

# OneLLM: One Framework to Align All Modalities with Language

Jiaming Han<sup>1,2</sup>, Kaixiong Gong<sup>1,2</sup>, Yiyuan Zhang<sup>1,2</sup>, Jiaqi Wang<sup>2</sup>, Kaipeng Zhang<sup>2</sup>  
Dahua Lin<sup>1,2</sup>, Yu Qiao<sup>2</sup>, Peng Gao<sup>2</sup>, Xiangyu Yue<sup>1†</sup>

<sup>1</sup>MMLab, The Chinese University of Hong Kong

<sup>2</sup>Shanghai Artificial Intelligence Laboratory

## Abstract

Multimodal large language models (MLLMs) have gained significant attention due to their strong multimodal understanding capability. However, existing works rely heavily on modality-specific encoders, which usually differ in architecture and are limited to common modalities. In this paper, we present **OneLLM**, an MLLM that aligns **eight** modalities to language using a unified framework. We achieve this through a unified multimodal encoder and a progressive multimodal alignment pipeline. In detail, we first train an image projection module to connect a vision encoder with LLM. Then, we build a universal projection module (UPM) by mixing multiple image projection modules and dynamic routing. Finally, we progressively align more modalities to LLM with the UPM. To fully leverage the potential of OneLLM in following instructions, we also curated a comprehensive multimodal instruction dataset, including **2M** items from image, audio, video, point cloud, depth/normal map, IMU and fMRI brain activity. OneLLM is evaluated on **25** diverse benchmarks, encompassing tasks such as multimodal captioning, question answering and reasoning, where it delivers excellent performance. Code, data, model and online demo are available at <https://github.com/csuhan/OneLLM>.

## 1. Introduction

Large Language Models (LLMs) are getting increasingly popular in the research community and industry due to their powerful language understanding and reasoning capabilities. Notably, LLMs such as GPT4 [60] have reached performance nearly on par with humans in various academic exams. The progress in LLMs has also inspired researchers to employ LLMs as an interface for multimodal tasks, such as vision-language learning [3, 42], audio and speech recognition [23, 94], video understanding [10, 43, 95], etc.

<sup>†</sup> Corresponding author

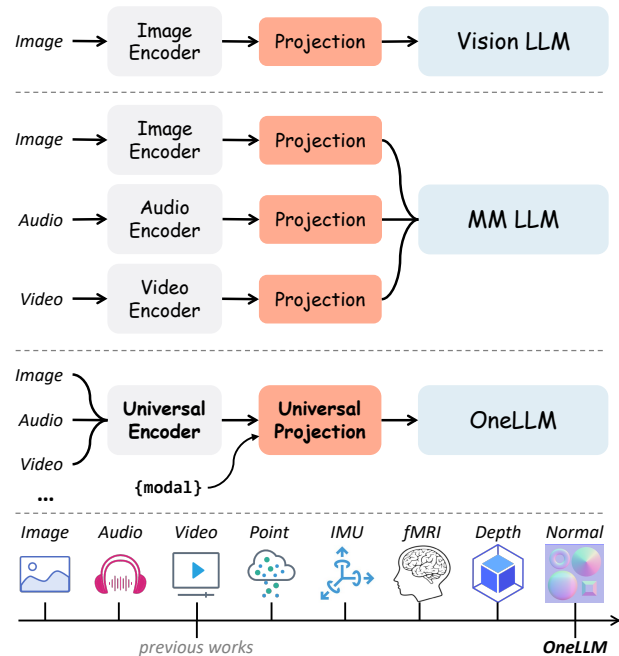


Figure 1. **Comparisons of Different Multimodal LLMs.** Vision LLM: one image encoder and projection module. Multimodal (MM) LLM: modality-specific encoder and projection module. **OneLLM**: a universal encoder, a universal projection module and modality tokens {modal} to switch between modalities. **Bottom**: OneLLM expands supported modalities from three to eight.

Among these tasks, vision-language learning is the most active field, with more than 50 vision LLMs proposed in the recent half-year alone [19]. Typically, a vision LLM comprises a visual encoder, an LLM, and a projection module connecting the two components. The vision LLM is first trained on massive paired image-text data [67] for vision-language alignment and then fine-tuned on visual instruction datasets, enabling it to complete various instructions tied to visual inputs. Beyond vision, significant efforts have been invested in developing other modality-specific LLMs, such as audio [23], video [43], and point clouds [26]. These models generally mirror the architectural framework and training methodology of vision LLMs, and rely on the

solid foundation of pretrained modality-specific encoders and well-curated instruction-tuning datasets for their effectiveness.

There are also several attempts to integrate multiple modalities into one MLLM [9, 29, 57, 99]. As an extension of vision LLM, most previous works align each modality with the LLM using modality-specific encoders and projection modules (middle of Fig. 1). For instance, X-LLM [9] and ChatBridge [99] connect pretrained image, video, and audio encoders with LLMs using separate Q-Former [42] or Perceiver [33] models. However, these modality-specific encoders usually differ in architecture and considerable effort is required to unify them into a single framework. Furthermore, pretrained encoders that deliver reliable performance are usually restricted to widely used modalities such as image, audio, and video. This limitation poses a constraint on MLLMs’ ability to expand to more modalities. Thus, a crucial challenge for MLLMs is *how to build a unified and scalable encoder capable of handling a wide range of modalities*.

We get inspiration from recent works on transferring pretrained transformers to downstream modalities [49, 55, 83, 98]. Lu *et al.* [49] proved that a frozen language-pretrained transformer can achieve strong performance on downstream modalities such as image classification. Meta-Transformer [98] demonstrated that a frozen visual encoder can achieve competitive results across 12 different data modalities. The insights from the works mentioned above suggest that pretrained encoders for each modality may not be necessary. Instead, a well-pretrained transformer may serve as a universal cross-modal encoder.

In this paper, we present **OneLLM**, an MLLM that aligns eight modalities to language using one unified framework. As shown in Fig. 1, OneLLM consists of lightweight modality tokenizers, a universal encoder, a universal projection module (UPM), and an LLM. In contrast to prior works, the encoder and projection module in OneLLM are shared across all modalities. The modality-specific tokenizers, each comprised of only one convolution layer, convert input signals into a sequence of tokens. Additionally, we add learnable modality tokens to enable modality switching and transform input tokens of diverse lengths into tokens of a fixed length.

Training a model of this complexity from scratch poses significant challenges. We start from a vision LLM and align other modalities to the LLM in a progressive way. Specifically, **(i)** we build a vision LLM with pretrained CLIP-ViT [64] as the image encoder, accompanied by several transformer layers as the image projection module, and LLaMA2 [73] as the LLM. After pretraining on massive paired image-text data, the projection module learns to map visual representations into the embedding space of LLM. **(ii)** To align with more modalities, we need a universal en-

coder and projection module. As discussed before, the pretrained CLIP-ViT is possible to serve as a universal encoder. For UPM, we propose to mix multiple image projection experts as a universal X-to-language interface. To increase the model capability, we also design a dynamic router to control the weight of each expert for the given inputs, which turns UPM into soft mixtures-of-experts [63]. Finally, we progressively align more modalities with the LLM based on their data magnitude.

We also curate a large-scale multimodal instruction dataset, including captioning, question answering, and reasoning tasks across eight modalities: image, audio, video, point clouds, depth/normal map, Inertial Measurement Unit (IMU), and functional Magnetic Resonance Imaging (fMRI). By finetuning on this dataset, OneLLM has strong multimodal understanding, reasoning, and instruction-following capabilities. We evaluate OneLLM on multimodal captioning, question answering and reasoning benchmarks where it achieves superior performance than previous specialized models and MLLMs. In conclusion, we summarize our contributions as:

- We propose a unified framework to align multimodal inputs with language. Different from existing works with modality-specific encoders, we show that a unified multimodal encoder, which leverages a pretrained vision-language model and a mixture of projection experts, can serve as a general and scalable component for MLLMs.
- To the best of our knowledge, OneLLM is the first MLLM that integrates eight distinct modalities within a single model. With the unified framework and progressive multimodal alignment pipeline, OneLLM can be easily extended to incorporate more data modalities.
- We curate a large-scale multimodal instruction dataset. OneLLM finetuned on this dataset achieves superior performance on multimodal tasks, outperforming both specialist models and existing MLLMs.

## 2. Related Work

**Large Vision-Language Models.** Large Language Models (LLMs) have gained a lot of attention recently. Therefore, extending LLMs to the vision domain is an emergent and rapidly growing research area. Flamingo [3] is a pioneer to inject frozen visual features into LLM with cross-attention layers, achieving superior performance on a wide range of vision-language tasks. BLIP2 [42] uses a Q-Former to aggregate visual features into a few tokens aligned with LLM. Recently, with the popularity of instruction-following LLMs, vision LLMs have experienced a new explosion. LLaMA-Adapter [20, 97] connects pretrained CLIP [64] and LLaMA [73] with parameter-efficient fine-tuning methods, which can tackle close-set visual question answering and image captioning tasks. Subsequent works [20, 46, 90, 100] propose to train such model on large-scale image-text

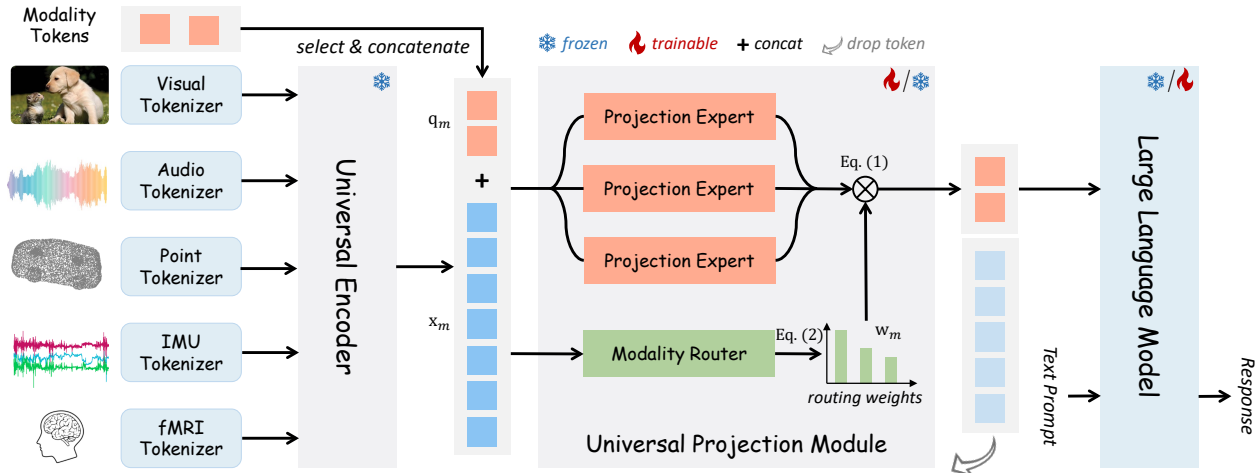


Figure 2. **The Architecture of OneLLM.** OneLLM consists of modality tokenizers, a universal encoder, a universal projection module (UPM) and an LLM. The modality tokenizer is a 2D/1D convolution layer to transform the input signal into a sequence of tokens. For simplicity, we omit video, depth/normal map tokenizers. The universal encoder is a frozen vision-language model (*i.e.* CLIP [64]) to extract high dimensional features. The UPM is composed of several projection experts and modality routers to align the input signal with language. For the alignment stage, we train modality tokenizers and UPM, and keep LLM frozen. For the instruction tuning stage, we only train the LLM and keep other models frozen. In a forward pass of UPM, we concatenate the input and modality tokens as input. Then we only take the modality tokens as a summary of the input signal and feed it into LLM for multimodal understanding.

data, enabling it to complete various instructions about images. Among them, LLaVA [46] adopt a linear layer to directly project visual tokens into LLMs, while MiniGPT-4 [100] and some other works [20, 90] resample visual tokens into fixed-length tokens, reducing the computation cost of LLMs. Our work also belongs to the later branch. We preset learnable tokens for each modality (*i.e.*, modality tokens), which are then used to aggregate input information and generate fixed-length tokens for all modalities.

**Multimodal Large Language Models.** In addition to vision LLMs, recent works proposed to extend LLMs to other modalities, such as audio [23, 94], video [10, 43, 95] and point cloud [26, 87]. These works make it possible to unify multiple modalities into one LLM. X-LLM [9] adopts modality-specific Q-Former [42] and adapters to connect pretrained image, audio and video encoders with LLMs. ChatBridge [99] and AnyMAL [57] follow a similar architecture with X-LLM but adopts Perceiver [33] and linear layers respectively to align modality encoders with LLMs. Meanwhile, PandaGPT [72] and ImageBind-LLM [29] utilize ImageBind [22] as the modality encoder and therefore naturally support multimodal inputs. However, current MLLMs are limited to supporting common modalities such as image, audio and video. It remains unclear how to expand MLLMs to more modalities with a unified framework. In this work, we propose a unified multimodal encoder to align all modalities with language. We show that one universal encoder and projection module can effectively map multimodal inputs to LLM. To our knowledge, OneLLM is first MLLM capable of supporting eight distinct modalities.

**Multimodal-Text Alignment.** Aligning multiple modalities into one joint embedding space is important for cross-modal tasks, which can be divided into two lines of works: discriminative alignment and generative alignment. The most representative work of discriminative alignment is CLIP [64], which utilize contrastive learning to align image and text. Follow-up works extend CLIP to audio-text [28, 80], video-text [51, 85], point-text [96] *etc.* Besides, ImageBind [22] proposes to bind various modalities to images with contrastive learning. On the other hand, generative alignment has attracted much attention in the era of LLM. GIT [77] aligns image and text using a generative image-to-text transformer. BLIP2 [42] proposes generative pretraining to connect frozen vision encoder and LLM. VALOR [11] and VAST [12] extends the training paradigm of BLIP2 to more modalities such as audio and video. Our work also belongs to generative alignment. In contrast to prior works, we directly align mutlimodal inputs to LLMs, thus getting rid of the stage of training modality encoders.

### 3. Method

In this section, we will first introduce the architecture of OneLLM (Sec. 3.1) and then present our two training phases: progressive multimodal alignment (Sec. 3.2) and unified multimodal instruction tuning (Sec. 3.3).

#### 3.1. Model Architecture

Fig. 2 depicts the four main components of OneLLM: modality-specific tokenizers, a universal encoder, a universal projection module (UPM) and an LLM. Detailed de-

scriptions are presented in the following sections.

**Lightweight Modality Tokenizers.** The modality tokenizer is to transform the input signal into a sequence of tokens, thereby a transformer-based encoder can process these tokens. We denote the input tokens as  $\mathbf{x} \in \mathbb{R}^{L \times D}$ , where  $L$  is the sequence length and  $D$  is the token dimension. Considering the variations inherent to different data modalities, we design a separate tokenizer for each modality. For visual inputs with 2D position information such as image and video, we directly utilize a single 2D convolution layer as the tokenizer. For other modalities, we transform the input into a 2D or 1D sequence, which is then tokenized using a 2D/1D convolution layer. For example, we transform audio signals into 2D spectrogram and sample a subset of point clouds with 2D geometric prior. Due to space limit, please refer to the appendix for more details.

**Universal Encoder.** As discussed in Sec. 1, frozen pretrained transformers demonstrate strong modality transfer capability [49, 98]. Therefore, we leverage pretrained vision-language models as the universal encoder for all modalities. Vision-language models, when trained on extensive image-text data, typically learn robust alignment between vision and language, so they can be easily transferred to other modalities. In OneLLM, we use CLIP-ViT [64] as a universal computation engine. Following previous works [49, 98], we keep the parameters of CLIP-ViT frozen during training. Note that for video signals, we will feed all video frames into the encoder in parallel and perform token-wise averaging between frames to speed up training. Other strategies, such as token concatenation, may further enhance the model’s video understanding capability.

**Universal Projection Module.** In contrast to existing works with modality-specific projection, we propose a Universal Projection Module (UPM) to project any modality into LLM’s embedding space. As shown in Fig. 2, UPM consists of  $K$  projection experts  $\{P_k\}$ , where each expert is a stack of transformer layers pretrained on image-text data (will discuss in Sec. 3.2). Although one expert can also realize any modality-to-LLM projection, our empirical findings suggest that multiple experts are more effective and scalable. When scaling to more modalities, we only need to add a few parallel experts.

To integrate multiple experts into one module, we propose a dynamic *modality router*  $R$  to control each expert’s contribution and increase the model capacity. The router  $R$  is structured as a straightforward Multi-Layer Perception that receives input tokens and calculates the routing weights for each expert, *i.e.*, a soft router [63]. We will also discuss other types of router in Sec. 4.3, such as constant router and sparse router. Besides, we add learnable modality tokens  $\{\mathbf{q}_m\}_{m \in \mathcal{M}}$  to switch between modalities, where  $\mathcal{M}$  is the set of modalities and  $\mathbf{q}_m \in \mathbb{R}^{N \times D}$  contains  $N$  tokens of dimension  $D$ . In a forward pass for modality  $m$ , we feed the concatenation of input tokens  $\mathbf{x}_m \in \mathbb{R}^{L \times D}$  and modality

tokens  $\mathbf{q}_m$  into UPM:

$$[\bar{\mathbf{q}}_m, \bar{\mathbf{x}}_m] = \text{UPM}([\mathbf{q}_m, \mathbf{x}_m]) = \sum_{k=1}^K \mathbf{w}_m \cdot P_k([\mathbf{q}_m, \mathbf{x}_m]), \quad (1)$$

$$\mathbf{w}_m = \sigma \circ R_m([\mathbf{q}_m, \mathbf{x}_m]), \quad (2)$$

where  $\mathbf{w}_m \in \mathbb{R}^{N \times K}$  is the routing weight and the SoftMax function  $\sigma$  is to ensure  $\sum_{k=1}^K \mathbf{w}_m \cdot k = 1$ . For any modality  $m$ , we *only* extract the projected modality tokens  $\bar{\mathbf{q}}_m$  as a summary of input signals, transforming  $\mathbf{x}_m$  from varying lengths into uniform, fixed-length tokens.

**LLM.** We employ the open-source LLaMA2 [74] as the LLM in our framework. The input to LLM includes projected modality tokens  $\bar{\mathbf{q}}_m$  and the text prompt after word embedding. Note we always put modality tokens at the beginning of the input sequence for simplicity. Then LLM is asked to generate appropriate response conditioned on modality tokens and text prompt.

### 3.2. Progressive Multimodal Alignment

Image-text alignment has been well investigated in previous works [20, 47, 100]. Therefore, a naive approach for multimodal alignment is to jointly train the model on multimodal-text data. However, training models directly on multimodal data can lead to biased representations between modalities due to the imbalance of data scale. Here we propose to train an image-to-text model as initialization and progressively ground other modalities into LLM.

**Image-Text Alignment.** We begin with a basic vision LLM framework, comprising an image tokenizer, a pretrained CLIP-ViT, an image projection module  $P_I$  and an LLM. Considering that image-text data is relatively abundant compared to other modalities, we first train the model on image-text data to well align CLIP-ViT and LLM, *i.e.*, learning a good image-to-text projection module. The pretrained  $P_I$  not only serves as a bridge connecting images and language, but also provides a good initialization for multimodal-text alignment. Then we build UPM by mixing multiple pretrained  $P_I$ :  $\text{UPM} = \{P_k\} = \{\text{Init}(P_I)\}$ , where  $\text{Init}$  is weight initialization, which effectively reduces the cost of aligning other modalities to language.

**Multimodal-Text Alignment.** We formulate multimodal-text alignment as a continual learning process [75]. At timestamp  $t$ , we have trained the model on a set of modalities  $\mathcal{M}_1 \cup \mathcal{M}_2 \cdots \mathcal{M}_{t-1}$ , and the current training data is from  $\mathcal{M}_t$ . To prevent catastrophic forgetting, we will sample evenly from both previous trained data and current data. In our case, we divide multimodal-text alignment into multiple training stages based on their data magnitude: stage I (image), stage II (video, audio and point cloud) and stage III (depth/normal map, IMU and fMRI). If we want to support new modalities, we can repeat the training episode, *i.e.*, sampling a similar amount of data from previous modalities and jointly training the model with the current modalities.

**Multimodal-Text Dataset.** We collect X-text pairs for each modality. The image-text pairs include LAION-400M [67] and LAION-COCO [66]. The training data for video, audio and point clouds are WebVid-2.5M [7], WavCaps [54] and Cap3D [52], respectively. Since there is no large-scale depth/normal map-text data, we use pretrained DPT model [18, 65] to generate depth/normal map. The source images and text and from CC3M [69]. For IMU-text pairs, we use the IMU sensor data of Ego4D [25]. For fMRI-text pairs, we use fMRI signals from the NSD [4] dataset and take the captions associated with the visual stimuli as text annotations. Note that the input to LLM is the concatenation of modality tokens and caption tokens. We do not add system prompts at this stage to reduce the number of tokens and speed up training.

### 3.3. Unified Multimodal Instruction Tuning

After multimodal-text alignment, OneLLM becomes a multimodal captioning model which can generate a short description for any input. To fully unleash OneLLM’s multimodal understanding and reasoning capabilities, we curate a large-scale multimodal instruction tuning dataset to further finetune OneLLM.

**Multimodal Instruction Tuning Dataset.** We collect instruction tuning (IT) dataset for each modality. Following previous works [14, 46], the *image* IT datasets are sampled from the following datasets: LLaVA-150K [47], COCO Caption [13], VQAv2 [24], GQA [32], OKVQA [53], A-OKVQA [68], OCRVQA [56], RefCOCO [34] and Visual Genome [36]. The *video* IT datasets include MSRVT-Cap [86], MSRVT-QA [84] and video instruction data from [99]. The *audio* IT datasets include AudioCaps [35] and audio conversation data from [99]. The *point cloud* IT dataset is a 70K point cloud description, conversation and reasoning dataset from [87]. The *depth/normal map* IT datasets are generated from image IT datasets: we random sample 50K visual instruction data from LLaVA-150K and generate depth/normal map using DPT model [18]. For *IMU* and *fMRI* IT datasets, we also random sample a subset from Ego4D [25] and NSD [4], respectively. Finally, our multimodal IT datasets have about **2M** items, covering multiple tasks such as detailed description/reasoning, conversation, short question answering and captioning.

**Prompt Design.** Given the diverse modalities and tasks within our multimodal IT datasets, we carefully design the prompts to avoid conflicts between them. (a) When utilizing IT datasets generated by GPT4 (e.g., LLaVA-150K), we adopt the original prompts provided by these datasets. (b) For captioning tasks, we employ the prompt: *Provide a one-sentence caption for the provided {modal}*. (c) For open-ended question answering tasks, we enhance the question with *Answer the question using a single word or phrase*. (d) For question answering tasks with options, the prompt

is:  $\{Question\} \{Options\}$  Answer with the option’s letter from the given choices directly. (e) For IMU and fMRI datasets, we apply prompt such as *Describe the motion* and *Describe this scene based on fMRI data*. Despite using these fixed prompts, our experiments indicate that OneLLM is capable of generalizing to open-ended prompts during inference. For detailed prompts on each task and modality, please check out the appendix.

In the instruction tuning stage, we organize the input sequence as:  $\{\bar{q}, Sys, [Inst_t, Ans_t]_{t=1}^T\}$  where  $\bar{q}$  is the modality tokens, *Sys* is the system prompt,  $[Inst_t, Ans_t]$  corresponds to the *t*-th instruction-answer pair in a conversation. Note that for multimodal inputs involving multiple modalities, such as audio-visual tasks [40], we position all modality tokens at the start of the input sequence.

We fully finetune the LLM and keep rest parameters frozen. Although recent works often employ parameter-efficient methods [31], we empirically show that the full finetuning approach more effectively harnesses the multimodal capabilities of OneLLM, particularly with the utilization of smaller LLMs (e.g., LLaMA2-7B).

## 4. Experiment

### 4.1. Implementation Details

**Architecture.** The universal encoder is CLIP VIT Large pretrained on LAION [67]. The LLM is LLaMA2-7B [74]. The UPM has  $K=3$  projection experts, where each expert has eight Transformer blocks and 88M parameters. The size of modality tokens for each modality is  $\mathbb{R}^{30 \times 1024}$ .

**Training Details.** We use AdamW optimizer with  $\beta_1=0.9$ ,  $\beta_2=0.95$  and weight decay of 0.1. We apply a linear learning rate warmup during the first 2K iterations. For stage I, we train OneLLM on 16 A100 GPUs for 200K iterations. The effective batch size (using gradient accumulation) is 5120. The maximum learning rate is  $5e-5$ . For stage II (resp. III), we train OneLLM on 8 GPUs for 200K (resp. 100K) with an effective batch size of 1080 and maximum learning rate of  $1e-5$ . In the instruction tuning stage, we train OneLLM on 8 GPUs for 1 epoch (96K) with an effective batch size of 512 and maximum learning rate of  $2e-5$ .

### 4.2. Quantitative Evaluation

We evaluate OneLLM on multimodal tasks and put evaluation details to the appendix.

**Image-Text Evaluation.** In Tab. 1, we evaluate OneLLM on visual question answering (VQA), image captioning and recent multimodal benchmarks. For VQA tasks, OneLLM-7B outperforms other MLLMs such as ChatBridge-13B [99] and AnyMAL-13B [57] by a large margin. Our 7B model is even better than AnyMAL with 70B parameters. For image captioning tasks, OneLLM-7B is on-par with ChatBridge-13B. Although OneLLM is not specifi-

Model	LLM	VQA						Image Caption		MM Benchmark			
		GQA	VQAv2	OKVQA	TVQA	SQA	Vizwiz	NoCaps	Flickr	MME	MMB	MMVet	SEED
vision specialized LLM													
Flamingo-9B [3]	Chinchilla-7B	-	51.8	44.7	30.1	-	28.8	-	61.5	-	-	-	-
Flamingo-80B [3]	Chinchilla-70B	-	56.3	50.6	31.8	-	31.6	-	67.2	-	-	-	-
BLIP-2 [42]	Vicuna-7B	-	-	-	40.1	53.8	-	107.5	74.9	-	-	-	-
BLIP-2 [42]	Vicuna-13B	41.0	41.0	-	42.5	61	19.6	103.9	71.6	1293.8	-	22.4	-
InstructBLIP [14]	Vicuna-7B	49.2	-	-	50.1	60.5	34.5	123.1	82.4	-	36	26.2	-
InstructBLIP [14]	Vicuna-13B	49.5	-	-	50.7	63.1	34.3	121.9	82.8	1212.8	-	25.6	-
IDEFICS-9B [37]	LLaMA-7B	38.4	50.9	38.4	25.9	-	35.5	-	27.3	-	48.2	-	-
IDEFICS-80B [37]	LLaMA-65B	45.2	60.0	45.2	30.9	-	36.0	-	53.7	-	54.5	-	-
LLaMA-Ad.v2 [20]	LLaMA-7B	43.9	-	55.9	43.8	54.2	-	42.7	30.5	972.7	38.9	31.4	32.7
Qwen-VL [6]	Qwen-7B	57.5	78.2	56.6	61.5	68.2	38.9	120.2	81.0	1487.5	60.6	-	58.2
LLaVA-v1.5 [46]	Vicuna-7B	62.0	78.5	-	58.2	66.8	50.0	-	-	1510.7	64.3	30.5	58.6
multimodal generalist LLM													
ImageBind-LLM [29]	LLaMA-7B	41.1	-	-	24.0	51.4	-	29.6	23.5	775.7	-	-	-
ChatBridge-13B [99]	Vicuna-13B	<u>41.8</u>	-	<u>45.2</u>	-	-	-	<u>115.7</u>	<b>82.5</b>	-	-	-	-
AnyMAL-13B [57]	LLaMA2-13B	-	59.6	33.1	24.7	52.7	24.4	-	-	-	-	-	-
AnyMAL-70B [57]	LLaMA2-70B	-	<u>64.2</u>	42.6	<u>32.9</u>	<b>70.8</b>	<u>33.8</u>	-	-	-	-	-	-
<b>OneLLM-7B (Ours)</b>	<b>LLaMA2-7B</b>	<b>59.5</b>	<b>71.6</b>	<b>58.9</b>	<b>34.0</b>	<u>63.4</u>	<b>45.9</b>	<b>115.9</b>	<u>78.6</u>	1392.0	60.0	29.1	<b>61.2</b>

Table 1. **Evaluation on 12 Image-Text Benchmarks**, including 6 VQA tasks (GQA [32], VQAv2 [24], OKVQA [53], TextVQA (TVQA) [70], ScienceQA (SQA) [50] and Vizwiz [27]), 2 image captioning tasks (Nocaps [1] and Flickr30K [62]), and 4 multimodal benchmarks (MME [19], MM Bench (MMB) [48], MMVet [93] and SEED [39]). The LLMs are Chinchilla [30], Vicuna [76], Qwen [5], LLaMA [73] and LLaMA2 [74]. The evaluation metrics for VQA and captioning tasks are accuracy and CIDEr, respectively. The results in **bold** and underline are the best and second-best results, respectively. -: Not reported result.

Model	0-shot	NextQA Acc.	How2QA Acc.	MSVD Acc.	VATEX CIDEr
HGQA [82]	✗	51.8	-	41.2	-
JustAsk [88]	✗	52.3	84.4	46.3	-
VALOR [11]	✗	-	-	60.0	95.1
SeViLA [92]	✗	73.8	83.6	-	-
FrozenBiLM [89]	✓	-	58.4	33.8	-
InternVideo [79]	✓	<u>49.1</u>	<u>62.2</u>	<u>55.5</u>	-
ChatBridge-13B [99]	✓	-	-	45.3	<b>48.9</b>
AnyMAL-13B [57]	✓	47.9	59.6	-	-
<b>OneLLM-7B (Ours)</b>	✓	<b>57.3</b>	<b>65.7</b>	<b>56.5</b>	<u>43.8</u>

Table 2. **Evaluation on Video-Text Tasks**, including video question answering (NextQA [81], How2QA [44] and MSVD [84]) and video captioning tasks (VATEX [78]). Acc.: Accuracy.

Model	0-shot	Clotho Caption CIDEr	Clotho Caption SPIDEr	Clotho AQA Acc.
FeatureCut [91]	✗	43.6	27.9	-
Wavcaps [54]	✗	48.8	31.0	-
MWAFM [41]	✗	-	-	22.2
Pengi [16]	✗	-	27.1	64.5
LTU-7B [23]	✓	-	11.9	-
ChatBridge-13B [99]	✓	26.2	-	-
<b>OneLLM-7B (Ours)</b>	✓	<b>29.1</b>	<b>19.5</b>	<b>57.9</b>

Table 3. **Evaluation on Audio-Text Tasks**, including audio captioning on Clotho Caption [17] and audio question answering on Clotho AQA [45].

cally designed for vision tasks, our results demonstrate that OneLLM can also reach the leading level in vision specialized LLMs, and the gap between MLLMs and vision LLMs has further narrowed.

**Video-Text Evaluation.** As shown in Tab. 2, we evaluate OneLLM on video QA and captioning tasks. Our model

Model	0-shot	MUSIC-AVQA Acc.	VALOR CIDEr	AVSD CIDEr
MAVQA [40]	✗	71.5	-	-
VALOR [11]	✗	78.9	61.5	-
VAST [12]	✗	80.7	62.2	-
FA+HRED [59]	✗	-	-	84.3
MTN [38]	✗	-	-	98.5
COST [61]	✗	-	-	108.5
ChatBridge-13B [99]	✓	<u>43.0</u>	<u>24.7</u>	<b>75.4</b>
<b>OneLLM-7B (Ours)</b>	✓	<b>47.6</b>	<b>29.2</b>	<u>74.5</u>

Table 4. **Evaluation on Audio-Video-Text Tasks**, including audio-visual question answering on MUSIC-AVQA [40] and audio-visual captioning on VALOR-32K [11] and dialog completion on AVSD [2].

Model	BLEU-1	Captioning ROUGE-L	Captioning METEOR	Classification GPT4-Acc.
InstructBLIP-7B [14]	11.2	13.9	14.9	38.5
InstructBLIP-13B [14]	12.6	15.0	16.0	35.5
PointLLM-7B [87]	8.0	11.1	15.2	<b>47.5</b>
PointLLM-13B [87]	9.7	12.8	15.3	45.0
<b>OneLLM-7B (Ours)</b>	<b>42.2</b>	<b>45.3</b>	<b>20.3</b>	44.5

Table 5. **Evaluation on Point Cloud-Text Tasks.** The evaluation dataset is from Objaverse [15], following the data split in PointLLM [87]. InstructBLIP takes single-view image as input, while PointLLM and OneLLM take point cloud as input. GPT4-Acc.: GPT4 as the accuracy evaluator [87].

outperforms both MLLMs [57, 99] and video-specific models [79, 89] in video QA tasks. Notably, our training datasets do not include video QA data like NextQA [81] and How2QA [44], which are video QA tasks that provide answer options. However, our model’s training on similar VQA datasets (*e.g.*, A-OKVQA [68]) has evidently en-

Model	0-shot	NYUv2 Acc.	SUN RGB-D Acc.
ImageBind [22]	✗	54.0	35.1
Omnivore [21]	✗	76.7	64.9
Random	✓	10.0	5.26
CLIP ViT-H* [64]	✓	41.9	25.4
<b>OneLLM-N (Ours)</b>	✓	<u>46.5</u>	21.2
<b>OneLLM-D (Ours)</b>	✓	<b>50.9</b>	<b>29.0</b>

Table 6. **Evaluation on Scene Classification Tasks Using Depth / Normal Map.** OneLLM-N/D: OneLLM with Depth / Normal map inputs. Note that NYUv2 [58] and SUN RGB-D [71] only have depth maps, we adopt pretrained DPT model [18] to generate normal maps. \*: The input to CLIP is depth rendered grayscale image. ImageBind is trained on image-depth pairs of SUN RGB-D and therefore is not zero-shot.

hanced its emergent cross-modal capabilities, contributing to the improved performance in video QA tasks.

**Audio-Text Evaluation.** We evaluate OnLLM on audio captioning and QA tasks. In Tab. 3, we outperforms both ChatBridge and LTU [23] on Clotho Caption [17]. Notably, our zero-shot result on Clotho AQA [45] is on par with fully finetuned Pengi [16]. Similar to our conclusion on video QA, we believe that the captioning task requires more dataset-specific training, while the QA task may be a more accurate measure of the model’s inherent zero-shot understanding capabilities.

**Audio-Video-Text Evaluation.** We evaluate OneLLM on audio-video-text tasks, such as QA (MUSIC AVQA [40]), captioning (VALOR-32K [11]) and dialog completion (AVSD [2]) based on the video and background audio. As shown in Tab. 4, OneLLM-7B surpasses ChatBridge-13B on all three datasets. Note that ChatBridge was trained on an audio-visual dataset [11], while OneLLM has not been trained on any audio-visual datasets. Since all modalities in OneLLM are well aligned with language, we can directly input video and audio signals to OneLLM during inference.

**Point Cloud-Text Evaluation.** In Tab. 5, We evaluate OneLLM on point cloud captioning and classification tasks. OneLLM can achieve excellent captioning results due to our carefully designed instruction prompts for switching between tasks (Sec. 3.3), while InstructBLIP [14] and PointLLM [87] struggle to generate short and accurate captions. On the classification task, OneLLM can also achieve comparable results to PointLLM.

**Depth/Normal Map-Text Evaluation.** Since there are currently no QA and captioning tasks using depth/normal maps, we evaluate OneLLM on two scene classification datasets [58, 71]. The performance, as displayed in Tab. 6, reveals that OneLLM achieves superior zero-shot classification accuracy compared to CLIP. These results affirm that OneLLM trained on synthetic depth/normal map data can adapt to real world scenarios.

Task	NoCaps	VQAv2	ClothoQA	MSVDQA
<b>(a) Training Mode</b>				
Separate	115.6(-0.2)	71.9(+0.3)	37.8(-19.6)	31.0(-25.8)
Joint	115.8	71.6	57.4	56.8
<b>(b) Weight Initialization</b>				
Random Init.	98.8(-17.0)	65.6(-6.0)	57.6(+0.2)	53.1(-3.7)
Image Init.	115.8	71.6	57.4	56.8
<b>(c) Number of Experts (Parameters)</b>				
1 (88M)	108.7(-7.1)	66.9(-4.7)	58.2(+0.8)	53.3(-3.5)
3 (264M)	115.8	71.6	57.4	56.8
5 (440M)	114.6	71.7	58.2	56.7
7 (616M)	114.9	71.6	58.8	56.0
<b>(d) Router Type</b>				
Constant Router	109.8(-6.0)	67.7(-3.9)	56.2(-1.2)	55.3(-1.5)
Sparse Router	112.8(-3.0)	71.1(-0.5)	56.7(-0.7)	55.7(-1.1)
Soft Router	115.8	71.6	57.4	56.8

Table 7. **Ablation Experiments.** We choose three modalities (image, audio, video) and four datasets (NoCaps [1], VQAv2 [24], ClothoQA [45] and MSVDQA [84]) for evaluation. The row with gray background is our default setting.

**IMU-Text and fMRI-Text Evaluation.** Since IMU/fMRI to text generation are seldom explored in previous literature, we solely report our results on IMU/fMRI captioning. For IMU captioning on Ego4D [25], we evaluate OneLLM on a held-out subset with 2000 items. The CIDEr and ROUGE-L score are 24.9 and 19.5, respectively. For fMRI captioning on NSD [4], we evaluate OneLLM on its testing set, where OneLLM achieves 31.7 CIDEr and 25.1 ROUGE-L.

### 4.3. Ablation Experiments

In this section, we explored some key designs of OneLLM. Our ablation experiments are conducted on three modalities: image, audio and video, except for studies on the number of experts. Other settings remain unchanged.

**Separate Training vs. Joint Training.** An important question for MLLMs is whether a jointly trained MLLM is better than modality-specific MLLM? To address this, we compare the performance of separately trained MLLMs against a jointly trained MLLM in Tab. 7 (a). In separate training, the model can only access its own data; in joint training, the model is jointly trained on all data. On two image-text tasks NoCaps and VQAv2, we can see that separately and jointly trained models achieve comparable results; While separately trained audio and video models are much worse than the jointly trained model on ClothoQA and MSVDQA, respectively. This suggest that joint training substantially benefits data-scarce modalities (e.g., audio and video), by allowing for the transfer of learned knowledge (e.g., question answering) across modalities.

**Image Alignment Benefits Multimodal Alignment.** Tab. 7 (b) demonstrates that OneLLM with image-text alignment can help multimodal-text alignment. If we di-

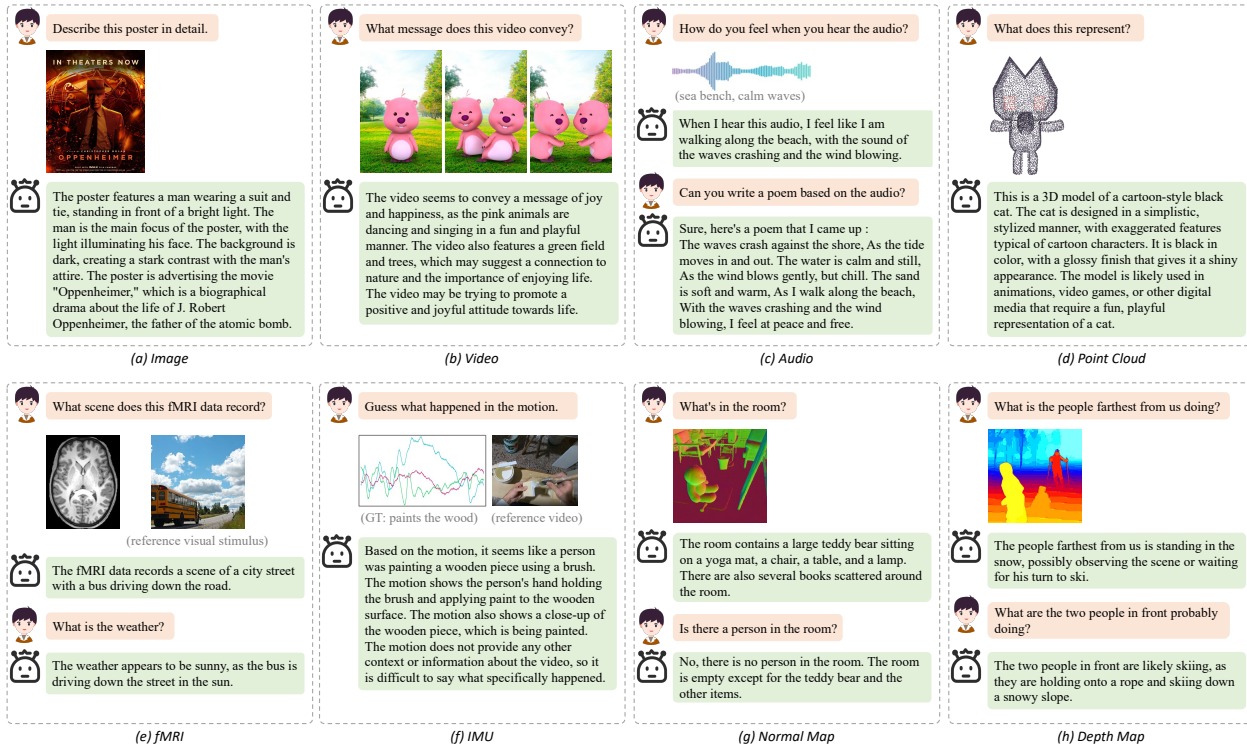


Figure 3. **Qualitative Results on Eight Modalities.** All demo inputs are from the web or the testing set of corresponding modalities.

rectly align all modalities with text using a random initialized model (*i.e.* UPM), the performance on image and video will drop significantly. Instead, OneLLM with image-text pretraining can better balance different modalities.

**Number of Projection Experts.** The number of projection experts in UPM is closely related to the number of modalities that OneLLM can accommodate. As shown in Tab. 7, OneLLM with three projection experts is enough to hold all modalities. Increasing the number of experts does not bring about the desired improvement, while the results with one expert is also not satisfactory.

**Router Type.** The modality router is to link multiple projection experts into a single module. Here we discuss three types of router: constant router, sparse router and the default soft router. (a) Constant router links  $K$  experts with a constant number  $1/K$ . The output of constant router is  $\sum_{k=1}^K \frac{1}{K} \cdot P_k(\mathbf{x})$ . (b) Sparse router only selects one expert with the maximum routing weight. The output is  $w_{k^*} P_{k^*}(\mathbf{x})$  where  $k^* = \arg \max_k w_k$ . As shown in Tab. 7 (d), soft router outperforms other two routers, indicating its effectiveness for dynamic routing of multimodal signals.

#### 4.4. Qualitative Analysis

Fig. 3 gives some qualitative results of OneLLM on eight modalities. We show OneLLM can (a) understand both visual and textual content in images, (b) leverage temporal information in videos, (c) do creative writing based on audio

content, (d) understand the details of 3D shapes, (e) analyze visual scenes recorded in fMRI data, (f) guess the person’s action based on motion data, and (g)-(h) scene understanding using depth/normal map. Due to space limit, we put more qualitative results to the appendix.

## 5. Conclusion

In this work, we introduce OneLLM, which aligns eight modalities with language using a unified framework. We first train a basic vision LLM and then design a multimodal framework with a universal encoder, a UPM and an LLM. By a progressive alignment pipeline, OneLLM can handle multimodal inputs with a single model. Besides, we also curate a large-scale multimodal instruction dataset and evaluate OneLLM on 25 diverse benchmarks, showing its excellent performance.

**Limitation and Future Work.** Our work faces two primary challenges: (i) The absence of large-scale, high-quality datasets for modalities beyond image. (ii) Fine-grained multimodal understanding in high-resolution images, long sequences video and audio *etc.* In the future, we will collect high-quality datasets and design new encoders to realize fine-grained multimodal understanding [8].

**Acknowledgements.** This work is partially supported by the National Natural Science Foundation of China (Grant No. 8326014), CUHK Direct Grants (Grant No. 4055190), National Key R&D Program of China (2022ZD0160201) and Shanghai Artificial Intelligence Laboratory.



## References

- [1] Harsh Agrawal, Karan Desai, Yufei Wang, Xinlei Chen, Rishabh Jain, Mark Johnson, Dhruv Batra, Devi Parikh, Stefan Lee, and Peter Anderson. nocraps: novel object captioning at scale. In *ICCV*, 2019. 6, 7
- [2] Huda Alamri, Vincent Cartillier, Abhishek Das, Jue Wang, Anoop Cherian, Irfan Essa, Dhruv Batra, Tim K. Marks, Chiori Hori, Peter Anderson, Stefan Lee, and Devi Parikh. Audio-visual scene-aware dialog. In *CVPR*, 2019. 6, 7
- [3] 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*, 35:23716–23736, 2022. 1, 2, 6
- [4] Emily J Allen, Ghislain St-Yves, Yihan Wu, Jesse L Breedlove, Jacob S Prince, Logan T Dowdle, Matthias Nau, Brad Caron, Franco Pestilli, Ian Charest, et al. A massive 7t fmri dataset to bridge cognitive neuroscience and artificial intelligence. *Nature neuroscience*, 25(1):116–126, 2022. 5, 7
- [5] Jinze Bai, Shuai Bai, Yunfei Chu, Zeyu Cui, Kai Dang, Xiaodong Deng, Yang Fan, Wenbin Ge, Yu Han, Fei Huang, et al. Qwen technical report. *arXiv preprint arXiv:2309.16609*, 2023. 6
- [6] Jinze Bai, Shuai Bai, Shusheng Yang, Shijie Wang, Sinan Tan, Peng Wang, Junyang Lin, Chang Zhou, and Jingren Zhou. Qwen-vl: A frontier large vision-language model with versatile abilities. *arXiv preprint arXiv:2308.12966*, 2023. 6
- [7] Max Bain, Arsha Nagrani, Gül Varol, and Andrew Zisserman. Frozen in time: A joint video and image encoder for end-to-end retrieval. In *ICCV*, pages 1728–1738, 2021. 5
- [8] Rohan Bavishi, Erich Elsen, Curtis Hawthorne, Maxwell Nye, Augustus Odena, Arushi Somani, and Saġnak Taşırılar. Introducing our multimodal models, 2023. 8
- [9] Feilong Chen, Minglun Han, Haozhi Zhao, Qingyang Zhang, Jing Shi, Shuang Xu, and Bo Xu. X-llm: Bootstrapping advanced large language models by treating multi-modalities as foreign languages. *arXiv preprint arXiv:2305.04160*, 2023. 2, 3
- [10] Guo Chen, Yin-Dong Zheng, Jiahao Wang, Jilan Xu, Yifei Huang, Junting Pan, Yi Wang, Yali Wang, Yu Qiao, Tong Lu, et al. Videollm: Modeling video sequence with large language models. *arXiv preprint arXiv:2305.13292*, 2023. 1, 3
- [11] Sihan Chen, Xingjian He, Longteng Guo, Xinxin Zhu, Weining Wang, Jinhui Tang, and Jing Liu. Valor: Vision-audio-language omni-perception pretraining model and dataset. *arXiv preprint arXiv:2304.08345*, 2023. 3, 6, 7
- [12] Sihan Chen, Handong Li, Qunbo Wang, Zijia Zhao, Mingzhen Sun, Xinxin Zhu, and Jing Liu. Vast: A vision-audio-subtitle-text omni-modality foundation model and dataset. *arXiv preprint arXiv:2305.18500*, 2023. 3, 6
- [13] Xinlei Chen, Hao Fang, Tsung-Yi Lin, Ramakrishna Vedantam, Saurabh Gupta, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco captions: Data collection and evaluation server. *arXiv preprint arXiv:1504.00325*, 2015. 5
- [14] Wenliang Dai, Junnan Li, Dongxu Li, Anthony Meng Huat Tiong, Junqi Zhao, Weisheng Wang, Boyang Li, Pascale Fung, and Steven Hoi. Instructblip: Towards general-purpose vision-language models with instruction tuning, 2023. 5, 6, 7
- [15] Matt Deitke, Dustin Schwenk, Jordi Salvador, Luca Weihs, Oscar Michel, Eli VanderBilt, Ludwig Schmidt, Kiana Ehsani, Aniruddha Kembhavi, and Ali Farhadi. Objaverse: A universe of annotated 3d objects. In *CVPR*, pages 13142–13153, 2023. 6
- [16] Soham Deshmukh, Benjamin Elizalde, Rita Singh, and Huaming Wang. Pengi: An audio language model for audio tasks. *arXiv preprint arXiv:2305.11834*, 2023. 6, 7
- [17] Konstantinos Drossos, Samuel Lipping, and Tuomas Virtanen. Clotho: An audio captioning dataset. In *ICASSP*, pages 736–740. IEEE, 2020. 6, 7
- [18] Ainaz Eftekhari, Alexander Sax, Jitendra Malik, and Amir Zamir. Omnidata: A scalable pipeline for making multi-task mid-level vision datasets from 3d scans. In *ICCV*, pages 10786–10796, 2021. 5, 7
- [19] Chaoyou Fu, Peixian Chen, Yunhang Shen, Yulei Qin, Mengdan Zhang, Xu Lin, Zhenyu Qiu, Wei Lin, Jinrui Yang, Xiaowu Zheng, et al. Mme: A comprehensive evaluation benchmark for multimodal large language models. *arXiv preprint arXiv:2306.13394*, 2023. 1, 6
- [20] Peng Gao, Jiaming Han, Renrui Zhang, Ziyi Lin, Shijie Geng, Aojun Zhou, Wei Zhang, Pan Lu, Conghui He, Xiangyu Yue, et al. Llama-adapter v2: Parameter-efficient visual instruction model. *arXiv preprint arXiv:2304.15010*, 2023. 2, 3, 4, 6
- [21] Rohit Girdhar, Mannat Singh, Nikhila Ravi, Laurens van der Maaten, Armand Joulin, and Ishan Misra. Omnivore: A single model for many visual modalities. In *CVPR*, pages 16102–16112, 2022. 7
- [22] Rohit Girdhar, Alaaeldin El-Nouby, Zhuang Liu, Mannat Singh, Kalyan Vasudev Alwala, Armand Joulin, and Ishan Misra. Imagebind: One embedding space to bind them all. In *CVPR*, pages 15180–15190, 2023. 3, 7
- [23] Yuan Gong, Hongyin Luo, Alexander H Liu, Leonid Karlinsky, and James Glass. Listen, think, and understand. *arXiv preprint arXiv:2305.10790*, 2023. 1, 3, 6, 7
- [24] Yash Goyal, Tejas Khot, Douglas Summers-Stay, Dhruv Batra, and Devi Parikh. Making the v in vqa matter: Elevating the role of image understanding in visual question answering. In *CVPR*, pages 6904–6913, 2017. 5, 6, 7
- [25] Kristen Grauman, Andrew Westbury, Eugene Byrne, Zachary Chavis, Antonino Furnari, Rohit Girdhar, Jackson Hamburger, Hao Jiang, Miao Liu, Xingyu Liu, et al. Ego4d: Around the world in 3,000 hours of egocentric video. In *CVPR*, pages 18995–19012, 2022. 5, 7
- [26] Ziyu Guo, Renrui Zhang, Xiangyang Zhu, Yiwen Tang, Xianzheng Ma, Jiaming Han, Kexin Chen, Peng Gao, Xianzhi Li, Hongsheng Li, et al. Point-bind & point-llm: Aligning point cloud with multi-modality for 3d understanding, generation, and instruction following. *arXiv preprint arXiv:2309.00615*, 2023. 1, 3
- [27] Danna Gurari, Qing Li, Abigale J Stangl, Anhong Guo, Chi Lin, Kristen Grauman, Jiebo Luo, and Jeffrey P Bigham.

- Vizwiz grand challenge: Answering visual questions from blind people. In *CVPR*, pages 3608–3617, 2018. 6
- [28] Andrey Guzhov, Federico Raue, Jörn Hees, and Andreas Dengel. Audioclip: Extending clip to image, text and audio. In *ICASSP*, pages 976–980. IEEE, 2022. 3
- [29] 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, 6
- [30] Jordan Hoffmann, Sebastian Borgeaud, Arthur Mensch, Elena Buchatskaya, Trevor Cai, Eliza Rutherford, Diego de Las Casas, Lisa Anne Hendricks, Johannes Welbl, Aidan Clark, et al. Training compute-optimal large language models. *arXiv preprint arXiv:2203.15556*, 2022. 6
- [31] 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. *arXiv preprint arXiv:2106.09685*, 2021. 5
- [32] Drew A Hudson and Christopher D Manning. Gqa: A new dataset for real-world visual reasoning and compositional question answering. In *CVPR*, 2019. 5, 6
- [33] Andrew Jaegle, Felix Gimeno, Andy Brock, Oriol Vinyals, Andrew Zisserman, and Joao Carreira. Perceiver: General perception with iterative attention. pages 4651–4664. PMLR, 2021. 2, 3
- [34] Sahar Kazemzadeh, Vicente Ordonez, Mark Matten, and Tamara Berg. Referitgame: Referring to objects in photographs of natural scenes. In *EMNLP*, pages 787–798, 2014. 5
- [35] Chris Dongjoo Kim, Byeongchang Kim, Hyunmin Lee, and Gunhee Kim. Audiocaps: Generating captions for audios in the wild. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 119–132, 2019. 5
- [36] Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, David A Shamma, et al. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *IJCV*, 123:32–73, 2017. 5
- [37] Hugo Laurençon, Lucile Saulnier, Léo Tronchon, Stas Bekman, Amanpreet Singh, Anton Lozhkov, Thomas Wang, Siddharth Karamcheti, Alexander M Rush, Douwe Kiela, et al. Obelisc: An open web-scale filtered dataset of interleaved image-text documents. *arXiv preprint arXiv:2306.16527*, 2023. 6
- [38] Hung Le, Doyen Sahoo, Nancy F Chen, and Steven CH Hoi. Multimodal transformer networks for end-to-end video-grounded dialogue systems. *arXiv preprint arXiv:1907.01166*, 2019. 6
- [39] Bohao Li, Rui Wang, Guangzhi Wang, Yuying Ge, Yixiao Ge, and Ying Shan. Seed-bench: Benchmarking multimodal llms with generative comprehension. *arXiv preprint arXiv:2307.16125*, 2023. 6
- [40] Guangyao Li, Yake Wei, Yapeng Tian, Chenliang Xu, Ji-Rong Wen, and Di Hu. Learning to answer questions in dynamic audio-visual scenarios. In *CVPR*, pages 19108–19118, 2022. 5, 6, 7
- [41] Guangyao Li, Yixin Xu, and Di Hu. Multi-scale attention for audio question answering. *arXiv preprint arXiv:2305.17993*, 2023. 6
- [42] Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. *arXiv preprint arXiv:2301.12597*, 2023. 1, 2, 3, 6
- [43] Kunchang Li, Yanan 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. 1, 3
- [44] Linjie Li, Yen-Chun Chen, Yu Cheng, Zhe Gan, Licheng Yu, and Jingjing Liu. Hero: Hierarchical encoder for video+ language omni-representation pre-training. *arXiv preprint arXiv:2005.00200*, 2020. 6
- [45] Samuel Lipping, Parthasaarathy Sudarsanam, Konstantinos Drossos, and Tuomas Virtanen. Clotho-aqa: A crowd-sourced dataset for audio question answering. In *2022 30th European Signal Processing Conference (EUSIPCO)*, pages 1140–1144. IEEE, 2022. 6, 7
- [46] Haotian Liu, Chunyuan Li, Yuheng Li, and Yong Jae Lee. Improved baselines with visual instruction tuning, 2023. 2, 3, 5, 6
- [47] Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. *NeurIPS*, 2023. 4, 5
- [48] Yuan Liu, Haodong Duan, Yuanhan Zhang, Bo Li, Songyang Zhang, Wangbo Zhao, Yike Yuan, Jiaqi Wang, Conghui He, Ziwei Liu, et al. Mmbench: Is your multimodal model an all-around player? *arXiv preprint arXiv:2307.06281*, 2023. 6
- [49] Kevin Lu, Aditya Grover, Pieter Abbeel, and Igor Mordatch. Pretrained transformers as universal computation engines. *arXiv preprint arXiv:2103.05247*, 1, 2021. 2, 4
- [50] Pan Lu, Swaroop Mishra, Tanglin Xia, Liang Qiu, Kai-Wei Chang, Song-Chun Zhu, Oyvind Tafjord, Peter Clark, and Ashwin Kalyan. Learn to explain: Multimodal reasoning via thought chains for science question answering. *NeurIPS*, 2022. 6
- [51] Huaishao Luo, Lei Ji, Ming Zhong, Yang Chen, Wen Lei, Nan Duan, and Tianrui Li. Clip4clip: An empirical study of clip for end to end video clip retrieval and captioning. *Neurocomputing*, 508:293–304, 2022. 3
- [52] Tiange Luo, Chris Rockwell, Honglak Lee, and Justin Johnson. Scalable 3d captioning with pretrained models. *arXiv preprint arXiv:2306.07279*, 2023. 5
- [53] Kenneth Marino, Mohammad Rastegari, Ali Farhadi, and Roozbeh Mottaghi. Ok-vqa: A visual question answering benchmark requiring external knowledge. In *CVPR*, 2019. 5, 6
- [54] Xinhao Mei, Chutong Meng, Haohe Liu, Qiuqiang Kong, Tom Ko, Chengqi Zhao, Mark D Plumbley, Yuexian Zou, and Wenwu Wang. Wavcaps: A chatgpt-assisted weakly-labelled audio captioning dataset for audio-language multimodal research. *arXiv preprint arXiv:2303.17395*, 2023. 5, 6
- [55] Suvir Mirchandani, Fei Xia, Pete Florence, Brian Ichter, Danny Driess, Montserrat Gonzalez Arenas, Kanishka Rao,

- Dorsa Sadigh, and Andy Zeng. Large language models as general pattern machines. In *CoRL*, 2023. 2
- [56] Anand Mishra, Shashank Shekhar, Ajeet Kumar Singh, and Anirban Chakraborty. Ocr-vqa: Visual question answering by reading text in images. In *ICDAR*, pages 947–952. IEEE, 2019. 5
- [57] Seungwhan Moon, Andrea Madotto, Zhaojiang Lin, Tushar Nagarajan, Matt Smith, Shashank Jain, Chun-Fu Yeh, Prakash Murugesan, Peyman Heidari, Yue Liu, et al. Any-mal: An efficient and scalable any-modality augmented language model. *arXiv preprint arXiv:2309.16058*, 2023. 2, 3, 5, 6
- [58] Pushmeet Kohli Nathan Silberman, Derek Hoiem and Rob Fergus. Indoor segmentation and support inference from rgb-d images. In *ECCV*, 2012. 7
- [59] Dat Tien Nguyen, Shikhar Sharma, Hannes Schulz, and Layla El Asri. From film to video: Multi-turn question answering with multi-modal context. *arXiv preprint arXiv:1812.07023*, 2018. 6
- [60] OpenAI. Gpt-4 technical report. *ArXiv*, abs/2303.08774, 2023. 1
- [61] Hoang-Anh Pham, Thao Minh Le, Vuong Le, Tu Minh Phuong, and Truyen Tran. Video dialog as conversation about objects living in space-time. In *ECCV*, pages 710–726. Springer, 2022. 6
- [62] Bryan A Plummer, Liwei Wang, Chris M Cervantes, Juan C Caicedo, Julia Hockenmaier, and Svetlana Lazebnik. Flickr30k entities: Collecting region-to-phrase correspondences for richer image-to-sentence models. In *ICCV*, pages 2641–2649, 2015. 6
- [63] Joan Puigcerver, Carlos Riquelme, Basil Mustafa, and Neil Houlsby. From sparse to soft mixtures of experts. *arXiv preprint arXiv:2308.00951*, 2023. 2, 4
- [64] 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*, pages 8748–8763. PMLR, 2021. 2, 3, 4, 7
- [65] René Ranftl, Alexey Bochkovskiy, and Vladlen Koltun. Vision transformers for dense prediction. *ArXiv preprint*, 2021. 5
- [66] C. Schuhmann, A. Köpf, R. Vencu, T. Coombes, and R. Beaumont. Laion coco: 600m synthetic captions from laion2b-en. 5
- [67] Christoph Schuhmann, Romain Beaumont, Richard Vencu, Cade Gordon, Ross Wightman, Mehdi Cherti, Theo Coombes, Aarush Katta, Clayton Mullis, Mitchell Wortsman, et al. Laion-5b: An open large-scale dataset for training next generation image-text models. *NeurIPS*, 35: 25278–25294, 2022. 1, 5
- [68] Dustin Schwenk, Apoorv Khandelwal, Christopher Clark, Kenneth Marino, and Roozbeh Mottaghi. A-okvqa: A benchmark for visual question answering using world knowledge. In *ECCV*, pages 146–162. Springer, 2022. 5, 6
- [69] Piyush Sharma, Nan Ding, Sebastian Goodman, and Radu Soricut. Conceptual captions: A cleaned, hypernymed, image alt-text dataset for automatic image captioning. In *Proceedings of ACL*, 2018. 5
- [70] Amanpreet Singh, Vivek Natarajan, Meet Shah, Yu Jiang, Xinlei Chen, Dhruv Batra, Devi Parikh, and Marcus Rohrbach. Towards vqa models that can read. In *CVPR*, pages 8317–8326, 2019. 6
- [71] Shuran Song, Samuel P Lichtenberg, and Jianxiong Xiao. Sun rgb-d: A rgb-d scene understanding benchmark suite. In *CVPR*, pages 567–576, 2015. 7
- [72] 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
- [73] 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. 2, 6
- [74] Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*, 2023. 4, 5, 6
- [75] Gido M Van de Ven and Andreas S Tolias. Three scenarios for continual learning. *arXiv preprint arXiv:1904.07734*, 2019. 4
- [76] Vicuna. Vicuna: An open-source chatbot impressing gpt-4 with 90%\* chatgpt quality. <https://vicuna.lmsys.org/>, 2023. 6
- [77] 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. *arXiv preprint arXiv:2205.14100*, 2022. 3
- [78] Xin Wang, Jiawei Wu, Junkun Chen, Lei Li, Yuan-Fang Wang, and William Yang Wang. Vatex: A large-scale, high-quality multilingual dataset for video-and-language research. In *ICCV*, pages 4581–4591, 2019. 6
- [79] Yi Wang, Kunchang Li, Yizhuo Li, Yanan He, Bingkun Huang, Zhiyu Zhao, Hongjie Zhang, Jilan Xu, Yi Liu, Zun Wang, et al. Internvideo: General video foundation models via generative and discriminative learning. *arXiv preprint arXiv:2212.03191*, 2022. 6
- [80] Yusong Wu, Ke Chen, Tianyu Zhang, Yuchen Hui, Taylor Berg-Kirkpatrick, and Shlomo Dubnov. Large-scale contrastive language-audio pretraining with feature fusion and keyword-to-caption augmentation. In *ICASSP*, pages 1–5. IEEE, 2023. 3
- [81] Junbin Xiao, Xindi Shang, Angela Yao, and Tat-Seng Chua. Next-qa: Next phase of question-answering to explaining temporal actions. In *CVPR*, pages 9777–9786, 2021. 6
- [82] Junbin Xiao, Angela Yao, Zhiyuan Liu, Yicong Li, Wei Ji, and Tat-Seng Chua. Video as conditional graph hierarchy for multi-granular question answering. In *AAAI*, pages 2804–2812, 2022. 6
- [83] Chenfeng Xu, Shijia Yang, Tomer Galanti, Bichen Wu, Xiangyu Yue, Bohan Zhai, Wei Zhan, Peter Vajda, Kurt

- Keutzer, and Masayoshi Tomizuka. Image2point: 3d point-cloud understanding with 2d image pretrained models. In *ECCV*, pages 638–656. Springer, 2022. 2
- [84] Dejing Xu, Zhou Zhao, Jun Xiao, Fei Wu, Hanwang Zhang, Xiangnan He, and Yueting Zhuang. Video question answering via gradually refined attention over appearance and motion. In *ACM Multimedia*, 2017. 5, 6, 7
- [85] Hu Xu, Gargi Ghosh, Po-Yao Huang, Dmytro Okhonko, Armen Aghajanyan, Florian Metze, Luke Zettlemoyer, and Christoph Feichtenhofer. Videoclip: Contrastive pre-training for zero-shot video-text understanding. *arXiv preprint arXiv:2109.14084*, 2021. 3
- [86] 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. 5
- [87] Runsen Xu, Xiaolong Wang, Tai Wang, Yilun Chen, Jiangmiao Pang, and Dahua Lin. Pointllm: Empowering large language models to understand point clouds. *arXiv preprint arXiv:2308.16911*, 2023. 3, 5, 6, 7
- [88] Antoine Yang, Antoine Miech, Josef Sivic, Ivan Laptev, and Cordelia Schmid. Just ask: Learning to answer questions from millions of narrated videos. In *ICCV*, pages 1686–1697, 2021. 6
- [89] Antoine Yang, Antoine Miech, Josef Sivic, Ivan Laptev, and Cordelia Schmid. Zero-shot video question answering via frozen bidirectional language models. *NeurIPS*, 35:124–141, 2022. 6
- [90] Qinghao Ye, Haiyang Xu, Guohai Xu, Jiabo Ye, Ming Yan, Yiyang Zhou, Junyang Wang, Anwen Hu, Pengcheng Shi, Yaya Shi, et al. mplug-owl: Modularization empowers large language models with multimodality. *arXiv preprint arXiv:2304.14178*, 2023. 2, 3
- [91] Zhongjie Ye, Yuqing Wang, Helin Wang, Dongchao Yang, and Yuexian Zou. Featurecut: An adaptive data augmentation for automated audio captioning. In *2022 Asia-Pacific Signal and Information Processing Association Annual Summit and Conference (APSIPA ASC)*, pages 313–318. IEEE, 2022. 6
- [92] Shoubin Yu, Jaemin Cho, Prateek Yadav, and Mohit Bansal. Self-chained image-language model for video localization and question answering. *arXiv preprint arXiv:2305.06988*, 2023. 6
- [93] Weihao Yu, Zhengyuan Yang, Linjie Li, Jianfeng Wang, Kevin Lin, Zicheng Liu, Xinchao Wang, and Lijuan Wang. Mm-vet: Evaluating large multimodal models for integrated capabilities. *arXiv preprint arXiv:2308.02490*, 2023. 6
- [94] Dong Zhang, Shimin Li, Xin Zhang, Jun Zhan, Pengyu Wang, Yaqian Zhou, and Xipeng Qiu. Speechgpt: Empowering large language models with intrinsic cross-modal conversational abilities. *arXiv preprint arXiv:2305.11000*, 2023. 1, 3
- [95] Hang Zhang, Xin Li, and Lidong Bing. Video-llama: An instruction-tuned audio-visual language model for video understanding. *arXiv preprint arXiv:2306.02858*, 2023. 1, 3
- [96] Renrui Zhang, Ziyu Guo, Wei Zhang, Kunchang Li, Xupeng Miao, Bin Cui, Yu Qiao, Peng Gao, and Hongsheng Li. Pointclip: Point cloud understanding by clip. In *CVPR*, pages 8552–8562, 2022. 3
- [97] Renrui Zhang, Jiaming Han, Aojun Zhou, Xiangfei Hu, Shilin Yan, Pan Lu, Hongsheng Li, Peng Gao, and Yu Qiao. Llama-adapter: Efficient fine-tuning of language models with zero-init attention. *arXiv preprint arXiv:2303.16199*, 2023. 2
- [98] Yiyuan Zhang, Kaixiong Gong, Kaipeng Zhang, Hongsheng Li, Yu Qiao, Wanli Ouyang, and Xiangyu Yue. Meta-transformer: A unified framework for multimodal learning. *arXiv preprint arXiv:2307.10802*, 2023. 2, 4
- [99] Zijia Zhao, Longteng Guo, Tongtian Yue, Sihao Chen, Shuai Shao, Xinxin Zhu, Zehuan Yuan, and Jing Liu. Chatbridge: Bridging modalities with large language model as a language catalyst. *arXiv preprint arXiv:2305.16103*, 2023. 2, 3, 5, 6
- [100] 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, 3, 4