

ModaVerse: Efficiently Transforming Modalities with LLMs

Xinyu Wang
University of Adelaide

Bohan Zhuang
Monash University

Qi Wu
University of Adelaide

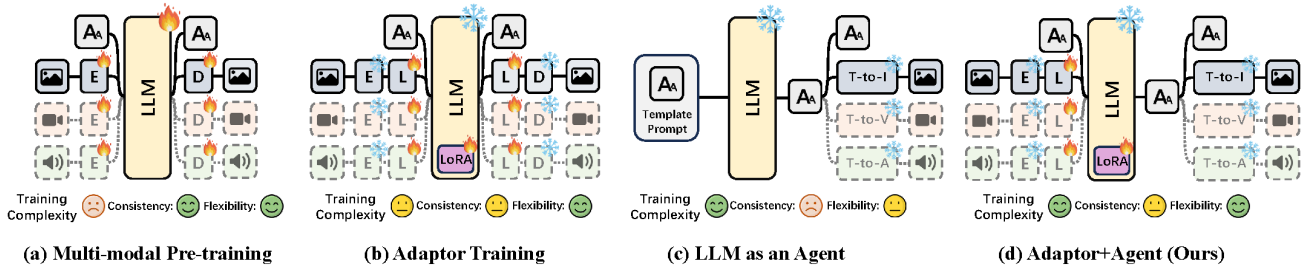


Figure 1. Comparative illustration of MLLM paradigms: (a) Multi-modal Pre-training, where new modules such as vision encoders and decoders are integrated within the standard LLM framework. (b) Adaptor Training, illustrating the use of projection layers to connect LLMs to pre-existing modules. (c) LLM as an Agent, highlighting the strategic application of prompts in conjunction with external tools. (d) Adaptor+Agent (ours), transforming modalities with efficient language-based Input/Output (I/O) alignment. E, D, and L represent the Encoder, Decoder, and Linear Layer respectively. T-to-x denotes a text-to-x generative model, where x can be Image, Video, and Audio.

Abstract

Humans possess the capability to comprehend diverse modalities and seamlessly transfer information between them. In this work, we introduce *ModaVerse*, a Multi-modal Large Language Model (MLLM) capable of comprehending and transforming content across various modalities including images, videos, and audio. Predominant MLLM frameworks have largely relied on aligning latent spaces of textual and non-textual features. This alignment process, which synchronizes a language model trained on textual data with encoders and decoders trained on multi-modal data, often necessitates extensive training of several projection layers in multiple stages. Inspired by LLM-as-agent methodologies, we propose a novel Input/Output (I/O) alignment mechanism that operates directly at the level of natural language. It aligns the LLM’s output with the input of generative models, avoiding the complexities associated with latent feature alignments, and simplifying the multiple training stages of existing MLLMs into a single, efficient process. By conducting experiments on several benchmarks, we demonstrate that our approach attains comparable performance with the state of the art while achieving considerable efficiencies in data usage. The code is available at <https://github.com/xinke-wang/ModaVerse>.

1. Introduction

Spanning from ancient inscriptions to contemporary online encyclopedias, texts have served as the quintessential medium for chronicling the expanse of human knowledge. Such a vast accumulation of textual data provides a fertile terrain for training Large Language Models (LLMs) [4, 31, 32, 43, 44]. Through extensive training on massive corpora, LLMs undergo a transformative process, a phenomenon captured by the concept where quantitative increases result in qualitative behavioral shifts [49], thus emerging with human-like reasoning abilities. This enables them to comprehend and respond to human instructions with remarkable precision. Such proficiency dramatically widens the scope of LLM applications across various domains, such as chatbots, programming copilots, and robotic agents.

Yet, the advent of richer communication forms calls for an evolution beyond the traditional confines of text. In the era where *a picture is worth a thousand words*, the capability to interpret and integrate complex visual and auditory data is invaluable. The pursuit of enabling LLMs to process and generate information beyond textual data reflects the natural progression of AI, aspiring to mimic the full breadth of human communication. This has spurred the evolution of Multi-modal LLMs (MLLMs), which are designed to understand, transform, and produce content across various modalities, such as images, audio, and video. This growing

interest has prompted a proliferation of research and innovation in the field [11, 17, 25, 40, 48, 50, 54, 56, 59]. Addressing the limitations of traditional text-only LLMs, multi-modal pretraining, adaptor training, and LLM as an agent emerge as three key paradigms for equipping LLMs with multi-modal capabilities. Figure 1 compares the existing paradigm’s overview schematic, assessing its performance and efficiency across three dimensions. Specifically, *Training Complexity* refers to the volume of training data, computational resources consumed, and the number of training stages involved. *Consistency* denotes the extent to which output is affected by modifications to inputs or prompts. *Flexibility* pertains to the capacity for a variety of interpreting and generating outputs under diverse conditions.

Multi-modal Pre-training (Figure 1 (a)) expands the traditional LLM framework to accommodate non-textual inputs and outputs by integrating additional modality encoders and decoders into the existing framework. Through custom-designed pre-training tasks, the LLM learns to represent multiple modalities effectively, achieving superior consistency and flexibility compared to existing paradigms. However, adapting a text-based LLM, which has been pre-trained on extensive textual data, to a multi-modal context often requires significant fine-tuning or even complete re-training. Therefore, this adaptation demands considerable computational resources. For example, Emu [40] combines Eva-CLIP [39], LLaMA [43], and Stable Diffusion [34] to develop a foundational multi-modal model. This development involves an intensive pre-training phase on a large-scale dataset and the use of hundreds of GPUs.

Adaptor Training (Figure 1 (b)) offers a computationally economical alternative. This strategy, demonstrated by BLIP-2 [19] and MiniGPT-4 [59], is typically built upon well-established LLMs and multi-modal encoders/decoders. It involves integrating these multi-modal modules with the LLM by training a set of projection layers while keeping the parameters of the LLM either frozen or fine-tuned using parameter-efficient techniques like LoRA [13]. These layers translate non-textual representations, such as image features, into the textual domain of LLMs, thus avoiding extensive training but preserving flexibility. However, despite reducing training data volume and time, these methods still require a complex training procedure. For example, NExT-GPT [50] employs a three-step training pipeline where the encode/decode-side projection layers and the LLM adaptor are each trained in distinct stages. This intricate setup substantially escalates the complexity and leads to redundancy in the training process.

LLM as an Agent (Figure 1 (c)) demonstrates a training free framework. These methods utilize the zero-shot inference capabilities of LLMs, emphasizing strategic prompt crafting and workflow design. This approach guides the interpreting and generation of multi-modal content through

interactions with external tools. For instance, Hugging-GPT [36] has developed a four-step pipeline that prompts OpenAI’s ChatGPT to select and execute models from the HuggingFace’s model zoo, thereby solving a variety of tasks. However, it is crucial to recognize that these methods, lacking targeted training, often rely on x-to-text or text-x models, for processing non-textual inputs. This reliance may result in limited flexibility in handling diverse data types. Additionally, the heavy dependence on the design of system prompts and the reasoning capabilities of LLMs can further lead to inconsistent results.

So far, each paradigm presents a specialized approach for achieving functionality in MLLMs, each with its advantages and limitations. Considering these trade-offs, exploring the integration of their strengths into a cohesive approach is compelling. Specifically, this paper proposes Adaptor+Agent, an approach that aims to find a harmonious balance between the efficiencies of the LLM-as-agent approaches and the flexibility of adaptor training methods.

Adaptor+Agent (Figure 1 (d)) aims to combine the benefits of adaptor training with LLM-as-agent methods. As shown in the figure, to maintain the flexibility of accepting arbitrary combinations of input modalities, we train a set of linear adaptors to map the input’s non-textual features into the LLM’s textual space. This approach allows the model to comprehend multi-modal inputs while preserving training efficiency by only tuning the adaptors. For the output, we adopt an LLM-as-agent design, using established text-to-x models for generating non-text outputs. This strategy avoids the need for tuning additional output-side projection layers, thus enhancing efficiency. The primary challenge in the Adaptor+Agent framework is *aligning the LLM’s output with the text-to-x models’ input*. To address this, we introduce Input/Output (I/O) Alignment. In contrast to previous adaptor-based approaches that focus on feature-level alignment between the LLM and generative models, our I/O Alignment strategy prompts the LLM to generate language-aligned meta-responses. These meta-responses contain detailed instructions for activating the generative models. We achieve this I/O Alignment through an instruction-following tuning process. As a result, in a single stage of tuning, the LLM is equipped to invoke external models for producing non-text outputs, thus bypassing the complex feature-level alignment typically required in the adaptor training paradigm.

In summary, the technical contributions of this paper are:

- We introduce a new Adaptor+Agent training paradigm for Multi-modal Large Language Models that synthesizes the strengths of both adaptor training and the LLM-as-Agent approach. This integration effectively reaps the benefits of training efficiency and model flexibility.
- To address the alignment challenges inherent in the LLM-as-Agent methodology, we propose an I/O Align-

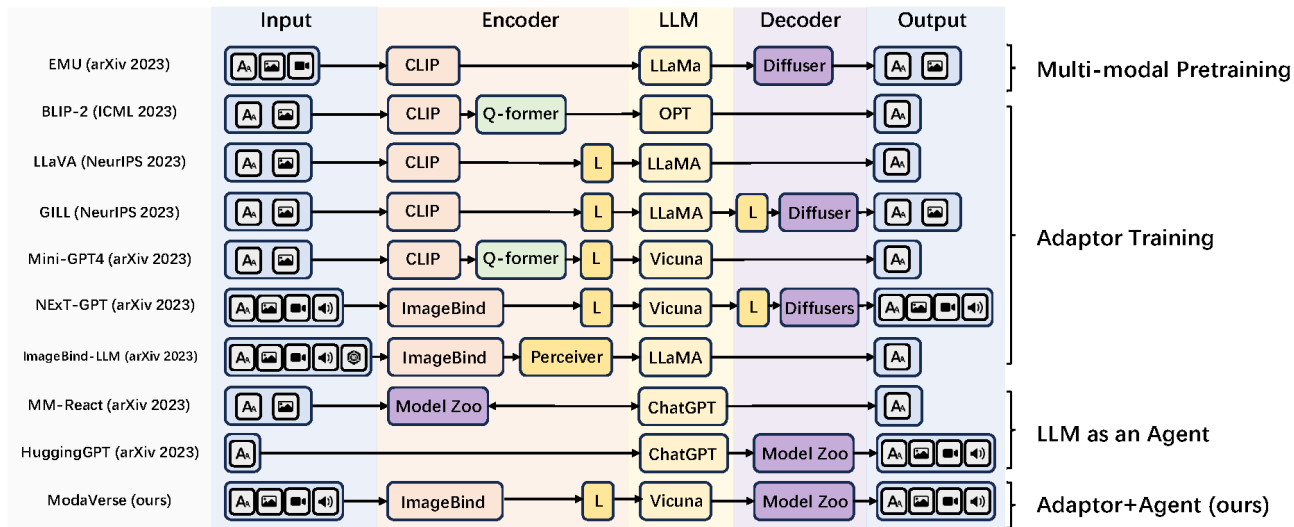


Figure 2. Comparison of the overview schematic of recent proposed MLLMs. L represents linear projection layers.

ment strategy. This strategy diverges from conventional feature-level alignment and instead operates at the natural language level, offering a more efficient alternative.

- Our final product, ModaVerse, demonstrates comparable performance to the current state of the arts on several widely used benchmarks while requiring fewer data and training resources, thereby offering a more efficient option without compromising effectiveness.

2. Related Work

Multi-modal Pretrained MLLM. Multi-modal pertaining is not a novel concept. Early efforts [21, 26, 37] have explored ways to extend the capabilities of language models to comprehend visual content. These models achieved promising performance on specific vision-language tasks [3, 47] but failed to generalize to more general scenarios. Recent advancements, however, have revealed that simpler model architectures can yield impressive outcomes when subjected to extensive large-scale pretraining, thanks to advanced computational resources and diverse datasets. For example, PaLI [6] demonstrates this by integrating a vision transformer with a language transformer and training on an extensive dataset of 10 billion image-text pairs. Similarly, CM3Leon [56] employs a straightforward decoder-only transformer architecture, trained on 340 million image-text pairs. This approach has enabled remarkable flexibility in generating and modifying both text and images, showing strong performance in image-to-text and text-to-image conversions. In addition, to enable an LLM that can generate non-text content, Emu [40] combines a stable diffusion model with the LLaMA as a decoder, trained on a diverse corpus of 82 million image-text and video-text pairs. This

integration marks a significant stride in the field, showcasing the growing versatility of LLM in multi-modal contexts.

Adaptor Trained MLLM: Leveraging recent advancements in parameter-efficient fine-tuning techniques [13] and data-efficient approaches [25], numerous studies have explored the feasibility of training adaptors for aligning features between LLMs and various non-textual modules. Flamingo [1] represents a pioneering effort in freezing the parameters of both visual encoders and LLMs, training a set of gated cross-attention layers to integrate visual knowledge into LLMs. However, it still necessitates extensive training on a massive dataset. Another notable example is BLIP-2 [18], which introduces a BERT-based Q-Former to translate image features into textual representations, thereby enabling LLMs to comprehend image content. This innovation has inspired subsequent research [5, 38], revealing that the Q-Former structure can be further simplified to a single linear layer, significantly reducing the number of trainable parameters. However, these advances, while showing considerable promise, have been predominantly applied to image-text tasks. In an effort to create a more versatile MLLM that can process a broader array of input types, subsequent works [11, 38] replaced the dual-modality encoder, CLIP [33], with the more versatile six-modality encoder, ImageBind [11]. In addition, other works [9, 17, 50, 57] try to align the output side of LLMs with generation models, enabling the utilization of latent features from LLMs to guide generative models in producing non-textual content. For example, a concurrent study, NExT-GPT [50], introduces a multi-stage training procedure. This procedure includes a series of adaptors that align the LLM with encoders and generative models at the feature level. Our work differs from NExT-GPT in that it aligns the generative mod-

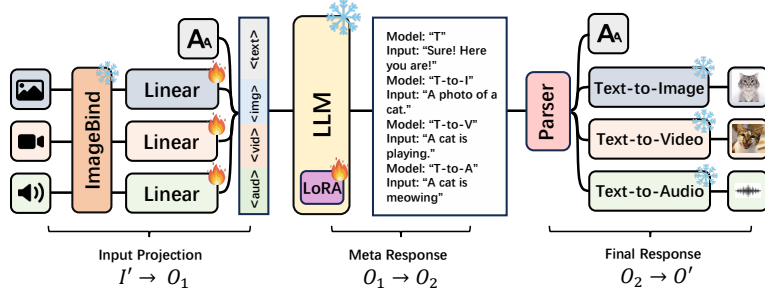


Figure 3. Overview of the Proposed ModaVerse Pipeline. In the input projection stage, multi-modal inputs I' are aligned to the LLM’s space O_1 using a series of trainable linear layers. During the meta-response generation stage, LLM is fine-tuned with a LoRA adaptor, prompting the generation of a meta-response O_2 . In the final response generation stage, additional pretrained text-to-x models are utilized to generate the ultimate multi-modal response O' based on the parsed meta response.

els at the language level instead of the feature level, thereby significantly reducing training complexity.

LLM-as-agent MLLM: The remarkable zero-shot inference capabilities of LLMs enable them to effectively utilize external tools [27, 35, 36, 41, 55]. This potential facilitates the creation of specialized pipelines and associated prompts, guiding LLMs to understand or produce multi-modal content. For instance, HuggingGPT [30] has developed a multi-step pipeline. In this process, ChatGPT initially interprets human instructions and selects appropriate models from a model zoo to accomplish the given tasks. Subsequently, the outputs from these external models are fed back into ChatGPT for parsing and generating the final response. Another notable example is MM-React [55], which introduces the integration of vision experts, such as OCR, image captioning, and object detection models, to extend the LLMs’ ability to process visual content. For each pair of input images and instructions, the LLM employs the relevant vision expert to extract pertinent information from the images, thereafter generating relevant responses.

Figure 2 presents a schematic overview of recently proposed MLLMs. It demonstrates the characteristic of adaptor training, wherein all models incorporate additional projection components, such as a linear layer or a Q-former, either before or after the LLM. These components are utilized to align textual and non-textual representations between the LLM and encoders/decoders. In contrast, multimodal pre-training methods usually feature a more straightforward and concise architecture. They direct the LLM itself to learn multimodal features, thus avoiding the projection structures between different modules. Furthermore, LLM-as-agent methods employ an external model zoo to assist in processing or producing non-textual content, without integrating trainable modules into the system. In comparison, the proposed Adaptor+Agent paradigm follows an adaptor structure on the encoder side, where linear projection layers are trained to align the input features with the LLM’s textual space. On the decoder side, the LLM is treated as an agent

to invoke external models for generating non-text content.

3. ModaVerse

3.1. Pipeline Overview

Figure 3 illustrates the comprehensive framework of the proposed ModaVerse, which contains three functional blocks, including input projection, meta-response generation, and final response generation.

Input Projection: To adapt a text-based LLM into an MLLM capable of interpreting non-textual inputs, it is essential to align the LLM’s textual features with various modalities during the input phase. Recent research [5, 38, 59] has demonstrated the feasibility of aligning these different modalities using a single linear layer. Following this, we employ ImageBind [11] as a unified encoder, which processes inputs from diverse data types, including images, videos, and audio, converting them into a specific embedding. Subsequently, for each modality, we learn a set of linear projection layers to map these encoded representations into the LLM’s text space. As a result, ModaVerse gains the capability to comprehend multi-modal inputs.

Meta Response Generation: Since the foundational LLM is pre-trained exclusively on text-only data, it lacks the capability to directly generate non-text outputs. To address this limitation, we treat the foundational LLM as an agent, designed to produce only meta-responses. As depicted in Figure 3, the meta-response comprises formatted information that includes the invocation details. For instance, according to the meta-response, the system might activate a text-to-image model to create an image based on the prompt “A photo of a cat”. This design circumvents the need for training an additional output-side projection layer to align the LLM’s feature space with that of generative models, thereby simplifying the training process.

Final Response Generation: This block incorporates several replaceable text-to-x models to generate the final response, which may include images [34], videos [28], and

audio [23]. Based on the invocation details parsed from the meta-responses, one or more models will be activated to produce the non-textual output.

So far, the Adaptor+Agent paradigm has become clear. In this paradigm, the input projection is designed with a set of linear adaptors that map multimodal features to the LLM’s textual space. The LLM itself is treated as an agent, invoking external models to generate the final responses. The benefits of such a design are:

- Training adaptors during the input phase preserve the details in the input data, while simultaneously reducing the training volume compared to full multimodal pretraining.
- Treating the LLM as an agent during the output phase not only decouples it from external generative models, enabling a plug-and-play approach but also eliminates the need for additional projection layers. This means that running generative models during the training stage is unnecessary, thereby reducing training complexity.

3.2. I/O Alignment

Consider a text-based LLM: $I \rightarrow O$ with its input and output sets defined as $I = O = \{text\}$, the objective of ModaVerse is to discover an efficient transformation that extends LLM to a multimodal model capable of handling $I' = O' = \{text, image, video, audio\}$, described as follows:

$$\begin{cases} P : \text{ImageBind}(I') \rightarrow O_1, \\ LLM' : O_1 \rightarrow O_2, \\ M : O_2 \rightarrow O' \end{cases} \quad (1)$$

Each line of this equation corresponds to a stage depicted in Figure 3. The first line denotes a trainable projection P from the multimodal feature - extracted by ImageBind [11] from the input I' - into the textual space of the LLM, O_1 . The second line involves tuning a LoRA [13] adaptor to prompt the adapted LLM, defined as LLM' , to generate a meta-response O_2 from the input feature O_1 . Finally, the third line, where M represents the frozen, established text-to-x model zoo, utilizes the parsed meta-response to generate the final multimodal outputs O' . To achieve these objectives, we propose an instruction-following I/O Alignment. This alignment aims to simultaneously fit the $I' \rightarrow O_1$ and $O_1 \rightarrow O_2$ alignments. As such, the trainable components, as depicted in Figure 3, consist of three linear layers and the LoRA adaptor. Specifically, the I/O Alignment involves two primary components: the construction of instructions, and the tuning of linear and LoRA adaptors.

To implement I/O Alignment, two issues must be addressed. First, since an exact representation of O_1 is not directly obtainable, existing adaptor-based methods typically train a direct mapping from $I' \rightarrow O_2$, using captions from paired datasets as O_2 to learn the projections, framing the process as a multi-modal captioning task. However, in our

case, O_2 is a meta response rather than just the text descriptions of I' . That is, given instructions and accompanying multi-modal inputs, the expected output should specifically identify *which* model to use and *how* to use it. For example, given the instruction, ‘Generate an image for an animal based on the provided audio clip of its vocalization’, along with an accompanying audio clip that records a cat’s meowing, the expected invocation information should be as follows: ‘{model: “text-to-image”, prompts: “a photo of a cat”}’. Therefore, simply using x-to-text datasets to train the projection layers under an image captioning task is not possible to facilitate such purposes. The second issue is the language-level misalignment due to the different training corpora of LLMs and generative models. For instance, to describe a landscape image, an LLM tends to produce coherent, literary paragraphs, whereas a text-to-image model typically prefers concise, descriptive prompts accompanied by attributes such as “4k” and “masterpiece”. Thus, I/O Alignment should also achieve $O_1 \rightarrow O_2$, ensuring language-level alignment between the meta-response and the input prompts required by generative models.

To address the issues mentioned above, I/O Alignment employs an instruction-following training approach. The adaptors are trained with an input that includes both a language instruction and accompanying multi-modal elements, with the aim to produce a meta-response that details the following invocation. We utilize training datasets from generative models to create pairs of instructions and their corresponding ground truths (see Section 3.3 for the instruction generation procedure). This method is beneficial for several reasons: 1. Instruction-following tuning compels the LLM to fully comprehend multi-modal inputs, thereby aiding in aligning the input projection layers between multi-modal input and LLM. 2. The training datasets of generative models typically provide both non-textual data, such as images, videos, or audio, and their corresponding textual descriptions, thereby offering a solid foundation of aligned data samples. 3. Most open-source generative models are trained on the same publicly available datasets. Aligning the meta-response with the text descriptions from these datasets means it is possible to seamlessly switch between generative models, thus facilitating a plug-and-play approach.

Based on this, we generate different types of instructions to fulfill the above objectives, which are as follows:

Input-side Alignment Instruction focuses on aligning the LLM’s capability to comprehend inputs comprising combinations of various modal data, such as text+image, image+video, or image+audio+video. For instance, when presented with a combination of image, audio, and video inputs, the instruction “Describe the given image, audio, and video” directs the MLLM to sequentially describe the content in the image, followed by the audio, and then the video.

Output-side Alignment Instruction aims to align the

Method	Type	Input	Output	Stage I		Stage II		Stage III	
				Data	GPU Time	Data	GPU Time	Data	GPU Time
Emu [40]	pretrain	T,I,V	T,I	82M	128×2d	24M	32×	1.28M	16×16h
LLaVA [25]	adaptor	T,I	T	595k	8×20h	158k	8×10h	N/A	N/A
BLIP-2 [18]	adaptor	T,I	T	129M	16×6d	129M	16×3d	N/A	N/A
NExT-GPT [50]	adaptor	T,I,V,A	T,I,V,A	13M+	-	13M+	-	20k+	-
ModaVerse	adaptor+agent	T,I,V,A	T,I,V,A	2M	4×20h	N/A	N/A	N/A	N/A

Table 1. Comparison of training complexity of ModaVerse with recent MLLMs. ‘N/A’ indicates stages that are not required, while ‘-’ denotes that the data was not disclosed by the authors. ‘T’, ‘I’, ‘V’, ‘A’ are the abbreviations of text, image, video, and audio, respectively.

LLM’s ability to generate meta-responses that include invocation details, such as the selected model and prompts. For example, the instruction “Generate an image based on the provided audio of an animal sound” teaches the model to utilize a text-to-image model to generate an image, potentially with a prompt like “A photo of a cat”.

Reasoning Boosting Instruction is designed to preserve and enhance the LLM’s reasoning capabilities through a diverse range of topics. For instance, an instruction like “Where might this audio clip have been recorded?” requires the LLM to make a reasoned inference based on the input data, thereby strengthening its reasoning skills.

3.3. Instruction Generation and Training

In this section, we introduce how to generate the instruction invocation pairs used in I/O Alignment.

```

1 {"instruction": ["Generate an image of
   an animal based on the provided
   vocalization.", "cat_meowing.wav", ]
2 "invocation": [("text-to-image", "A
   photo of a cat"), ]

```

As demonstrated in the code block above, an instruction-invocation pair consists of two parts: the instruction, which represents the input, and the invocation, which represents the expected output of the LLM. To obtain these forms of data samples, we constructed them from two sources. The first source utilizes components of existing instruction tuning works, such as LLaVA [25], VideoChat [20], and InstructBLIP [24]. However, it should be noted that the instructions in these works primarily fall into two categories: Input-side Alignment Instructions and Reasoning Boosting Instructions. Thus, to generate Output-side Alignment Instructions, which are crucial for the success of ModaVerse, we created specific templates to assist the OpenAI ChatGPT API in producing new instructions on a large scale. Specifically, each query sent to the ChatGPT API comprises three components: Seed Examples, Candidate Descriptions, and Language References.

Seed Examples consist of a set of standard instruction invocation pairs, randomly selected from a manually crafted collection. These examples serve as guides for the Chat-

GPT API, demonstrating how to generate samples in the given format and providing an illustration of the task.

Candidate Descriptions comprise randomly selected text descriptions from the paired datasets, which include descriptions of images, audio, and videos. These descriptions aim to mimic the true inputs, while the ChatGPT API is requested to generate appropriate instructions and invocation details based on these candidates and seed examples.

Language References include text descriptions randomly selected from the training set of generative models. These samples serve as a guide for the ChatGPT API to learn the language style of the prompts used in the generative models, helping to generate language-aligned invocation details.

For training, we use Vicuna [7] as the foundation LLM, the trainable parts of the proposed ModaVerse (see Figure 3) consist only of three linear layers and the LoRA adaptor of the LLM. Table 1 compares the training complexity of ModaVerse with that of some recently proposed MLLMs. It shows that the proposed method enjoys lower training complexity.

4. Experiments

4.1. Quantitative Results

To evaluate the proposed ModaVerse, we follow previous works [42] to assess the model’s understanding ability ($x \rightarrow \text{text}$) and its generation ability ($\text{text} \rightarrow x$), where x can be image, audio, and video. Tables 2 and 3 illustrate the text-to-image and image-to-text performance on the COCO caption dataset [22]. The results demonstrate that our method achieved an 11.28 FID score in the image generation task, comparable to recent methods. In terms of image understanding capability, the proposed ModaVerse outperforms the any-to-any diffuser CoDi [42] in both B@4 and METEOR metrics, though it is slightly lower than NExT-GPT [50] and OFA [46]. The text-to-audio and audio-to-text performances on the AudioCaps dataset [15] are showed in Tables 4 and 5. The proposed method outperforms NExT-GPT and is narrowly eclipsed by CoDi. In audio captioning, our approach secures the second-best performance, rivaling the state-of-the-art models. Tables 6 and 7 showcase the text-to-video and video-to-text perfor-

Method	FID (\downarrow)
CogVideo [8] (NeurIPS’21)	27.10
GLIDE [29] (ICML’22)	12.24
CoDi [42] (NeurIPS’23)	11.26
SD [34] (CVPR’22)	11.21
NExT-GPT [50] (arXiv’23)	11.28
ModaVerse	11.24

Table 2. Text-to-image performance on COCO-caption [22].

Method	FD (\downarrow)	IS (\uparrow)
DiffSound [53] (TASLP’23)	47.68	4.01
AudioLDM-S [23] (ICML’23)	29.48	6.90
AudioLDM-L [23] (ICML’23)	23.31	8.13
CoDi [42] (NeurIPS’23)	22.90	8.77
NExT-GPT [50] (arXiv’23)	23.58	8.35
ModaVerse	23.40	8.22

Table 4. Text-to-audio performance on AudioCaps [15].

Method	FID (\downarrow)	CLIPSIM (\uparrow)
CogVideo [12] (NeurIPS’21)	23.59	0.2631
MakeVideo [14] (ICML’23)	13.17	0.3049
Latent-VDM [34] (CVPR’22)	14.25	0.2756
Latent-Shift [2] (arXiv’23)	15.23	0.2773
CoDi [42] (NeurIPS’23)	—	0.2890
NExT-GPT [50] (arXiv’23)	13.04	0.3085
ModaVerse	13.35	0.3014

Table 6. Text-to-video performance on MSR-VTT [52].

mances on MSR-VTT [52]. Our method achieved a 13.35 FID score and a 0.3014 CLIPSIM score in video generation, demonstrating parity with top-tier methods. Similarly, our approach shows competitive results in video-to-text tasks, as evidenced by its B@4 and METEOR scores.

Although our method does not outperform all state-of-the-art methods (including two concurrent arXiv submissions) across the six benchmarks, it is important to note the efficiency of the proposed ModaVerse. First, our method is capable of converting a variety of modalities, whereas some state-of-the-art methods, such as SD [34] and OFA [46], are specifically designed for single-route conversions like text-to-image. Second, as Table 1 illustrates, ModaVerse benefits from a more efficient training paradigm. It requires less data and fewer computational resources. Specifically, in contrast to NExT-GPT which necessitates three stages to independently train the projection layers, our method streamlines this process into a single stage. Regarding training data, our approach uses less than 2% of the data volume required by Emu [17] and BLIP-2 [18].

4.2. Qualitative Results

Since publicly available datasets are limited to certain common modality combinations, such as image-to-text and text-to-video, this limitation may not comprehensively capture

Method	B@4 (\uparrow)	METEOR (\uparrow)	CIDEr (\uparrow)
Oscar [21] (ECCV’20)	36.58	30.4	124.12
BLIP-2 [19] (ICML’23)	43.7	—	145.8
OFA [46] (ICML’22)	44.9	32.5	154.9
CoDi [42] (NeurIPS’23)	40.2	31.0	149.9
NExT-GPT [50] (arXiv’23)	44.3	32.9	156.7
ModaVerse	43.9	31.8	151.4

Table 3. Image-to-text performance on COCO-caption data [22].

Method	SPIDEr (\uparrow)	CIDEr (\uparrow)
AudioCaps [15] (NAACL’19)	0.369	0.593
BART [10] (DCASE’21)	0.465	0.753
AL-MixGen [16] (ArXiv’22)	0.466	0.755
CoDi [42] (NeurIPS’23)	0.480	0.789
NExT-GPT [50] (arXiv’23)	0.521	0.802
ModaVerse	0.494	0.792

Table 5. Audio-to-text performance on AudioCaps [15].

Method	B@4 (\uparrow)	METEOR (\uparrow)
ORG-TRL [58] (CVPR’20)	43.6	28.8
GIT [45] (TMLR’22)	54.8	33.1
mPLUG-2 [51] (ICML’23)	57.8	34.9
CoDi [42] (NeurIPS’23)	52.1	32.5
NExT-GPT [50] (arXiv’23)	58.4	38.5
ModaVerse	56.5	35.2

Table 7. Video-to-text performance on MSR-VTT [52].

the full extent of ModaVerse’s capabilities. Therefore, Figure 4 showcases various qualitative results of ModaVerse across different modalities. For instance, examples (a), (c), (f), and (l) emphasize the model’s conditioned generative abilities. In addition, examples (g), (h), and (i) demonstrate its proficiency in answering questions with inputs from a variety of modalities. Moreover, examples (d) and (e) demonstrate the potential for style transfer.

4.3. Limitations and Failure Cases

In exploring the capabilities of ModaVerse, Figure 5 includes some challenging scenarios where the model’s performance can be further enhanced. Specifically, Example (a) illustrates a current limitation of the model in image editing tasks, where it fails to retain the original background and layout of the input images. Instead of modifying the existing image, the model generates an entirely new one. This limitation highlights a specific challenge in our approach, particularly for tasks that require fidelity to the original image’s resolution and details. However, this can potentially be addressed by integrating an additional editing model into the model zoo at the final response generation stage, a development we leave for future work. Another notable case is Example (b), where, in the absence of language clues at the input phase, the model tends to produce random, ir-

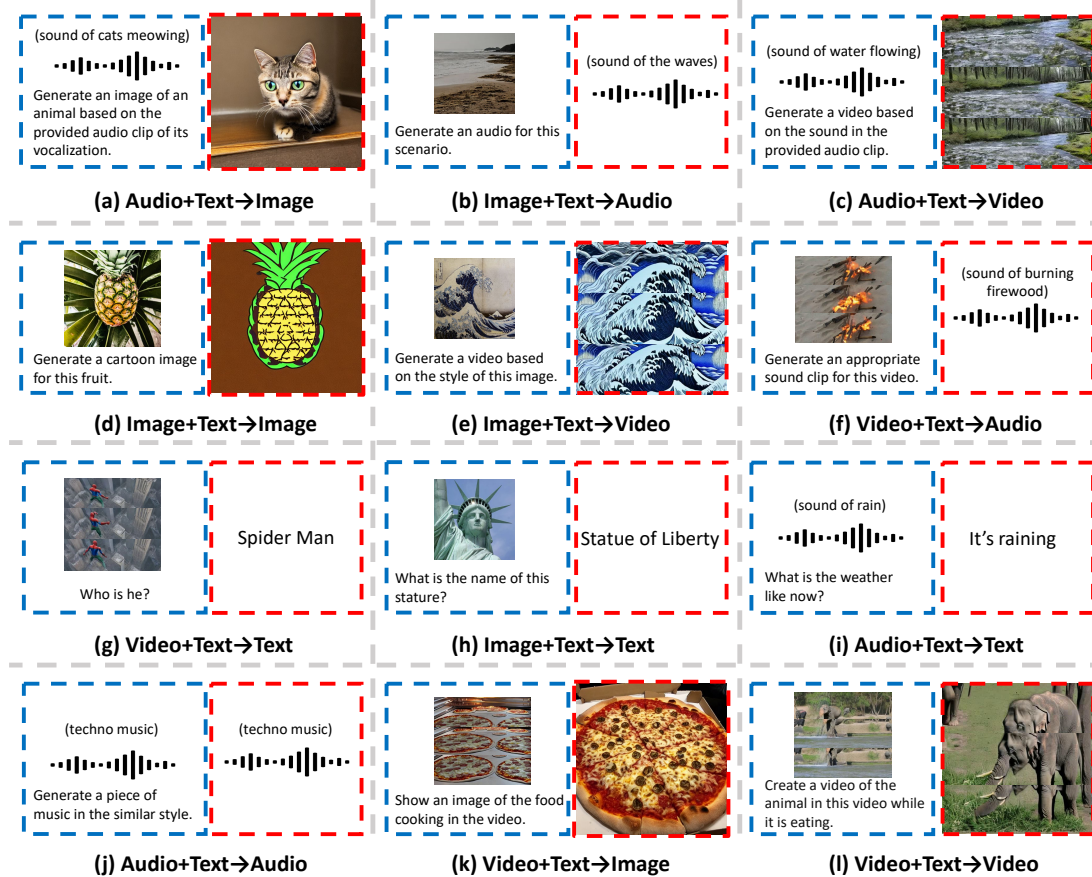


Figure 4. Qualitative examples of the proposed ModaVerse interpreting and producing data presented in combinations of various modalities. Blue and Red dashed boxes represent input and output respectively.




Figure 5. Failure cases of ModaVerse. (a) The model can only generate entirely new images and cannot modify the original pixels. (b) The model tends to generate irrelevant outputs in the absence of language instructions during the input phase.

relevant outputs. This issue arises because the instruction-following trained model relies on given language instructions for reasoning out the expected response. Without such clues, it may struggle to produce appropriate responses.

5. Conclusion

In this paper, we have presented ModaVerse, a MLLM capable of interpreting and generating data in various modalities. This model diverges from existing MLLM frameworks by adopting a synergistic approach that merges adaptor train-

ing with the LLM-as-agent methodology. By employing adaptors, ModaVerse effectively aligns the text-based LLM with multi-modal inputs through a set of linear projection layers. This enhances its capability to interpret a diverse array of input modalities. On the output side, instead of training additional projection layers to align the output space with generative models, we treat the LLM as an agent. This agent produces a meta-response containing invocation details, which are then parsed to activate generative models for generating the final response. This integrative Adaptor+Agent training paradigm not only streamlines the complex multi-stage feature alignment process but also significantly boosts the efficiency of the training process, offering an alternative for the training of MLLMs. For future work, we aim to address the current framework’s limitations and weaknesses, such as preserving the original layout information of inputs, thereby broadening its applicability to scenarios requiring original information, like image and video editing.

References

- [1] Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katherine Millican, Malcolm Reynolds, et al. Flamingo: a visual language model for few-shot learning. *NeurIPS*, 2022. 3
- [2] Jie An, Songyang Zhang, Harry Yang, Sonal Gupta, Jia-Bin Huang, Jiebo Luo, and Xi Yin. Latent-shift: Latent diffusion with temporal shift for efficient text-to-video generation. *arXiv preprint arXiv:2304.08477*, 2023. 7
- [3] Stanislaw Antol, Aishwarya Agrawal, Jiasen Lu, Margaret Mitchell, Dhruv Batra, C Lawrence Zitnick, and Devi Parikh. Vqa: Visual question answering. In *ICCV*, 2015. 3
- [4] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *NeurIPS*, 2020. 1
- [5] Jun Chen, Deyao Zhu, Xiaoqian Shen, Xiang Li, Zechun Liu, Pengchuan Zhang, Raghuraman Krishnamoorthi, Vikas Chandra, Yunyang Xiong, and Mohamed Elhoseiny. Minigt-v2: large language model as a unified interface for vision-language multi-task learning. *arXiv preprint arXiv:2310.09478*, 2023. 3, 4
- [6] 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. *ICLR*, 2023. 3
- [7] Wei-Lin Chiang, Zhuohan Li, Zi Lin, Ying Sheng, Zhanghao Wu, Hao Zhang, Lianmin Zheng, Siyuan Zhuang, Yonghao Zhuang, Joseph E Gonzalez, et al. Vicuna: An open-source chatbot impressing gpt-4 with 90%* chatgpt quality. See <https://vicuna.lmsys.>, 2023. 6
- [8] Ming Ding, Zhuoyi Yang, Wenyi Hong, Wendi Zheng, Chang Zhou, Da Yin, Junyang Lin, Xu Zou, Zhou Shao, Hongxia Yang, et al. Cogview: Mastering text-to-image generation via transformers. *NeurIPS*, 2021. 7
- [9] Runpei Dong, Chunrui Han, Yuang Peng, Zekun Qi, Zheng Ge, Jinrong Yang, Liang Zhao, Jianjian Sun, Hongyu Zhou, Haoran Wei, et al. Dreamllm: Synergistic multimodal comprehension and creation. *arXiv preprint arXiv:2309.11499*, 2023. 3
- [10] Félix Gontier, Romain Serizel, and Christophe Cerisara. Automated audio captioning by fine-tuning bart with audioset tags. In *DCASE*, 2021. 7
- [11] Jiaming Han, Renrui Zhang, Wenqi Shao, Peng Gao, Peng Xu, Han Xiao, Kaipeng Zhang, Chris Liu, Song Wen, Ziyu Guo, et al. Imagebind-llm: Multi-modality instruction tuning. *arXiv preprint arXiv:2309.03905*, 2023. 2, 3, 4, 5
- [12] Wenyi Hong, Ming Ding, Wendi Zheng, Xinghan Liu, and Jie Tang. Cogvideo: Large-scale pretraining for text-to-video generation via transformers. *ICLR*, 2023. 7
- [13] Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. Lora: Low-rank adaptation of large language models. *ICLR*, 2022. 2, 3, 5
- [14] Rongjie Huang, Jiawei Huang, Dongchao Yang, Yi Ren, Luping Liu, Mingze Li, Zhenhui Ye, Jinglin Liu, Xiang Yin, and Zhou Zhao. Make-an-audio: Text-to-audio generation with prompt-enhanced diffusion models. *ICML*, 2023. 7
- [15] Chris Dongjoo Kim, Byeongchang Kim, Hyunmin Lee, and Gunhee Kim. Audiocaps: Generating captions for audios in the wild. In *NAACL*, 2019. 6, 7
- [16] Eungbeom Kim, Jinhee Kim, Yoori Oh, Kyungsu Kim, Minju Park, Jaehoon Sim, Jinwoo Lee, and Kyogu Lee. Improving audio-language learning with mixgen and multi-level test-time augmentation. *arXiv preprint arXiv:2210.17143*, 2022. 7
- [17] Jing Yu Koh, Daniel Fried, and Ruslan Salakhutdinov. Generating images with multimodal language models. *NeurIPS*, 2023. 2, 3, 7
- [18] Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. *ICML*, 2023. 3, 6, 7
- [19] Junnan Li, Dongxu Li, Silvio Savarese, and Steven C. H. Hoi. BLIP-2: bootstrapping language-image pre-training with frozen image encoders and large language models. In *ICML*, 2023. 2, 7
- [20] KunChang Li, Yinan He, Yi Wang, Yizhuo Li, Wenhai Wang, Ping Luo, Yali Wang, Limin Wang, and Yu Qiao. Videochat: Chat-centric video understanding. *arXiv preprint arXiv:2305.06355*, 2023. 6
- [21] Xiujuan Li, Xi Yin, Chunyuan Li, Pengchuan Zhang, Xiaowei Hu, Lei Zhang, Lijuan Wang, Houdong Hu, Li Dong, Furu Wei, Yejin Choi, and Jianfeng Gao. Oscar: Object-semantics aligned pre-training for vision-language tasks. In *ECCV*, 2020. 3, 7
- [22] 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 *ECCV*. Springer, 2014. 6, 7
- [23] Haohe Liu, Zehua Chen, Yi Yuan, Xinhao Mei, Xubo Liu, Danilo Mandic, Wenwu Wang, and Mark D Plumbley. Audioldm: Text-to-audio generation with latent diffusion models. *ICML*, 2023. 5, 7
- [24] Haotian Liu, Chunyuan Li, Yuheng Li, and Yong Jae Lee. Improved baselines with visual instruction tuning. *arXiv preprint arXiv:2310.03744*, 2023. 6
- [25] Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. *NeurIPS*, 2023. 2, 3, 6
- [26] Jiasen Lu, Dhruv Batra, Devi Parikh, and Stefan Lee. Vilbert: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. *NeurIPS*, 2019. 3
- [27] Pan Lu, Baolin Peng, Hao Cheng, Michel Galley, Kai-Wei Chang, Ying Nian Wu, Song-Chun Zhu, and Jianfeng Gao. Chameleon: Plug-and-play compositional reasoning with large language models. *NeurIPS*, 2023. 4
- [28] Zhengxiong Luo, Dayou Chen, Yingya Zhang, Yan Huang, Liang Wang, Yujun Shen, Deli Zhao, Jingren Zhou, and Tieniu Tan. Videofusion: Decomposed diffusion models for high-quality video generation. In *CVPR*, 2023. 4
- [29] Alex 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. *ICML*, 2022. 7
- [30] OpenAI. Chatgpt. <https://openai.com/blog/chatgpt>, 2023. 4
- [31] Alec Radford, Karthik Narasimhan, Tim Salimans, Ilya Sutskever, et al. Improving language understanding by generative pre-training. 2018. 1
- [32] Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. Language models are unsupervised multitask learners. *OpenAI blog*, 2019. 1
- [33] 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*. PMLR, 2021. 3
- [34] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *CVPR*, 2022. 2, 4, 7
- [35] Timo Schick, Jane Dwivedi-Yu, Roberto Dessì, Roberta Raïleanu, Maria Lomeli, Luke Zettlemoyer, Nicola Cancedda, and Thomas Scialom. Toolformer: Language models can teach themselves to use tools. *NeurIPS*, 2023. 4
- [36] Yongliang Shen, Kaitao Song, Xu Tan, Dongsheng Li, Weiming Lu, and Yueting Zhuang. Hugginggpt: Solving ai tasks with chatgpt and its friends in huggingface. *NeurIPS*, 2023. 2, 4
- [37] Weijie Su, Xizhou Zhu, Yue Cao, Bin Li, Lewei Lu, Furu Wei, and Jifeng Dai. Vi-bert: Pre-training of generic visual-linguistic representations. *ICLR*, 2020. 3
- [38] 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. 3, 4
- [39] Quan Sun, Yuxin Fang, Ledell Wu, Xinlong Wang, and Yue Cao. Eva-clip: Improved training techniques for clip at scale. *arXiv preprint arXiv:2303.15389*, 2023. 2
- [40] Quan Sun, Qiyang Yu, Yufeng Cui, Fan Zhang, Xiaosong Zhang, Yueze Wang, Hongcheng Gao, Jingjing Liu, Tiejun Huang, and Xinlong Wang. Generative pretraining in multimodality. *arXiv preprint arXiv:2307.05222*, 2023. 2, 3, 6
- [41] Dídac Surís, Sachit Menon, and Carl Vondrick. Vipergpt: Visual inference via python execution for reasoning. *ICCV*, 2023. 4
- [42] Zineng Tang, Ziyi Yang, Chenguang Zhu, Michael Zeng, and Mohit Bansal. Any-to-any generation via composable diffusion. *NeurIPS*, 2023. 6, 7
- [43] 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. 1, 2
- [44] Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruiti Bhosale, et al. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*, 2023. 1
- [45] Jianfeng Wang, Zhengyuan Yang, Xiaowei Hu, Linjie Li, Kevin Lin, Zhe Gan, Zicheng Liu, Ce Liu, and Lijuan Wang. GIT: A generative image-to-text transformer for vision and language. *TMLR*, 2022. 7
- [46] 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 *ICML*, 2022. 6, 7
- [47] Xinyu Wang, Yuliang Liu, Chunhua Shen, Chun Chet Ng, Canjie Luo, Lianwen Jin, Chee Seng Chan, Anton van den Hengel, and Liangwei Wang. On the general value of evidence, and bilingual scene-text visual question answering. In *CVPR*, pages 10126–10135, 2020. 3
- [48] Xinyu Wang, Bohan Zhuang, and Qi Wu. Switchgpt: Adapting large language models for non-text outputs. *arXiv preprint arXiv:2309.07623*, 2023. 2
- [49] Jason Wei, Yi Tay, Rishi Bommasani, Colin Raffel, Barret Zoph, Sebastian Borgeaud, Dani Yogatama, Maarten Bosma, Denny Zhou, Donald Metzler, et al. Emergent abilities of large language models. *TMLR*, 2022. 1
- [50] Shengqiong Wu, Hao Fei, Leigang Qu, Wei Ji, and Tat-Seng Chua. Next-gpt: Any-to-any multimodal llm. *arXiv preprint arXiv:2309.05519*, 2023. 2, 3, 6, 7
- [51] Haiyang Xu, Qinghao Ye, Ming Yan, Yaya Shi, Jiabo Ye, Yuanhong Xu, Chenliang Li, Bin Bi, Qi Qian, Wei Wang, et al. mplug-2: A modularized multi-modal foundation model across text, image and video. *ICML*, 2023. 7
- [52] Jun Xu, Tao Mei, Ting Yao, and Yong Rui. Msr-vtt: A large video description dataset for bridging video and language. In *CVPR*, pages 5288–5296, 2016. 7
- [53] Dongchao Yang, Jianwei Yu, Helin Wang, Wen Wang, Chao Weng, Yuexian Zou, and Dong Yu. Diffsound: Discrete diffusion model for text-to-sound generation. *TASLP*, 2023. 7
- [54] Zhengyuan Yang, Linjie Li, Kevin Lin, Jianfeng Wang, Chung-Ching Lin, Zicheng Liu, and Lijuan Wang. The dawn of Imms: Preliminary explorations with gpt-4v (ision). *arXiv preprint arXiv:2309.17421*, 2023. 2
- [55] Zhengyuan Yang, Linjie Li, Jianfeng Wang, Kevin Lin, Ehsan Azarnasab, Faisal Ahmed, Zicheng Liu, Ce Liu, Michael Zeng, and Lijuan Wang. Mm-react: Prompting chatgpt for multimodal reasoning and action. *arXiv preprint arXiv:2303.11381*, 2023. 4
- [56] Lili Yu, Bowen Shi, Ramakanth Pasunuru, Benjamin Muller, Olga Golovneva, Tianlu Wang, Arun Babu, Binh Tang, Brian Karrer, Shelly Sheynin, et al. Scaling autoregressive multimodal models: Pretraining and instruction tuning. *arXiv preprint arXiv:2309.02591*, 2023. 2, 3
- [57] Pan Zhang, Xiaoyi Dong Bin Wang, Yuhang Cao, Chao Xu, Linke Ouyang, Zhiyuan Zhao, Shuangrui Ding, Songyang Zhang, Haodong Duan, Hang Yan, et al. Internlm-xcomposer: A vision-language large model for advanced text-image comprehension and composition. *arXiv preprint arXiv:2309.15112*, 2023. 3
- [58] Ziqi Zhang, Yaya Shi, Chunfeng Yuan, Bing Li, Peijin Wang, Weiming Hu, and Zheng-Jun Zha. Object relational graph with teacher-recommended learning for video captioning. In *CVPR*, 2020. 7

- [59] 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. [2](#), [4](#)