (Sharma et al., 2018; Changpinyo et al., 2021) as a case study. The details of the construction process can be found in Appx. F. In this section, we present comprehensive analyses of the constructed dataset.

Statistics As depicted in Tab. 1, our constructed dataset comprises 25,629 conversations. The average number of turns per conversation is 3.36 (each turn is defined as a pair of instruction and response).

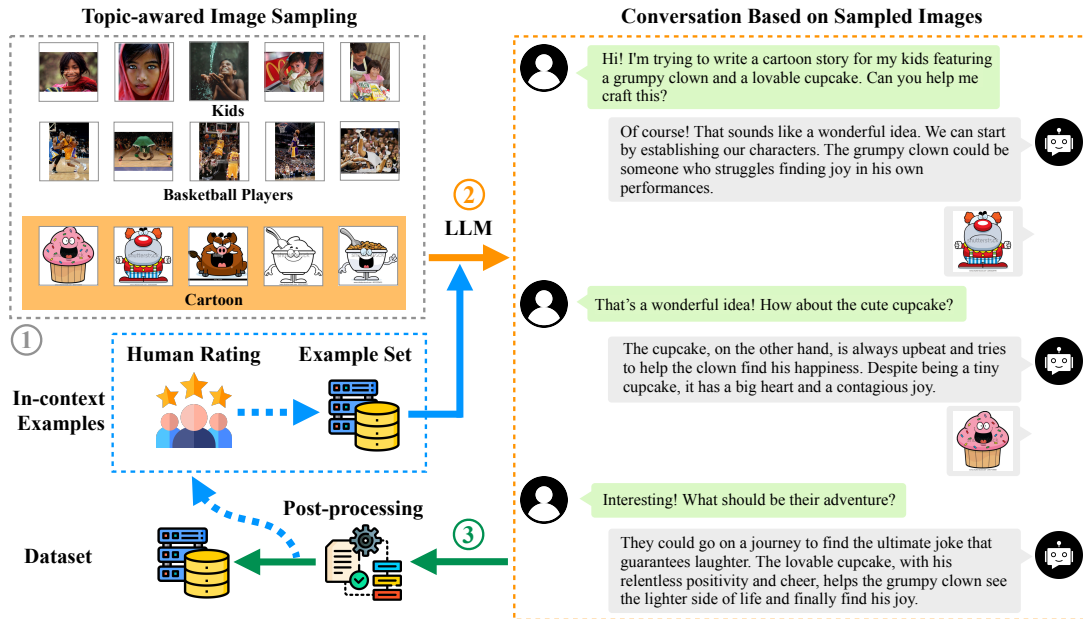


Figure 1: Illustration of the TEXTBIND method. In the top-left corner, we display five representative images from each of the three example clusters obtained via unsupervised clustering. On the right-hand side, a conversation is showcased and constructed using two randomly sampled images from the cartoon cluster.

The mean number of images in each conversation is 2.46. The length of each dialogue is about 285.90.

Diversity To understand the lexical and task diversity of our constructed data, we identify four types of required visual abilities: image creation, image comparison, intrinsic understanding, extrinsic understanding (See Tab. 8 for detailed definitions). The distribution is shown in Fig. 2b. We observe that a significant portion of conversations in our dataset focuses on more insightful and informative tasks, such as extrinsic understanding and image comparison. For topic diversity, we display three randomly sampled clusters in Fig. 1. The distribution of images across different turns is depicted in Fig. 2c. We also compare the lexical diversity of our dataset and existing datasets in Tab. 2. It can be seen that our datasets exhibit higher lexical diversity than existing datasets.

Quality To check the quality of the generated data, we randomly sample 100 conversations and perform an in-depth error analysis. The percentages of “Excellent”, “Good”, and “Poor” data are 57%, 34%, and 9% respectively. Note that we label the whole conversation as “Poor” if any of its turns has a problem. We identify three major sources: image-caption mismatch, incoherence, and hallucination. Detailed analysis is in Appx. G.

5 Augmenting LLMs with Visual I/O

5.1 Model

To support interleaved multimodal inputs and outputs, we supplement LLMs with visual input and output modules. Specifically, Llama2-Chat¹ (Touvron et al., 2023) is employed as the backbone LM. For visual input, we use the vision encoder from BLIP2 (Li et al., 2023b)², followed by a pretrained Q-Former model (Li et al., 2023b) that maps the features from the vision model into the embedding space of the LM. Inspired by GILL (Koh et al., 2023a), we attempt to learn a mapping from the output space of the LM to the input space of a stable diffusion (SD) model (Rombach et al., 2022) (in this work, the embeddings produced by the text encoder of Stable Diffusion XL (Podell et al., 2023)). To this end, we explore three model variants in our preliminary experiments.

Q-Former as Medium. We add a special token to the vocabulary of the LM, indicating that an image should be generated when it is emitted. We then use a Q-Former (Li et al., 2023b) that takes all previous hidden states of the LM as input and outputs the SD embeddings.

Q-Former with Prompt Tokens as Medium. To further leverage the reasoning abilities of the LM,

¹<https://huggingface.co/meta-llama/Llama-2-7b-chat-hf>

²<https://huggingface.co/Salesforce/blip2-flan-t5-xxl>

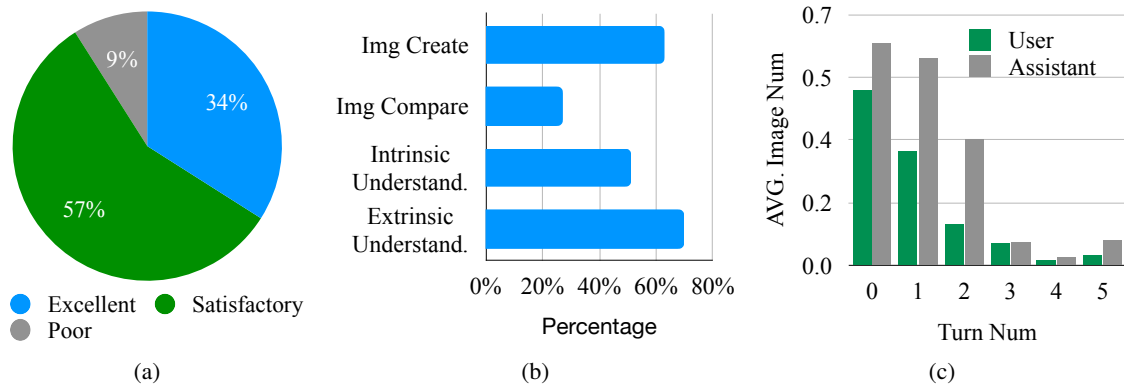


Figure 2: Statistics of data quality and diversity. The results in Fig. 2a and 2b are based on the human annotations on 100 randomly sampled conversations.

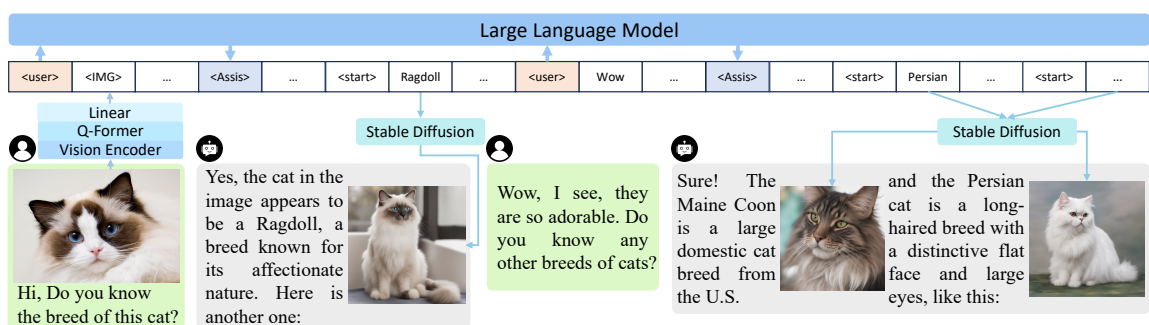


Figure 3: The architecture of MIM. It integrates a vision model, a language model, and a stable diffusion model. MIM is able to process multi-turn interleaved multimodal inputs and outputs. Notably, the images generated by stable diffusion model will be fed back to the MIM model to enhance the image understanding.

Dataset	Instruct	Response	Overall
LLAVA	1.56	1.84	1.70
MINIGPT-4	0.00	1.11	0.89
MULTIINSTRUCT	0.51	1.69	0.51
PLATYPUS	0.98	0.75	0.78
SHIKRA	0.89	1.08	0.87
TEXTBIND	1.76	1.92	1.84

Table 2: Averaged diversity scores of roles in various datasets. Details of this analysis are in Appx. D.

we incorporate a series of special tokens (`<img1>`, ..., `<IMG{r}>`), instead of a single token (``), to the LM. When `<img1>` is emitted, the generation of the special token sequence is enforced, serving as additional reasoning steps for predicting the forthcoming image. Subsequently, the Q-Former only accepts the hidden states of special tokens as input.

Language Description as Medium. The previous two variants try to align the continuous hidden spaces of different models. An alternative is to use discrete language descriptions for information exchange, as depicted in Fig. 3. Specifically, we add two special tokens, `<start>` and `<end>`, and

encode the generated text between these two tokens using the text encoder in the SD model.

The first two variants are optimized by minimizing the mean squared error (MSE) loss between the output and SD embeddings. For the third variant, we employ the standard cross-entropy loss. We empirically find that only the last method demonstrates satisfactory performance on **multi-turn interleaved multimodal instruction-following**, for which we name it MIM. The key difference between MIM and GILL is that MIM uses discrete text as medium rather than continuous embeddings. See §6.3 for detailed comparisons.

5.2 Training

Our training process consists of two stages, namely, the multimodal alignment stage and the multimodal instruction tuning stage.

Multimodal Alignment The first stage aims to align the feature spaces of the vision model and the language model. We utilize massive image-caption pairs for training, drawing from datasets such as Conceptual Captions (Changpinyo et al.,

2021; Sharma et al., 2018) and SBU (Ordonez et al., 2011). During training, only the Q-Former connecting the vision and language models is optimized while other model components remain frozen.

Multimodal Instruction Following The second stage further trains the joint model on multimodal instruction tuning data to improve its instruction-following capabilities. The Q-Former model and LLM are optimized in this stage. In addition to TEXTBIND, we also explore existing multimodal instruction data, including MultiInstruct (Xu et al., 2023b), MiniGPT-4 (Zhu et al., 2023), LLaVA (Liu et al., 2023b), and Shikra (Chen et al., 2023a).

6 Experiments

To verify the effectiveness of the proposed methods, we carry out quantitative evaluations against a set of recent baselines. Our quantitative evaluations are divided into three parts: textual response generation, image generation, and a holistic evaluation of multimodal instruction-following.

6.1 TEXTBINDEVAL

To facilitate comprehensive and dedicated evaluation for instruction-following in realistic scenarios, we construct a new dataset named TEXTBINDEVAL. TEXTBINDEVAL is initially generated through the automatic pipeline of TEXTBIND (§3) and subsequently refined by human annotators. These annotators are tasked with discarding low-quality examples or rectifying amendable issues such as revising incoherent or hallucinated content. After a rigorous review, we establish an evaluation dataset comprising 278 conversations in total.

6.2 Textual Response Generation

Setup We consider each assistant turn of each conversation in TEXTBINDEVAL as a test point. All its preceding context is treated as input (which may contain interleaved images and text), and the goal is to generate a coherent and helpful response. We measure the response quality using a set of reference-based evaluation metrics such as BLEU (Papineni et al., 2002), ROUGE (Lin, 2004), and BERTScore (Zhang et al., 2020). We also report the Diversity (Su et al., 2022) scores of the generated responses. For simplicity, we replace any image in the responses with a special token <image>.

For a fair comparison, we compare different MIM models trained on different datasets (Xu et al., 2023b; Zhu et al., 2023; Liu et al., 2023b;

Chen et al., 2023a)⁴ and GILL (Koh et al., 2023a)⁵. The implementation details are shown in Appx. H.

Results As shown in Tab. 3, the MIM model trained on TEXTBIND outperforms all other baselines by wide margins across all evaluation metrics. The results suggest that more realistic and diverse training data such as TEXTBIND is necessary for tackling open-world tasks, which cannot be well-supported by existing template-based and VQA-like datasets. Nevertheless, we also find that the performance can be further improved when combining different datasets (i.e., the mixture of MultiInstruct, MiniGPT-4, LLaVA, Shikra, and TEXTBIND, denoted by MIX), indicating that there is a complementary relationship between TEXTBIND and existing datasets.

Additional Results on Conventional Benchmarks

Finally, we report the results on three popular multimodal benchmarks, MME (Fu et al., 2023), MMBench (Liu et al., 2023c), and MM-Vet (Yu et al., 2023). As shown in Tab. 4, TEXTBIND gets a relatively lower score than other datasets. The reason stems from the intrinsic difference between TEXTBIND and the two benchmarks. TEXTBIND focuses more on realistic instructions (i.e., free-form text generation based on arbitrarily interleaved images and text across diverse user-assistant conversations). In contrast, MME, MMBench and MM-Vet focus more on VQA questions (e.g., who is this person, what is the color of the object), which are more similar to the data in MultiInstruct, LLaVA, and Shikra. For example, the model trained on MultiInstruct achieves the *best* performance on MME, though it displays the *worst* performance in open-world scenarios in Tab. 3. Another interesting observation is that the mix of all datasets (MIX) attains the best overall performance on MMBench and MM-Vet, indicating that different datasets are complementary. This again confirms that the capabilities that TEXTBIND can bring are almost orthogonal to existing multimodal instruction-following datasets.

6.3 Image Generation

Setup The models trained on existing datasets, i.e., the baselines in §6.2 except for GILL, are in-

⁴The original papers of these datasets used distinct model architectures such as different pretrained language models. Nevertheless, all of them do not support image generation.

⁵For a fair comparison, we replicate GILL using the same image-captioning data to train by our models.

Methods	BLEU-2	BLEU-4	ROUGE-2	ROUGE-L	BERTScore	Diversity
GILL (Koh et al., 2023a)	3.97	1.44	4.61	13.97	0.847	0.902
MultiInstruct (Xu et al., 2023b) ³	7.16	2.27	3.16	10.60	0.830	0.654
MiniGPT-4 (Zhu et al., 2023)	9.24	3.29	6.77	17.56	0.858	0.658
LLaVA (Liu et al., 2023b)	12.16	4.41	8.66	19.79	0.872	0.852
Shikra (Chen et al., 2023a)	10.37	3.83	7.79	18.63	0.864	0.722
TEXTBIND	24.45	11.83	15.45	28.69	0.891	0.927
MIX	27.64	14.49	17.90	31.22	0.896	<u>0.912</u>

Table 3: Evaluation results of textual response generation on TEXTBINDEVAL.

Training Dataset	MME		MMBench	MM-Vet
	Perception	Cognition	Overall	Overall
MultiInstruct (2023b)	1099.16	302.50	31.54	17.2
MiniGPT-4 (2023)	0.00	0.00	31.87	9.8
LLaVA (2023b)	683.28	<u>267.86</u>	<u>42.10</u>	<u>23.4</u>
Shikra (2023a)	166.87	2.86	41.10	19.9
TEXTBIND	549.00	226.43	22.64	19.4
MIX	<u>1023.33</u>	255.00	44.94	23.9

Table 4: Evaluation Results on MME (Fu et al., 2023), MMBench (Liu et al., 2023c), MM-Vet (Yu et al., 2023).

capable of generating images. To showcase the image generation capabilities of our model, we compare it with SDXL (Podell et al., 2023) and GILL (Koh et al., 2023a). In addition, we present the results of the two model variants described in §5.1, namely, **Q-former as Medium** and **Q-former with Prompt Tokens as Medium**.

We take each image from the assistant in TEXTBINDEVAL as a test point. All its preceding context is taken as input, and the models are enforced to output an image. We take the original images in TEXTBINDEVAL as references. Following Koh et al. (2023a), we evaluate image generation with two reference-based metrics: (1) **CLIP Similarity**. We use the CLIP vision encoder to produce image representations and compute the cosine similarity between generated images and reference images. A higher score means better semantic similarity. (2) **Learned Perceptual Image Path Similarity (LPIPS)**. LPIPS (Zhang et al., 2018) measures the distance between generated images and reference images. A lower score means that images are more similar in perceptual space. (3) **Frechet Inception Distance (FID)**. FID measures the distributional difference between the generated images and reference images. A lower score indicates better resemblance to reference images.

Results To gain further insights into the multi-turn instruction-following abilities, we group different test points by the number of previous conversation turns. The results are shown in Tab. 5. As seen,

MIM generally achieves better performance than SDXL and GILL across different turns and evaluation metrics. Importantly, the performance gaps are enlarged as the number of turns increases. This indicates that our model exhibits a better understanding ability of multi-turn conversations. Compared to the two model variants, MIM is substantially better. Our case study reveals that the disparity stems from the *one-to-many* nature of image generation in real-world conversations. Unlike generating images for explicit descriptions, there can exist numerous distinct images for a given conversation context. Operating in the hidden space may inadvertently average all possibilities, resulting in ambiguous or noisy images. However, MIM mitigates the *one-to-many* issue by taking full advantage of the autoregressive generation of language models for decision-making.

6.4 Holistic Evaluation

In addition to the above automatic evaluation, we also conduct a holistic evaluation of instruction-following abilities through human annotation.

Setup We randomly sample 100 contexts from TEXTBINDEVAL and evaluate the responses generated by MIM and two representative baselines, LLaVA (Liu et al., 2023b) and GILL (Koh et al., 2023a). We instruct three human annotators to score the quality of each generated response on a Likert scale from 1 to 4 (major error, minor error, acceptable, and excellent). The details of evaluation guideline are in Appx. I. For inter-annotator agreement, the Krippendorff’s $\alpha = 0.75$, which indicates a high consistency across annotators.

Results As depicted in Table 6, MIM achieves substantially higher human scores than GILL and LLaVA, in terms of both average score and acceptance rate, indicating its remarkable superiority in open-world multimodal conversations. To further elucidate where the derived dataset and training helps, we ask human annotators to evaluate the

Model	CLIP Similarity (\uparrow)			LPIPS (\downarrow)			FID (\downarrow)
	Turn-1	Turn-2	Turn-3	Turn-1	Turn-2	Turn-3	All
SDXL (Podell et al., 2023)	0.612	0.599	0.608	0.712	0.735	0.735	144.76
GILL (Koh et al., 2023a)	0.569	0.550	0.530	0.712	0.734	0.742	158.64
Q-Former as Medium	0.558	0.568	0.592	0.717	0.728	0.729	155.01
Q-Former with Prompt Tokens as Medium	0.566	0.571	0.606	0.718	0.727	0.732	152.23
MIM (Language Description as Medium)	0.640	0.645	0.673	0.712	0.720	0.726	139.46

Table 5: Evaluation results of image generation on TEXTBINDEVAL.

Methods	Overall	Acc. rate	IR	CU	RI
GILL	1.71	19%	-	-	-
LLaVA	2.93	70%	3.59	3.56	3.78
MIM	3.39	89%	3.99	3.82	3.72

Table 6: Human evaluation results. Acc. rate denotes the percentage of cases with an average score ≥ 3 . IR, CU, and RI represent intent recognition, context understanding and response informativeness, respectively.

quality of the generated responses across three fine-grained dimensions: intent recognition, context understanding and response informativeness. As seen, MIM is more adept at following human instructions and leveraging the multimodal context, while the informativeness scores of models trained on TEXTBIND and LLaVA remain comparable.

7 Conclusion

In conclusion, the introduction of the TEXTBIND framework has opened new doors for enhancing large language models with multi-turn interleaved multimodal instruction-following capabilities. By requiring only image-caption pairs, our approach significantly reduces the need for high-quality exemplar data, making it a more accessible and scalable solution for various real-world tasks. The MIM architecture seamlessly integrates image encoder and decoder models, enabling the model to effectively handle interleaved image-text inputs and outputs. Comprehensive quantitative and qualitative experiments demonstrate the remarkable performance of MIM, trained on TEXTBIND, when compared to recent baselines in open-world multimodal conversations.

Limitations

Despite the amazing instruction-following capabilities demonstrated by our models trained on TEXTBIND, there are many avenues for further improvement.

1. Our best-performing models use textual descriptions as a medium for image generation. Nevertheless, it is impossible to describe everything in an image through words. Thus, our current models have difficulties completing tasks such as precise image editing. Investigating fine-grained multimodal feature combinations may enhance performance in these areas.
2. The multimodal capabilities of TEXTBIND can be further enriched by incorporating additional modalities. For example, future work may use audio-text pairs for data synthesis.
3. Similar to existing multimodal language models, models trained on TEXTBIND exhibit some common deficiencies such as hallucination, toxicity, and stereotypes.

References

- Harsh Agrawal, Peter Anderson, Karan Desai, Yufei Wang, Xinlei Chen, Rishabh Jain, Mark Johnson, Dhruv Batra, Devi Parikh, and Stefan Lee. 2019. nocaps: novel object captioning at scale. In *2019 IEEE/CVF International Conference on Computer Vision, ICCV 2019, Seoul, Korea (South), October 27 - November 2, 2019*.
- Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katherine Millican, Malcolm Reynolds, et al. 2022. Flamingo: a visual language model for few-shot learning. *Advances in Neural Information Processing Systems*, 35.
- Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei.

2020. Language models are few-shot learners. In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.
- Soravit Changpinyo, Piyush Sharma, Nan Ding, and Radu Soricut. 2021. Conceptual 12m: Pushing web-scale image-text pre-training to recognize long-tail visual concepts. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2021, virtual, June 19-25, 2021*.
- Keqin Chen, Zhao Zhang, Weili Zeng, Richong Zhang, Feng Zhu, and Rui Zhao. 2023a. Shikra: Unleashing multimodal llm’s referential dialogue magic. *ArXiv preprint*, abs/2306.15195.
- Liangyu Chen, Bo Li, Sheng Shen, Jingkang Yang, Chunyuan Li, Kurt Keutzer, Trevor Darrell, and Ziwei Liu. 2023b. Language models are visual reasoning coordinators. In *ICLR 2023 Workshop on Mathematical and Empirical Understanding of Foundation Models*.
- Wei-Ge Chen, Irina Spiridonova, Jianwei Yang, Jianfeng Gao, and Chunyuan Li. 2023c. Llava-interactive: An all-in-one demo for image chat, segmentation, generation and editing. *arXiv preprint arXiv:2311.00571*.
- Xinlei Chen, Hao Fang, Tsung-Yi Lin, Ramakrishna Vedantam, Saurabh Gupta, Piotr Dollár, and C Lawrence Zitnick. 2015. Microsoft coco captions: Data collection and evaluation server. *ArXiv preprint*, abs/1504.00325.
- Wenliang Dai, Junnan Li, Dongxu Li, Anthony Meng Huat Tiong, Junqi Zhao, Weisheng Wang, Boyang Li, Pascale Fung, and Steven Hoi. 2023. Instructblip: Towards general-purpose vision-language models with instruction tuning. *ArXiv preprint*, abs/2305.06500.
- Runpei Dong, Chunrui Han, Yuang Peng, Zekun Qi, Zheng Ge, Jinrong Yang, Liang Zhao, Jianjian Sun, Hongyu Zhou, Haoran Wei, Xiangwen Kong, Xiangyu Zhang, Kaisheng Ma, and Li Yi. 2023. Dream-llm: Synergistic multimodal comprehension and creation. *ArXiv preprint*, abs/2309.11499.
- Zhengxiao Du, Yujie Qian, Xiao Liu, Ming Ding, Jiezhong Qiu, Zhilin Yang, and Jie Tang. 2022. GLM: General language model pretraining with autoregressive blank infilling. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*.
- Chaoyou Fu, Peixian Chen, Yunhang Shen, Yulei Qin, Mengdan Zhang, Xu Lin, Zhenyu Qiu, Wei Lin, Jinrui Yang, Xiawu Zheng, Ke Li, Xing Sun, and Rongrong Ji. 2023. Mme: A comprehensive evaluation benchmark for multimodal large language models. *ArXiv preprint*, abs/2306.13394.
- Yuying Ge, Sijie Zhao, Ziyun Zeng, Yixiao Ge, Chen Li, Xintao Wang, and Ying Shan. 2023. Making llama see and draw with seed tokenizer. *ArXiv preprint*, abs/2310.01218.
- Yash Goyal, Tejas Khot, Douglas Summers-Stay, Dhruv Batra, and Devi Parikh. 2017. Making the V in VQA matter: Elevating the role of image understanding in visual question answering. In *2017 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017, Honolulu, HI, USA, July 21-26, 2017*.
- Danna Gurari, Qing Li, Abigale J. Stangl, Anhong Guo, Chi Lin, Kristen Grauman, Jiebo Luo, and Jeffrey P. Bigham. 2018. Vizwiz grand challenge: Answering visual questions from blind people. In *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018*.
- Drew A. Hudson and Christopher D. Manning. 2019. GQA: A new dataset for real-world visual reasoning and compositional question answering. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16-20, 2019*.
- Jeff Johnson, Matthijs Douze, and Hervé Jégou. 2019. Billion-scale similarity search with GPUs. *IEEE Transactions on Big Data*, 7(3).
- Jing Yu Koh, Daniel Fried, and Ruslan Salakhutdinov. 2023a. Generating images with multimodal language models. *ArXiv preprint*, abs/2305.17216.
- Jing Yu Koh, Ruslan Salakhutdinov, and Daniel Fried. 2023b. Grounding language models to images for multimodal inputs and outputs. *ICML*.
- Bo Li, Yuanhan Zhang, Liangyu Chen, Jinghao Wang, Jingkang Yang, and Ziwei Liu. 2023a. Otter: A multi-modal model with in-context instruction tuning. *ArXiv preprint*, abs/2305.03726.
- Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. 2023b. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. *ArXiv preprint*, abs/2301.12597.
- Junnan Li, Dongxu Li, Caiming Xiong, and Steven C. H. Hoi. 2022. BLIP: bootstrapping language-image pre-training for unified vision-language understanding and generation. In *International Conference on Machine Learning, ICML 2022, 17-23 July 2022, Baltimore, Maryland, USA*, volume 162 of *Proceedings of Machine Learning Research*.
- Lei Li, Yuwei Yin, Shicheng Li, Liang Chen, Peiyi Wang, Shuhuai Ren, Mukai Li, Yazheng Yang, Jingjing Xu, Xu Sun, et al. 2023c. M3it: A large-scale dataset towards multi-modal multilingual instruction tuning. *ArXiv preprint*, abs/2306.04387.
- Chin-Yew Lin. 2004. ROUGE: A package for automatic evaluation of summaries. In *Text Summarization Branches Out*.

- Fuxiao Liu, Kevin Lin, Linjie Li, Jianfeng Wang, Yaser Yacoob, and Lijuan Wang. 2023a. Mitigating hallucination in large multi-modal models via robust instruction tuning. *arXiv preprint arXiv:2306.14565*, 1.
- Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. 2023b. Visual instruction tuning. *ArXiv preprint*, abs/2304.08485.
- Yuan Liu, Haodong Duan, Yuanhan Zhang, Bo Li, Songyang Zhang, Wangbo Zhao, Yike Yuan, Jiaqi Wang, Conghui He, Ziwei Liu, et al. 2023c. Mmbench: Is your multi-modal model an all-around player? *ArXiv preprint*, abs/2307.06281.
- Pan Lu, Swaroop Mishra, Tanglin Xia, Liang Qiu, Kai-Wei Chang, Song-Chun Zhu, Oyvind Tafjord, Peter Clark, and Ashwin Kalyan. 2022. Learn to explain: Multimodal reasoning via thought chains for science question answering. *Advances in Neural Information Processing Systems*, 35.
- Kenneth Marino, Mohammad Rastegari, Ali Farhadi, and Roozbeh Mottaghi. 2019. OK-VQA: A visual question answering benchmark requiring external knowledge. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16-20, 2019*.
- David Marr. 2010. *Vision: A computational investigation into the human representation and processing of visual information*.
- Seungwhan Moon, Andrea Madotto, Zhaojiang Lin, Tushar Nagarajan, Matt Smith, Shashank Jain, Chun-Fu Yeh, Prakash Murugesan, Peyman Heidari, Yue Liu, et al. 2023. Anymal: An efficient and scalable any-modality augmented language model. *ArXiv preprint*, abs/2309.16058.
- OpenAI. 2022. Introducing chatgpt.
- OpenAI. 2023. Gpt-4 technical report. *ArXiv preprint*, abs/2303.08774.
- Vicente Ordonez, Girish Kulkarni, and Tamara L. Berg. 2011. Im2text: Describing images using 1 million captioned photographs. In *Advances in Neural Information Processing Systems 24: 25th Annual Conference on Neural Information Processing Systems 2011. Proceedings of a meeting held 12-14 December 2011, Granada, Spain*.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. 2022. Training language models to follow instructions with human feedback. *Advances in Neural Information Processing Systems*, 35.
- Xichen Pan, Li Dong, Shaohan Huang, Zhiliang Peng, Wenhui Chen, and Furu Wei. 2023. Kosmos-g: Generating images in context with multimodal large language models. *ArXiv preprint*, abs/2310.02992.
- 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 of the Association for Computational Linguistics*.
- Dustin Podell, Zion English, Kyle Lacey, Andreas Blattmann, Tim Dockhorn, Jonas Müller, Joe Penna, and Robin Rombach. 2023. Sdxl: Improving latent diffusion models for high-resolution image synthesis. *ArXiv preprint*, abs/2307.01952.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. 2021. Learning transferable visual models from natural language supervision. In *Proceedings of the 38th International Conference on Machine Learning, ICML 2021, 18-24 July 2021, Virtual Event*, volume 139 of *Proceedings of Machine Learning Research*.
- Jeff Rasley, Samyam Rajbhandari, Olatunji Ruwase, and Yuxiong He. 2020. Deepspeed: System optimizations enable training deep learning models with over 100 billion parameters. In *KDD '20: The 26th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, Virtual Event, CA, USA, August 23-27, 2020*.
- Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. 2022. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*.
- Victor Sanh, Albert Webson, Colin Raffel, Stephen H. Bach, Lintang Sutawika, Zaid Alyafeai, Antoine Chaffin, Arnaud Stiegler, Arun Raja, Manan Dey, M Saiful Bari, Canwen Xu, Urmish Thakker, Shanya Sharma Sharma, Eliza Szczechla, Taewoon Kim, Gunjan Chhablani, Nihal V. Nayak, Debajyoti Datta, Jonathan Chang, Mike Tian-Jian Jiang, Han Wang, Matteo Manica, Sheng Shen, Zheng Xin Yong, Harshit Pandey, Rachel Bawden, Thomas Wang, Trishala Neeraj, Jos Rozen, Abheesht Sharma, Andrea Santilli, Thibault Févry, Jason Alan Fries, Ryan Teehan, Teven Le Scao, Stella Biderman, Leo Gao, Thomas Wolf, and Alexander M. Rush. 2022. Multi-task prompted training enables zero-shot task generalization. In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*.
- Piyush Sharma, Nan Ding, Sebastian Goodman, and Radu Soricut. 2018. Conceptual captions: A cleaned, hypernymed, image alt-text dataset for automatic image captioning. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*.
- Yongliang Shen, Kaitao Song, Xu Tan, Dongsheng Li, Weiming Lu, and Yueting Zhuang. 2023. Hugging-gpt: Solving ai tasks with chatgpt and its friends in huggingface. *ArXiv preprint*, abs/2303.17580.

- Amanpreet Singh, Vivek Natarajan, Meet Shah, Yu Jiang, Xinlei Chen, Dhruv Batra, Devi Parikh, and Marcus Rohrbach. 2019. Towards VQA models that can read. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16-20, 2019*.
- Yixuan Su, Tian Lan, Yan Wang, Dani Yogatama, Lingpeng Kong, and Nigel Collier. 2022. A contrastive framework for neural text generation. In *Advances in Neural Information Processing Systems*.
- Qingfeng Sun, Yujing Wang, Can Xu, Kai Zheng, Yaming Yang, Huang Hu, Fei Xu, Jessica Zhang, Xiubo Geng, and Daxin Jiang. 2022. Multimodal dialogue response generation. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*.
- Dídac Surís, Sachit Menon, and Carl Vondrick. 2023. Vipergpt: Visual inference via python execution for reasoning. *ArXiv preprint*, abs/2303.08128.
- InternLM Team. 2023. Internlm: A multilingual language model with progressively enhanced capabilities. <https://github.com/InternLM/InternLM-techreport>.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. 2023. Llama: Open and efficient foundation language models. *ArXiv preprint*, abs/2302.13971.
- Yizhong Wang, Yeganeh Kordi, Swaroop Mishra, Alisa Liu, Noah A. Smith, Daniel Khashabi, and Hannaneh Hajishirzi. 2023. Self-instruct: Aligning language models with self-generated instructions. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*.
- Jason Wei, Maarten Bosma, Vincent Y. Zhao, Kelvin Guu, Adams Wei Yu, Brian Lester, Nan Du, Andrew M. Dai, and Quoc V. Le. 2022. Finetuned language models are zero-shot learners. In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. 2020. Transformers: State-of-the-art natural language processing. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*.
- Chenfei Wu, Shengming Yin, Weizhen Qi, Xiaodong Wang, Zecheng Tang, and Nan Duan. 2023a. Visual chatgpt: Talking, drawing and editing with visual foundation models. *ArXiv preprint*, abs/2303.04671.
- Shengqiong Wu, Hao Fei, Leigang Qu, Wei Ji, and Tat-Seng Chua. 2023b. Next-gpt: Any-to-any multimodal llm.
- Peng Xu, Wenqi Shao, Kaipeng Zhang, Peng Gao, Shuo Liu, Meng Lei, Fanqing Meng, Siyuan Huang, Yu Qiao, and Ping Luo. 2023a. Lvlm-ehub: A comprehensive evaluation benchmark for large vision-language models. *ArXiv preprint*, abs/2306.09265.
- Zhiyang Xu, Ying Shen, and Lifu Huang. 2023b. Multi-Instruct: Improving multi-modal zero-shot learning via instruction tuning. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*.
- Zhengyuan Yang, Linjie Li, Jianfeng Wang, Kevin Lin, Ehsan Azarnasab, Faisal Ahmed, Zicheng Liu, Ce Liu, Michael Zeng, and Lijuan Wang. 2023. Mm-react: Prompting chatgpt for multimodal reasoning and action. *ArXiv preprint*, abs/2303.11381.
- Zhewei Yao, Xiaoxia Wu, Conglong Li, Minjia Zhang, Heyang Qin, Olatunji Ruwase, Ammar Ahmad Awan, Samyam Rajbhandari, and Yuxiong He. 2023. DeepSpeed-VisualChat: Multi-Round Multi-Image Interleave Chat via Multi-Modal Causal Attention. *ArXiv preprint*, abs/2309.14327.
- Qinghao Ye, Haiyang Xu, Guohai Xu, Jiabo Ye, Ming Yan, Yiyang Zhou, Junyang Wang, Anwen Hu, Pengcheng Shi, Yaya Shi, et al. 2023. mplug-owl: Modularization empowers large language models with multimodality. *ArXiv preprint*, abs/2304.14178.
- Peter Young, Alice Lai, Micah Hodosh, and Julia Hockenmaier. 2014. From image descriptions to visual denotations: New similarity metrics for semantic inference over event descriptions. *Transactions of the Association for Computational Linguistics*, 2.
- Weihao Yu, Zhengyuan Yang, Linjie Li, Jianfeng Wang, Kevin Lin, Zicheng Liu, Xinchao Wang, and Lijuan Wang. 2023. Mm-vet: Evaluating large multimodal models for integrated capabilities. *ArXiv preprint*, abs/2308.02490.
- Renrui Zhang, Jiaming Han, Aojun Zhou, Xiangfei Hu, Shilin Yan, Pan Lu, Hongsheng Li, Peng Gao, and Yu Qiao. 2023. Llama-adapter: Efficient fine-tuning of language models with zero-init attention. *ArXiv preprint*, abs/2303.16199.
- Richard Zhang, Phillip Isola, Alexei A. Efros, Eli Shechtman, and Oliver Wang. 2018. The unreasonable effectiveness of deep features as a perceptual metric. In *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018*.
- Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q. Weinberger, and Yoav Artzi. 2020. BERTscore: Evaluating text generation with BERT. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*.

Kaizhi Zheng, Xuehai He, and Xin Eric Wang. 2023. Minigpt-5: Interleaved vision-and-language generation via generative vokens.

Luowei Zhou, Chenliang Xu, and Jason J. Corso. 2018. Towards automatic learning of procedures from web instructional videos. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence, (AAAI-18), the 30th innovative Applications of Artificial Intelligence (IAAI-18), and the 8th AAAI Symposium on Educational Advances in Artificial Intelligence (EAAI-18), New Orleans, Louisiana, USA, February 2-7, 2018*.

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

Xueyan Zou, Zi-Yi Dou, Jianwei Yang, Zhe Gan, Linjie Li, Chunyuan Li, Xiyang Dai, Harkirat Behl, Jianfeng Wang, Lu Yuan, et al. 2022. Generalized decoding for pixel, image, and language. *ArXiv preprint*, abs/2212.11270.

A Demonstrations

The four high-level characteristics of the TEXTBIND dataset equips MIM with a variety of capabilities. We demonstrate those capabilities with concrete user cases.

Image Creation One core innovation of TEXTBIND is that it enables the model to create images based on the conversation context without explicit dictations from the users. This characteristic is extremely useful for open-world scenarios, because in many cases people may just have an implicit intention and have no clear thoughts about what the images should be. We observe that our model can explain concepts and ideas for users with vivid images (Figure 5a), creating images with correct emotions (Figure 5b), and editing images based on the whole context (Figure 5c and 5d). Furthermore, as shown in Figure 4, we discover that our model is proficient in generating long stories featuring interleaved text and images while maintaining exceptional coherence.

Image Comparison Another interesting feature of TEXTBIND is that it can compare or relate the information in multiple images. For example, our model can correctly explain the different and common parts in multiple images in Figure 6.

Intrinsic & Extrinsic Image Understanding The model trained on TEXTBIND can understand the content in images precisely in a multi-turn conversation. In all the three sub-figures of Figure 7, the model precisely follows the human instructions and explains the details of the images to users. Moreover, TEXTBIND also enables the model to explore the meaning of an image beyond the symbols in it. For example, the model also explains the the influence of Bob Dylan’s album in Figure 7b and the impact of iPhone in Figure 7c.

B Prompt of TEXTBIND

The prompt is shown in Fig. 8.

C Examples of Constructed Conversations

The examples are given in Fig. 9.

D Comparison with previous datasets.

We further investigate the lexical diversity of generated data, including both instructions and responses. The diversity of text in a dataset is defined as

$\sum_{n=2}^4 \left(\frac{\# \text{Unique } n\text{-grams}}{\# \text{Total } n\text{-grams}} \right)$, in line with previous works (Su et al., 2022). As shown in Tab. 2, our dataset achieves better diversity scores than most existing visual-language datasets, indicating that the language used in our dataset is more informative.

E Human annotation guideline.

The comprehensive guideline for human evaluation is shown in Table 8.

F Implementation Details (Data)


We construct our TEXTBIND dataset based on the CONCEPTUAL CAPTIONS 3M (CC3M) (Sharma et al., 2018; Changpinyo et al., 2021) dataset, which only provides image-caption pairs. In our experiments, we employ the clip-vit-base-patch16 model released by OpenAI⁶ (Radford et al., 2021) to filter out image-caption pairs with matching scores lower than 30. We use the k -means clustering algorithm implemented by FAISS (Johnson et al., 2019) toolkit to classify the cleaned CC3M dataset into 4096 clusters. The features used for k -means clustering are the hidden representations of images encoded by clip-vit-base-patch16 model. In addition, clusters with less than 32 images are regarded as outliers and will not be considered. The number of images in each conversation is sampled from $\{2, 3, 4\}$. We access the GPT-4 model through the OpenAI API⁷, and set top_p and temperature hyper-parameters to 1.0.

G Constructed Conversations with “Poor” Label

In Table 9, we identify three typical errors present in the constructed dataset. Despite setting a high threshold to filter out mismatched image-caption pairs, some mismatched cases cannot be detected by the CLIP model (Radford et al., 2021). A few conversations suffering from incoherence and hallucinations may be attributed to the GPT-4 model. Overall, while a small number of conversations are affected by errors that are difficult to detect using rules, most generated conversations exhibit high quality. We present several cases labeled with “Poor”. We can find that most of those “Poor” cases only have minor and non-obvious problems.


You
I really like marvel movies. Can you help me create a new superhero story?

Assistant
Sure! Let's start with a character who is a skilled martial artist and a master of disguise. They could be a secret agent who uses their skills to infiltrate criminal organizations. Here's an image to inspire you:




You
Great start! What's his super power?

Assistant
The character's superpower could be the ability to transform into a powerful, mythical creature. This would give them the strength and abilities to take on powerful enemies. Here's an image to represent this:




You
Will he fall in love with someone in the story?

Assistant
Absolutely! The character could fall in love with a woman who is also a skilled martial artist and a master of disguise. They could meet while working together to take down a criminal organization. This would create an interesting dynamic, as they both have the same skills and goals, but also have a personal connection. Here's an image to represent this:



You
Wow, cool! BTW, who is the leader of the criminal organizations in the story?

Assistant
The leader of the criminal organization could be a powerful and ruthless villain. They could be a former martial artist who turned to crime after a tragic event in their past. They could have a personal vendetta against the protagonist, making their conflict more intense and personal. Here's an image to represent this:



You
What's the ending of this story?

Assistant
The story could end with the protagonist and the woman they love successfully taking down the criminal organization, but at a great personal cost. The protagonist could sacrifice their life to save the woman they love, or they could be gravely injured in the final battle. The woman could be left to carry on the legacy of the protagonist, using their skills and knowledge to continue the fight against crime. This would create a sense of closure and a new beginning for the story.

Figure 4: Generation of a long story with interleaved text and images.

Dataset	Multi-turn	Arbitrarily Interleaved	Multimodal	Instruction Following	Image Generation
PLATYPUS	X	X	X	✓	X
FLAMINGO	X	✓	✓	X	✓
MULTIINSTRUCT	X	X	✓	✓	X
MINIGPT-4	X	X	✓	✓	X
LLAVA	✓	X	✓	✓	X
TEXTBIND	✓	✓	✓	✓	✓

Table 7: Comparison with previous datasets.

You
Can you explain the difference between soccer and football?

Assistant
Sure! Soccer and football are two different sports with distinct rules and gameplay. Soccer, also known as football in many parts of the world, is played with a round ball and involves two teams of 11 players each. The objective is to score goals by kicking or heading the ball into the opponent's goal. Here's an image of a soccer player kicking the ball:




. On the other hand, football, also known as American football, is played with an oval-shaped ball and involves two teams of 11 players each. The objective is to score points by carrying or throwing the ball into the opponent's end zone. Here's an image of a football player running with the ball:



(a) Explaining concepts with multiple images.

You
I am really sad today. Can you show me something happy?


Assistant
Of course, I'm sorry to hear that you're feeling sad. Here's an image that might cheer you up:



. It's a cute image of a little girl holding a red balloon, which can symbolize hope and joy.

You
Wow, so cute! I really love children and animals.

Assistant
I'm glad to hear that! Here's an image that combines both:




. This image shows a little girl and her dog enjoying a walk together in a beautiful field.

(b) Creating images with correct emotions.

You
I just opened a fried chicken restaurant, can you generate an ad for me? I hope to put it online to attract customers.


Assistant
Sure! Here's an ad for your fried chicken restaurant:



This ad showcases the delicious and crispy fried chicken that your restaurant offers.

You
We also sell burgers, can you include burgers in your ad?

Assistant
Of course! Here's an updated ad that includes burgers:




This ad highlights the variety of food options available at your restaurant, including fried chicken and burgers.

(c) Editing images based on context.


You
I am planning for my summer vacation, but I am not sure which city to visit. I have two initial candidates, New York and Seattle. Can you give me some suggestions?

Assistant
Sure! Both cities have unique attractions. New York City is known for its iconic skyline, bustling streets, and world-class museums. It's a great place to explore the arts, culture, and history. Here's an image of the New York City skyline at night:



You
What about Seattle?

Assistant
Seattle is a city with a rich history and culture. It's known for its coffee culture, music scene, and outdoor activities. The city is surrounded by mountains and water, making it a great place for hiking, kayaking, and exploring the outdoors. Here's an image of the city skyline at night:




(d) Creating images based on context.

Figure 5: User cases of creating images.

You

What's the difference between the two types of music?



and


Assistant

The two images depict different types of music performances. The first image shows a rock artist performing on stage, which is a common form of live music. Rock music is characterized by its strong emphasis on rhythm, powerful vocals, and often, electric instruments. The second image, on the other hand, shows a country artist performing on stage. Country music is a genre that originated in the southern United States and is characterized by its blend of folk, blues, and western music. It often features acoustic instruments, storytelling lyrics, and a strong emphasis on the singer's voice. In summary, the main difference between the two images is the genre of music being performed. Rock music is typically associated with electric instruments and a more energetic, rebellious vibe, while country music is often associated with acoustic instruments and a more laid-back, storytelling style.

(a) Comparing music styles.

You

What's the common part between the following two images:



and


Assistant

The common part between the two images is that they both depict a white horse. The first image shows a shark-like horse, which is a unique and imaginative representation of a horse. The second image, on the other hand, shows a more traditional representation of a horse, with a white coat and a calm demeanor. Both images highlight the versatility of horses as a subject in art, whether it's through realistic depictions or creative, imaginative representations.

(b) Relating images.

You

My little cousin really like super heroes. However, he is a little confused about the two groups super heroes:



and

Can you tell him the differences between the two groups of super heroes?

Assistant

Sure! The first image shows a group of superheroes from the Avengers, which is a team of superheroes from the Marvel Comics universe. The Avengers are known for their diverse range of powers and skills, and they often work together to protect the world from various threats. The second image shows a group of superheroes from the Justice League, which is a team of superheroes from the DC Comics universe. The Justice League is also a team of superheroes, but they are from different universes. They often work together to protect their respective universes and the world from threats. The main difference between the two groups is the universe they come from. The Avengers are from the Marvel Universe, while the Justice League is from the DC Universe. Additionally, the Avengers are known for their more light-hearted and humorous tone, while the Justice League is often more serious and dramatic.

You

Which universe do you think my cousin may like most as a young children.

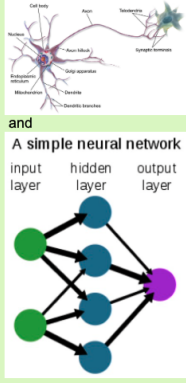
Assistant

As a young child, your cousin may be more drawn to the Avengers, as they are often portrayed as fun and exciting characters. The Avengers have a more light-hearted and humorous tone, which can be appealing to children. However, it's important to note that both universes have their own unique stories and characters, and it's ultimately up to your cousin's personal preference.

(c) Comparing movies.

You

I am new to deep learning. Can you explain the relationship between the two concepts:



and

Assistant

Sure! The first image shows a simple nervous system, which is a fundamental concept in biology. It consists of a network of interconnected cells that transmit information. The second image is a diagram of a simple artificial neural network, which is a fundamental concept in deep learning. It consists of interconnected nodes or "neurons" that process information. The key difference between the two is that the first image is a biological system, while the second is a man-made system. The biological system is more complex and has a more intricate structure, while the artificial system is simpler and more easily controlled. In deep learning, the artificial neural network is designed to mimic the structure and function of the biological nervous system. By doing so, it can learn to perform tasks that are difficult or impossible for traditional machine learning algorithms.

(d) Comparing different concepts.

Figure 6: User cases of comparing images.

You
What this person is doing?



Assistant

The image shows a person playing a mobile game on a smartphone. The game appears to be a multiplayer online battle arena (MOBA) game, which is a popular genre of video games that involves teams of players competing against each other in real-time. The player is likely using the touchscreen controls on the smartphone to navigate the game and interact with other players.

You
Do you know name of this game?

Assistant

Unfortunately, I don't have access to the specific game's name. However, based on the description of the game, it seems to be a popular genre of mobile games.

You
Given this picture, do you know the name of this game?




Assistant

Yes, the image shows a game called "League of Legends." It's a popular multiplayer online battle arena game that was first released in 2009. The game has a large player base and is known for its competitive gameplay and strategic depth.

(a)

You
It seems that this old picture has lots of stories:




Assistant

Indeed, this image has a rich history. It's a cover of a 1960s album by the artist Bob Dylan, who is considered one of the most influential figures in popular music. The album features some of his most iconic songs, such as "Blowin' in the Wind" and "Like a Rolling Stone".

You
Amazing! Do you have pictures about his live show?


Assistant

Yes, I do. Here's an image of a live performance by Bob Dylan during his 1966 tour.



(b)

You
Say his name!



Assistant

Sure! The man in the image is Steve Jobs, a renowned entrepreneur, inventor, and designer who co-founded Apple Inc. He was known for his innovative ideas and his ability to merge technology and art.

You
What is he holding?


Assistant

Steve Jobs is holding an iPhone, which is a product he helped design and launch. The iPhone is a revolutionary smartphone that has had a significant impact on the way people communicate and access information.

You
Can you generate a image when he is young?

Assistant


Certainly! Here's an image of Steve Jobs when he was younger:



. This image shows him at a time when he was still in the early stages of his career, but already demonstrating his innovative spirit and passion for technology.

(c)


You
Who is in the picture?




Assistant

The image shows a young man working at a desk. The man is likely a software engineer or a computer programmer, as he is working on a computer. This profession requires a high level of technical knowledge and problem-solving skills.

You
Awesome! Which of the following books is better for his career from a long-term viewpoint:



or



Assistant

Both books can be beneficial for a software engineer's career. The first book, "15-Minute Fitness," focuses on physical fitness and wellness, which is important for maintaining a healthy lifestyle and reducing stress. The second book, "C++ Primer," is a comprehensive guide to the C++ programming language, which is a fundamental skill for software engineers. However, the choice between the two books depends on the individual's priorities and goals. If the person is more interested in physical fitness and wellness, the first book might be more beneficial. If they are more focused on improving their programming skills, the second book would be more relevant.

(d)

Figure 7: User cases of understanding both intrinsic & extrinsic information in the images.
9070

GPT-4 Prompt

Please construct a dialogue between a human and a helpful, honest and harmless assistant. The dialogue contains interleaved text and images. Each image is represented by `<imgX>` **DESCRIPTION** `</imgX>`, where **DESCRIPTION** is a textual description of the image and **X** is an index of the image. Please do not assume any further visual information beyond the description.

The constructed dialogues must and can only contain the following input images:

`<img0>` museum - the 1st nuclear submarine `</img0>`
`<img1>` response to the production of heavy `</img1>`

Characteristics about the assistant:

1. The assistant is trained to understand text, images, and their combinations.
2. The assistant can reply to the human with images and/or text.
3. The assistant has exceptional world knowledge and commonsense reasoning capabilities.
4. The assistant does not have access to the Internet or any other external tools.
5. If the assistant is asked to create an image, it can only show the image in the provided image list.
6. Please do not copy the images appearing in the dialogue. The assistant should refer to the previously mentioned image by natural language.

Characteristics about the human:

1. The human may send images and/or text to the assistant.
2. The human may ask questions requiring visual reasoning and/or understanding the relations between multiple images.
3. The human may ask the assistant to show images based on his/her intention.
4. The human may ask the assistant to do interesting things, rather than simply describing the content of the image.

Properties of a bad dialogue:

1. Simply describing or analyzing the content in the image.
2. Dialogue without a good logic.

Properties of a good dialogue:

1. Introducing extrinsic and attractive information of stuff shown in the image.
2. Discovering the connection between multiple images.
3. The dialogue happens in a reasonable context.

Examples:

=====

Case 1:

Input Images: ...

Output Dialogue: ...

Case 2:

Input Images: ...

Output Dialogue: ...

Case 3:

Input Images: ...

Output Dialogue: ...

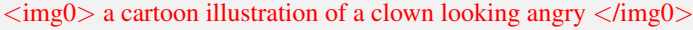
=====

Please directly give the dialogue if you understand. **The number of turns of the dialogue should be less than 6.** The dialogue should be self-contained. **Do NOT** assume any previous dialogue between the human and the assistant. Please use the same format `<imgX>` **DESCRIPTION** `</imgX>` to denote images in the dialogue and do not modify the description of the image. Note that the dialogue can stop at any time without an explicit ending.

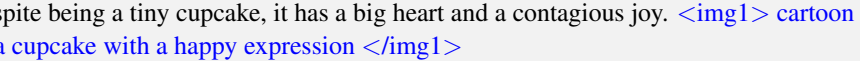
Figure 8: Illustration of the prompt used for TEXTBIND.

Example #1

Human: Hi! I'm trying to write a **cartoon story** for my kids featuring a grumpy clown and a lovable cupcake. Can you help me craft this?

Assistant: Of course! That sounds like a wonderful idea. We can start by establishing our characters. The grumpy clown could be someone who struggles finding joy in his own performances. 

Human: That's a wonderful idea! How about **the cute cupcake**?

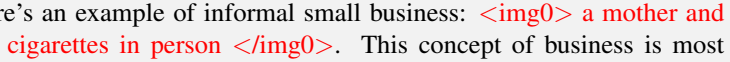
Assistant: The cupcake, on the other hand, is always upbeat and tries to help the clown find his happiness. Despite being a tiny cupcake, it has a big heart and a contagious joy. 

Human: Interesting! What should be their adventure?

Assistant: They could go on a journey to find the ultimate joke that guarantees laughter. The lovable cupcake, with his relentless positivity and cheer, helps the grumpy clown see the lighter side of life and finally find his joy.

Example #2


Human: Hi, I am researching small business operations. Can you show me an example of informal small business?

Assistant: Of course. Here's an example of informal small business: . This concept of business is most commonly found in rural or underdeveloped areas where formal establishments might not be as prevalent.

Human: Interesting! On a somewhat related note, what sells well in such arrangements?

Assistant: It depends on the location and the demand. Mostly, such businesses sell products of daily use such as food items, candy, toiletries etc. For example an ice cream cart is a common sight in many areas, as it's a popular treat among people of all ages.



Human: Can you show me an image of a kid indulging in something like that?

Assistant: Absolutely! Here's a delightful image of . Ice cream and similar snacks are often popular products for small informal businesses.

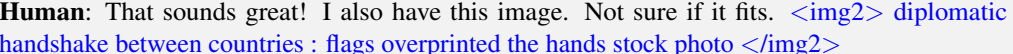
Example #3

Human: Hi there! I'm preparing a speech on unity and cooperation in global issues. I have a few images but I'm not sure how to describe them in a way that fits my theme.

Assistant: I'd be happy to help with that! Could you show me the images?

Human: Sure, here they are.  and 

Assistant: The first image can be seen as a symbol of unity. Here, numerous individuals come together to form the shape of a flag, showing that collective action can lead to a powerful representation, like a flag - often a symbol of identity and unity. The second image illustrates the idea of global responsibility and commitment. The male boxer, suggesting strength and determination, has a globe - representing the world - in his hands, which could indicate a keen sense of global responsibility. The flag painted on his back can symbolize allegiance or commitment to his country, and by extension to global harmony if we're expanding the theme to encompass international cooperation.

Human: That sounds great! I also have this image. Not sure if it fits. 

Assistant: This image very much fits your theme. It can represent a formal agreement or cohesion between countries. The flags overprinted on the shaking hands signify the countries involved in the meet. This image can be used to illustrate diplomatic dialogue or mutual understanding which is an integral component to solving global issues through cooperation and unity.

Figure 9: Illustration of in-context examples of TEXTBIND.

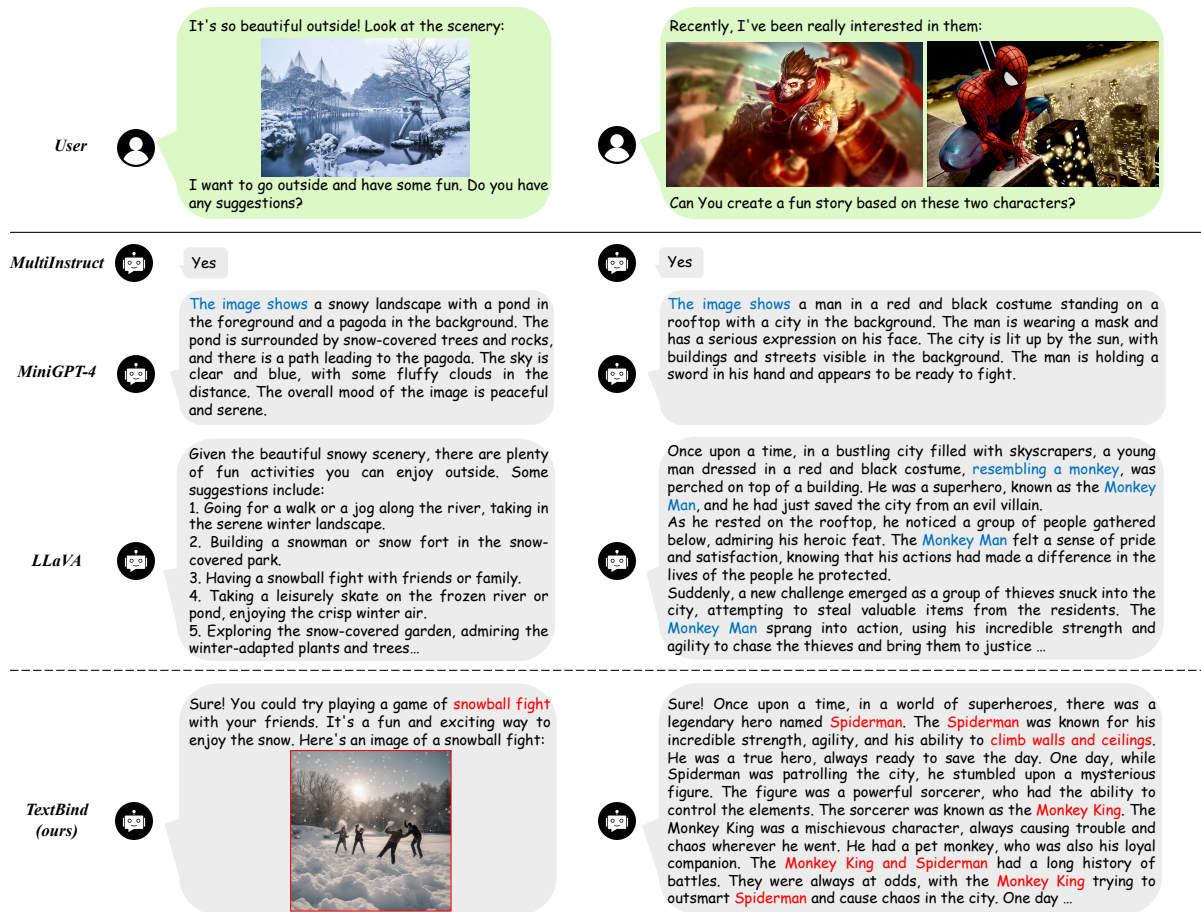


Figure 10: Comparison among models trained on different datasets.

Annotation	Labels	Description
Overall Quality	Excellent	This conversation is very interesting, practical, or intricate.
	Satisfactory	This conversation is coherent and reasonable without any factual errors.
	Poor	At least one turn in the conversation is unreasonable in some aspects, e.g., unrealistic content, illegal formats, etc.
Required Abilities	Image Creation	To create new images in appropriate contexts.
	Image Comparison	To combine, relate, or compare the information in different images.
	Intrinsic Image Understanding	To identify and recognize the objects, colors, shapes, and patterns in images.
	Extrinsic Image Understanding	To interpret the underlying meaning of images, e.g., the context, emotions, symbolism, or narrative conveyed by the images. It goes beyond merely recognizing the elements in the images and often requires external knowledge and/or deep analysis.

Table 8: Human annotation guideline.

Error Type	Percentage
Img-Cap Mismatch	0.03
Incoherence	0.03
Hallucination	0.03

Table 9: Error types and percentage in dataset constructed by TEXTBIND.

H Implementation Details (Model)

Our experiments are based on Huggingface Transformers⁸ (Wolf et al., 2020) and DeepSpeed⁹

⁶<https://huggingface.co/openai/clip-vit-base-patch16>


⁷<https://openai.com/blog/openai-api>

⁸<https://huggingface.co/docs/transformers/index>

⁹<https://github.com/microsoft/DeepSpeed>

Human: Hi! I am looking for some trendy rain boots. Can you show me an example?


Assistant: Sure! Here's an image of a rain boot design which might interest you.



'design on rain boots available.'

Human: That's a nice design. What else can you show me?

Assistant: I found a selection of fashionable rain boots in this image. These are some stylish options you can consider.




Beautiful Brown Boots

(a) Hallucination Case: The topic is about rain boot. However, not all the boots shown in the second image are rain boots.

Human: Hi! I'm planning a vacation and I want to find a peaceful place to relax. Can you recommend a location based on the images provided?


Assistant: Sure! Based on the images, I would recommend going to a place with a beautiful lake surrounded by mountains. It's a serene environment that allows you to relax and enjoy the natural beauty. Here's an image of a boat on a lake at sunset with mountains in the background.



www.alamy.com - 182096
"boat on a lake at sunset surrounded by mountains"

Human: That looks amazing! I also enjoy outdoor activities. Can you show me an image that represents an outdoor activity I can do in such a location?

Assistant: Absolutely! In a location like this, one of the classic outdoor activities you can do is rowing. Here's an image of a person in a rowboat on a lake with majestic mountains in the background:



www.alamy.com - 23410
"person in a row boat on a lake set against the backdrop of majestic mountains"
It's a great way to get some exercise while enjoying the stunning scenery.

(b) Incoherence Case: The first turn discusses about relaxing while the second turn is about outdoor activities. There are no connections.

Human: Hi, I am looking for a cozy accommodation for my weekend getaway, and I found an image of a cottage. Can you suggest what type of room would be suitable and comfortable?



'double room with private bathroom in a cottage'

Assistant: In the image you provided, a double room with a private bathroom in a cottage seems to be an excellent choice for your weekend getaway. This kind of room offers the perfect balance of privacy and comfort, ensuring you have a relaxing stay.

Human: That sounds lovely! What kind of bathroom can I expect in a room like this?



(c) Image-caption Mismatch Case: The first image only shows a bathroom, but the caption is "double room with private bathroom in a cottage".

Figure 11: Constructed conversations with "Poor" Label. The caption is shown below the image with gray color.

Training Stage	Epoch	Learning Rate	Batch Size	Max Sequence Length	Training Modules
Multimodal Alignment	2	1e-4	256	256	Q-Former, Linear
Multimodal Instruction Following	3	1e-5	64	768	Q-Former, Linear, LLM

Table 10: Training Configures of our Experiments

(Rasley et al., 2020). We use the filtered synthetic captions given by BLIP¹⁰ (Li et al., 2022), including Conceptual Captions (Changpinyo et al., 2021; Sharma et al., 2018) and SBU (Ordonez et al., 2011), totally 12M image-caption pairs. We employ the same visual encoder and Q-Former as used in BLIP-2 (Li et al., 2023b) and use their weights for initialization. LLama2-Chat¹¹ (Touvron et al., 2023) is utilized as the backbone language model. For the image generation model, we use Stable Diffusion XL¹² (Podell et al., 2023). The training configures are shown in Table 10. We use 8 NVIDIA A100 (40G) GPUs for all experiments.

I Human Evaluation Guideline

For overall quality assessment, we adopt the following criteria.

- Score 4: The response is excellent.
- Score 3: The response is acceptable but may not be very informative and interesting.
- Score 2: The response has minor issues, such as slight hallucinations when describing the images in the context.
- Score 1: The response is invalid and has significant drawbacks, e.g., irrelevant to the context.

For fine-grained annotations, intent recognition evaluates whether the response fulfills the the users’ intent, context understanding examines whether the response correctly comprehends the information in text and images, and the response informativeness measures the usefulness of the response. For each dimension, a human annotator will assign a score in {1, 2, 3, 4}. The four scores ranging from 1 to 4 indicate “major error”, “minor error”, “acceptable”, and “perfect”, respectively.

J Data Format of Model Input

The data format of model input is given in Fig. 12.

¹⁰<https://github.com/salesforce/BLIP>

¹¹<https://huggingface.co/meta-llama/LLama-2-7b-chat-hf>

¹²<https://github.com/Stability-AI/generative-models>

K TEXTBIND Usage

TEXTBIND is completely open-source and can be used for academic research and commercial purposes.

```
1 {
2   "conversation" {
3     {
4       "role": "user",
5       "content": "<image> Do you...",
6       "image_list": ["1235.png"]
7       "caption_list": ["a runing husky ..."]
8     },
9     {
10      "role": "assistant",
11      "content": "Yes, I do! ...",
12      "image_list": [],
13      "caption_list": []
14    },
15    ...
16  }
17 }
```

Figure 12: Data Format of Model Input.