Survey of Large Multimodal Model Datasets, Application Categories and Taxonomy

Priyaranjan Pattnayak[1], Hitesh Laxmichand Patel[2], Bhargava Kumar[3], Amit Agarwal[4], Ishan Banerjee[5], and Srikant Panda[6] Tejaswini Kumar[7]

[1] University of Washington, Seattle, USA
[2] New York University, New York, USA
[3] Columbia University, New York, USA
[4] Liverpool John Moores University, Liverpool, UK
[5] Chennai Mathematical Institute, Chennai, India
[6] Birla Institute of Technology, Pilani, USA
[7] Columbia University, New York, USA

Abstract. Multimodal learning, a rapidly evolving field in artificial intelligence, seeks to construct more versatile and robust systems by integrating and analyzing diverse types of data, including text, images, audio, and video. Inspired by the human ability to assimilate information through many senses, this method enables applications such as text-to-video conversion, visual question answering, and image captioning. Recent developments in datasets that support multimodal language models (MLLMs) are highlighted in this overview. Large-scale multimodal datasets are essential because they allow for thorough testing and training of these models. With an emphasis on their contributions to the discipline, the study examines a variety of datasets, including those for training, domain-specific tasks, and real-world applications. It also emphasizes how crucial benchmark datasets are for assessing models' performance in a range of scenarios, scalability, and applicability. Since multimodal learning is always changing, overcoming these obstacles will help AI research and applications reach new heights.

Keywords: Multimodal · LMM · LLM · Video · Audio · VLM.

1 Introduction to Multimodal Learning and Large Language Models

Multimodal learning, a growing field in AI, focuses on integrating and processing multiple data types like text, images, and audio, aiming to replicate human cognition, which naturally combines sensory inputs. This approach enables more robust and intelligent systems compared to single-modality methods.

Large language models (LLMs) such as GPT-3, BERT, and T5 excel in text-based tasks like question answering and summarization [36]. However, they struggle with non-text data, driving interest in multimodal large language models (MLLMs) that combine LLMs' language capabilities with computer vision's

strengths. MLLMs have achieved state-of-the-art results in tasks like image captioning and visual question answering [18]. Challenges remain, including limited high-quality datasets, high computational costs, and ethical concerns like bias and privacy [28]. Despite these hurdles, MLLMs hold transformative potential in healthcare, education, and research, making them a key focus in advancing AI.

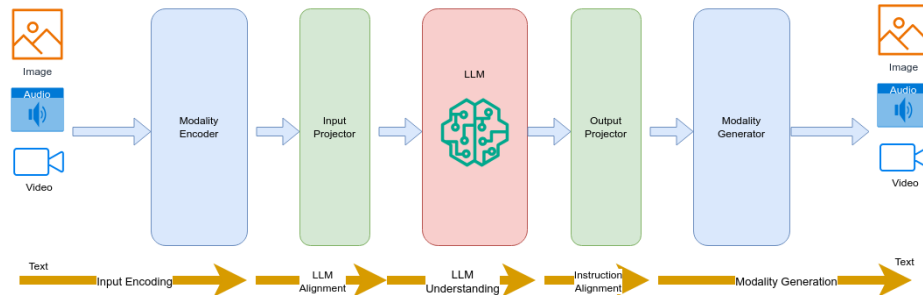## 1.1   Multimodal Learning: Foundations and Concepts

Fig. 1. Flowchart describing the multimodal language model pipeline

Multimodal learning consists of building models that can process and combine information from various data modalities, such as text, images, audio, and video. This is due to the fact that real-world experiences are inherently multimodal in nature, and the different types of information carried by the various modalities provide a way to understand such complex environments more thoroughly [28].

Multimodal learning incorporates multiple types of data: texts, images, audio, and video. These create representations, each unique to one modality. Given the diverse nature of various types of data, different methods have traditionally been used to capture their properties. Text, for example, is represented with word embeddings emphasizing meaning and structure [3], while image data most often depend on a convolutional neural network that would extract details from visual scenes. Similarly, audio data is often transformed into spectrograms or mel-frequency cepstral coefficients to capture patterns in time and frequency [41]. A typical pipeline for a large multimodal model (MLLM) is shown in Fig. 1, where inputs are first processed through a modality encoder to unify their representations. These are then refined by an input projector and passed into a Large Language Model (LLM) for deeper alignment and understanding. Finally, the output projector and modality generator transform the model's results into meaningful outputs, enabling tasks like generating multimodal content or translating between data types.

Fusion of modality representations is a key focus in multimodal learning. The widely used methods are early fusion, in which the concatenation or combination

of representations is carried out at an initial stage of processing [38], and late fusion, where modality-specific representations are combined later in the process, often with attention or gating mechanism [38].

Beyond representation and fusion, there are more challenges involved with multimodal learning, like alignment, translation, and co-learning. Alignment allows temporal or semantic synchronization across modalities, which is an essential task for video understanding or audio-visual speech recognition [5]. Translation helps in modality transformation, as generating images from text, for example [57]. Co-learning allows to learn in the conditions when some of the data modalities are not available or damaged, by transferring knowledge from available modalities [41].

Recent progress on LLMs, such as BERT, GPT, and DALL-E, has considerably accelerated progress in multimodal learning. These models are very good at both understanding and generating text; their extension to multiple data types now also enables answering questions about images, creating picture descriptions, or even generating images based on text [32].

In short, multimodal learning has been a very critical factor in the development of intelligent systems that effectively process and integrate information from varied sources. The complementary strengths of multiple modalities ensure that this area continuously creates innovation across domains like NLP, computer vision, and robotics, among others, with ever-widening scopes of applications and research directions.

1.2 Multimodal Large Language Models: Opportunities and Challenges

Recent advances in LLMs have laid the path for multimodal large language models that combine data across modalities, such as text, images, audio, and video [59]. MLLMs hold the potential to transform various domains by enhancing understanding and representation through a mix of different modalities.

MLLMs expand the capability of LLMs to wider ranges of tasks beyond traditional text-only models. This class of models is very strong on tasks like image captioning, visual question answering, and text-to-video generation-all requiring an in-depth understanding of language-visual relationships [63].

The integration of multi-modal data opens the way for scientific research and domain-specific applications for MLLMs by pushing the boundary. Some critical domains like medical imaging, autonomous driving, geospatial intelligence combine textual, visual, and sensor data to yield more realistic decision-making processes.

Despite the potential of MLLMs, there are significant challenges in developing them. Among the primary issues is the absence of large-scale, high-quality multimodal datasets [49]. Complex, unbiased data covering the richness of reality is a necessary ingredient to train robust MLLMs[28].

Another challenge is the increase in computational demands and complexity in integrating these various modalities. The training and deployment of MLLMs

require considerable resources; thus, there is a need to develop novel model architectures, efficient training strategies, and hardware capabilities [28].

Finally, ensuring the reliability, interpretability, and ethical alignment of MLLMs is important. With increased sophistication of these models, there is a growing need to provide insights into their decision-making processes to reduce biases and align them more closely with human values. The development of robust evaluation frameworks and interpretability tools are necessary to engender trust in MLLMs [45].

Despite this, the prospects for MLLMs are enormous. While incorporating multimodal data, the models pave the way for a better comprehension of complicated scenarios, hence giving birth to new applications and promoting scientific research accordingly. In addition, future interdisciplinary collaboration and emphasis on ethical considerations are the critical elements toward the transformation that can be brought by MLLMs [28].

In the following sections, we classify datasets that are critical for MLLMs into three major types: training-specific datasets, task-specific datasets, and domain-specific datasets, as illustrated in Fig. 2.


Fig. 2. An illustration representing the high-level classification of the datasets mentioned in the survey under Training specific (datasets under MM-IT and MM-PT), Task specific and Domain specific.

2   Multimodal Datasets for Training Specific Needs

The development of multimodal datasets is essential to advancing MLLMs. These datasets span a wide range of modalities and applications, enabling researchers to train models that can integrate and reason across various data types. Challenges

in dataset design include ensuring scale, modality diversity, annotation quality, and applicability to real-world scenarios. In this section, we discuss few datasets that are used for training (both pre-training and instruction tuning) multimodal models, as illustrated in Fig. 2 under training specific types.

To train large-scale multimodal models for tasks like picture captioning, visual question answering, and audio-text understanding, the MLLMs dataset [59] consists of text, images, video, and audio. For academics looking into multimodal integration, its scope across several modalities offers a cohesive framework. In Fig. 3, datasets made available as part of MLLM models are displayed.

Multimodal Pre-Training (MM-PT) and Multimodal Instruction Tuning (MM-IT) are the two primary tenets of the MLLM development pipeline. Image-Text, Video-Text, and Audio-Text datasets are utilized in these phases and are further separated into:

– Image-Text Pairs: Simple <img1> <txt1> format.
– Interleaved Image-Text Corpus: Mixed sequences like <txt1> <img1> <txt2> <txt3> <img2>.

The process of pre-training a model to comprehend and align modalities, including text and a picture, is known as MM-PT. Along with offering a basis for comprehending and aligning various data kinds, MM-PT enables the capture of links between various modalities, such as connecting verbal descriptions to visual attributes. This is crucial for activities like visual reasoning or image captioning. MM-IT is an approach that uses datasets formatted as instructions to adapt pre-trained MM-LLMs. In exploring the landscape of datasets, we leveraged information from existing literature [4,60,52]. Table 1 provides a comprehensive list of such datasets used for MM-PT, including modality, size, and year of release.

The instruction-formatted datasets used for tuning in the MM-IT stage improve zero-shot performance by increasing task flexibility using the following non-exhaustive list of methods:

– Supervised Fine Tuning (SFT) improves MM-LLMs on visual question answering through the conversion of data into templates that may be utilized-for example, single-turn or multi-turn.
– Reinforcement Learning with Human Feedback (RLHF) aligns model responses with human purpose by using natural language input.

Notable Datasets in MM-PT and MM-IT

– **LAION-5B:** The LAION-5B[43] dataset contains 5.85 billion CLIP-filtered image-text pairs (2.32 billion in English) supporting large-scale multi-modal research with tools for watermark detection and NSFW filtering.
– **MS-COCO:** The MS-COCO[9] contains over 330K images with five human-written captions each, serving as a benchmark for image recognition, segmentation, and captioning.

Fig. 3. The datasets released as part of MLLMs

– **Flickr30k:** Flickr30k[54] contains 31,000 images, each annotated with five captions, for image captioning and cross-modal retrieval.
– **COYO-700M:** The COYO-700M[4] dataset includes 747 million image-text pairs, extracted from Common Crawl and filtered for quality for model training.
– **CC12M:** CC12M(Conceptual Captions 12M)[8] is a large dataset with 12M image-text pairs for training vision-language models. It uses web images and their captions, offering varied data to support robust learning.
– **OpenVid-1M:** OpenVid-1M[35] is a dataset of 1M high-quality text-video pairs, created to support text-to-video generation research.
– **AISHELL-2:** AISHELL-2 [11] is a Mandarin speech corpus with 1 million transcribed segments from 400 speakers, offering high-quality data for automatic speech recognition tasks.
– **LAMM:** It provides 186K language-image and 10K language-3D instruction-response pairs for MLLMs, with a modular framework for 2D and 3D vision tasks [53].
– **Wukong:** With 100 million image-text pairs, the Wukong dataset is a comprehensive Chinese multimodal resource that focuses on vision-language pre-training [60].

Table 1. Overview of datasets under MM-PT.

Dataset	Modality	Description	Size	Year
ALIGN	Image	Large-scale image-text pairs for training vision-language models	1.8B pairs	2021
LTIP	Image	Large-scale image-text pairs dataset	312M pairs	2022
Visual Genome	Image	Dataset with object relationships and attributes for VQA	108K images	2017
CC3M	Image	Conceptual Captions dataset	3.3M pairs	2018
CC12M	Image	Large-scale captioning dataset for vision-language tasks	12.4M pairs	2021
SBU	Image	Image-caption pairs dataset	1M pairs	2011
LAION-5B	Image	Large-scale web-crawled image-text pairs	5.9B pairs	2022
LAION-400M	Image	Subset of LAION dataset with 400M image-text pairs	400M pairs	2021
MS-COCO	Image	Object detection, segmentation, and captioning dataset	124K images	2014
YFC100M	Image	Large-scale collection of web-crawled images and videos	99.2M pairs	2014
LAION-COCO	Image	COCO-aligned subset of LAION dataset	600M pairs	2022
Flickr30k	Image	Image-captioning dataset with real-world scenes	31K images	2014
COYO-700M	Image	Large-scale web-crawled dataset for image-text pairs	747M pairs	2022
Wukong	Image	Chinese language image-text pairs dataset	101M pairs	2022
WebLI	Image	Large-scale dataset for multilingual image-text pairs	10B images	2022
FILIP	Image	Fine-grained vision-language dataset	300M images	2021
CLIP	Image	Large-scale image-text pairs, collected from diverse sources	400M images	2021
WebVid	Video	Large-scale video dataset with captions	10M videos	2021
Youku-mPLUG	Video	Large-scale Chinese Video-Language Dataset	10M videos	2023
VideoCC3M	Video	Large-scale Video-Language Dataset	6.3M videos	2022
YouTube8M	Video	Large and Diverse Labeled Video Dataset for Video	8M videos	2022
Vript	Video	Large and fine-grained video-text dataset	420K videos	2024
OpenVid-1M	Video	Large-Scale High-Quality Dataset for Text-video	1M videos	2024
VidGen-1M	Video	Large-Scale High-Quality Dataset for Text-video	1M videos	2024
VTP	Video	Dataset with video-text pairs	27M videos	2022
AudioSet	Audio	Dataset with audio-video pairs	2M audio-videos	2011
YODAS	Audio	Dataset with audio-text pairs	500 K audio	2024
LAION-Audio	Audio	large-scale audio-text dataset	630K audio-text	2023
AVSET-10M	Audio	large-scale audio visual dataset	10M audio-video	2023
Auto-ACD	Audio	large-scale audio-text dataset	1.5M audio-text	2023
AISHELL-1	Audio	Chinese language audio dataset for ASR	128K utterances	2023
AISHELL-2	Audio	Expanded AISHELL audio dataset for ASR	1M utterances	2023
WaveCaps	Audio	Dataset with audio-caption pairs	403K pairs	2023
VSDial-CN	Image, Audio	Chinese language visual dialogue dataset	1.2M audio utterances	2024

- **VRIPT:** Each of the 420,000 clips and 12,000 high-resolution videos in the VRIPT dataset [51] has a comprehensive caption that describes the content, shot kinds, and camera motions. It is intended for sophisticated video comprehension and production, including transcriptions.
- **SBU:** The SBU Caption Dataset consists of 1 million image-text pairs, where image captions are sourced from user-provided descriptions. It is widely used for training and evaluating vision-and-language models[52].
- **LLaVA-Instruct:** The dataset contains 150K set of GPT-generated multi-modal instruction-following data to train LLaVa[29] model.
- **SVIT:** SVIT[61] contains 4.2 million examples for visual captioning and QA tasks, created with GPT-4 and manual annotations, enabling models like SVIT-v1.5 to excel in visual understanding benchmarks.
- **MIMIC-IT:** MIMIC-IT[24] includes 2.8 million multimodal samples for instruction-following tasks across multiple languages, supporting models like OpenFlamingo in understanding complex multimodal instructions.
- **SoM-LLaVA's IT:** SoM-LLaVA's[50] Instruct dataset consists of 695,000 image-text pairs with structured visual prompts, enhancing model reasoning in tasks like visual question answering using set-of-mark prompting.

We utilized information from existing literature [4,52] to curate the list of MM-IT datasets. Table 2 shows a comprehensive list of all such datasets that

have been used for MM-IT, including modality, size, and year of release. These further align with human purpose by feedback and instruction-aware formatting to develop better interaction capabilities for the MLLMs.

Table 2. Overview of datasets under MM-IT.

Dataset	Modality	Description	Size	Year
MiniGPT-4 IT	Image + Text	Dataset for training generative models with images and text pairs.	5K	2023
StableLLaVA	Image + Text	Robust dataset for vision-language tasks with diverse image-text pairs.	126K	2023
Q-Pathway	Image + Text	Dataset for improving low-level visual perception.	19K	2023
LLaVA-Instruct-150K	Image + Text	Image-text instruction tuning dataset.	150K	2023
ShareGPT4V's IT	Image + Text	Interactive dataset integrating multiple visual sources with text.	100K	2023
MultiModalGPT IT	Image + Text	Dataset designed to visual inputs.	6K	2023
SoM-LLaVA's IT	Image + Text	Dataset for fine-tuning models on complex visual prompts.	695K	2024
VideoChat's IT	Video + Text	Dataset for training conversational agents with video inputs.	11K	2023
Video-ChatGPT's IT	Video + Text	Dataset for generating dialogues based on video content.	100K	2023
Video-LLaMA's IT	Image/Video + Text	Integrative dataset combining image and video for text generation.	171K	2023
InstructBLIP's IT	Image/Video + Text	Comprehensive dataset for instruction-based visual understanding.	1.6M	2023
X-InstructBLIP's IT	Image/Video/Audio/3D + Text	Dataset combining multiple modalities for complex interactions.	1.8M	2023
MIMIC-IT	Image/Video + Text	Medical dataset integrating image and video data with annotations.	2.8M	2023
VSTaR-1M	Image/Video + Text	Video instruction tuning dataset.	1M	2023
PandaGPT's IT	Image + Text	Dataset for training generative conversational models with visuals.	160K	2023
MGVLID	Image + Bounding Box + Text	Visual dataset with bounding box annotations for training.	108K	2023
M³IT	Image/Video/Bounding Box + Text	Multi-modal dataset for rich visual-text interactions.	2.4M	2023
LAMM	Image + 3D + Text	Dataset focusing on 3D interactions and image-text pairs.	196K	2023
BuboGPT's IT	(Image + Audio)/Audio + Text	Mixed dataset for training with audio and visual inputs.	9K	2023
mPLUG-DocOwl'sIT	Image/Table/Web + Text	Dataset for document understanding across various formats.	–	2023
T2M	Text to Image/Video/Audio + Text	Transformative dataset for generating visuals from text prompts.	14.7K	2023
MosIT	Image + Video + Audio + Text	Complex dataset for multi-modal training and interactions.	5K	2023
Osprey's IT	Image + Text	Diverse dataset for fine-tuning models with image inputs.	724K	2023
LLaVA-RLHF	Image + Text	Dataset focusing on reinforcement learning from human feedback.	10K	2023
RLHF-V's IT	Image + Text	Dataset designed for video input interactions in RL training.	1.4K	2023
Shikra	Image + Text	Language vision question answeing dataset	156K	2022
Vision-Flan	Image + Text	Language vision instruction tuning dataset	1.6M	2024
RTVLM	Image + Text	Recent dataset focusing on real-time visual language modeling.	5K	2024
LVIS-Instruct4V	Image + Text	fine-grained visual instruction dataset.	220K	2024
SVIT	Image + Text	Large scale visual instruction dataset.	4.2M	2024
MultiInstruct	Image + Text	Multi-modal dataset for instruction tuning.	235K	2023
MAmmoTH-VL-Instruct	Image + Text	Multi-modal dataset for instruction tuning.	12M	2024
Leopard-Instruct	Image + Text	text rich dataset for instruction tuning.	925K	2023
MMInstruct	Image + Text	High-Quality data Extensive Diversity	973K	2024

3   Multimodal Datasets for Task Specific Applications

The datasets discussed in this section form the core of designing versatile models that can work on a wide range of tasks, including sentiment analysis, emotion detection, and visual question answering. They form the key building blocks for robust and adaptable multimodal frameworks. Several of these datasets are highlighted in Fig. 4, while a detailed overview is given in Table 3, including dataset modality, size, and release year.

SlideVQA: SlideVQA [46] expands on single-image VQA datasets by introducing a multi-image framework for complex reasoning tasks. It includes over 2,600 slide decks (52,000+ images) and 14,500+ questions, focusing on multi-hop reasoning and numerical analysis with annotated arithmetic expressions. The dataset presents challenges in real-world document understanding, such as evidence identification and sequence-based reasoning across multiple images.

Peacock dataset: The Peacock dataset [1] addresses a gap in multilingual, culturally relevant multimodal resources by focusing on the Arabic language and its cultural context. It enables the training of multilingual large language models (MLLMs) tailored for Arabic-speaking populations, combining text, visuals, and culturally significant elements. This enhances the models' applicability across various languages and cultures.

OmniACT: OmniACT [21] serves as a benchmark for evaluating agents' ability to generate programs and test models for computer task automation. It includes a wide range of tasks, from simple commands like "Play the next song" to complex ones like "Send an email to John Doe with meeting details." The dataset combines screen visuals with task descriptions, supporting studies in automated task execution through visually grounded interfaces.

InternVid: InternVid [48] is a large-scale video-text dataset with 7 million videos (760,000 hours, 234 million snippets, and 4.1 billion words) designed for learning video-text representations. Leveraging the ViCLIP model and contrastive learning, it excels in zero-shot action recognition, video retrieval, and text-to-video/video-to-text tasks. This dataset significantly advances video-based multimodal tasks, including dialogue systems.

ImageNet: Tens of millions of annotated photos (150Gb) that are in line with the WordNet semantic hierarchy are available through ImageNet [10], a large-scale image ontology. Applications that show the value of sizable, well-organized image datasets include object identification, image categorization, and clustering.

CMU MultimodalSD dataset: The CMU MultimodalSD dataset [56] contains over 1,000 YouTube video clips, each accompanied by text transcripts and sentiment annotations. This dataset is designed for multimodal sentiment analysis, where visual and textual information are integrated to train models for predicting sentiment in real-time video scenarios.

MELD: The MELD dataset [39] contains 13,000+ utterances from 1,433 dialogues, annotated into seven emotion categories. It integrates acoustic and textual data for emotion recognition in real-world conversations. The dataset ensures accuracy through timestamp alignment and majority-vote annotation. Compared to EmotionLines, it reduces dialogues by constraining timestamps and shared scenes. After resolving annotation disagreements, final dataset consists of audiovisual clips, 16-bit PCM WAV audio, and aligned visual, audio, and text modalities.

Fig. 4. An illustration of the datasets as per the survey under Task Specific Needs

HowTo100M: The HowTo100M dataset [33] contains 136 million video clips sourced from over 1.22 million narrated instructional videos, spanning more than 23,000 tasks. Models trained on this dataset excel in text-to-video retrieval and action localization, outperforming benchmarks like YouCook2 and CrossTask, and generalize effectively to other datasets, including MSR-VTT and LSMDC, after fine-tuning. The dataset, along with associated tools and models, is openly available for research.

AudioCaps: Using the AudioSet dataset [15], Audio Captioning in the Wild [22] offers a collection of 46K audio clips with crowdsourced human-generated captions. It investigates sophisticated captioning models and audio representation approaches. Notably, it adds two new elements to attention-based models for better captioning performance: aligned semantic attention and a top-down multi-scale encoder.

CommonVoice: The Common Voice dataset [2] is a large-scale, multilingual resource for Automatic Speech Recognition (ASR) and language identification, built through crowdsourcing. By November 2019, it included contributions from over 50,000 participants, amassing 2,500 hours of audio across 29 supported languages, with data collection underway for 38 languages. This makes Common Voice the largest publicly accessible audio corpus for speech recognition tasks.

EmotionCaps: EmotionCaps [31] introduces a dataset of 120,000 audio clips with synthetic descriptions enriched with emotional content, addressing the gap

in emotion-aware audio-language pretraining. By incorporating emotional data, models trained on EmotionCaps generate captions that better reflect audio emotions, surpassing baseline models and driving research in audio captioning.

M3ED: The M3ED dataset [62] comprises 990 dialogues and 24,449 utterances from 56 TV series, annotated with seven emotions—happy, surprise, sad, disgust, anger, fear, and neutral—across acoustic, visual, and textual modalities. As the first Chinese multimodal emotional dialogue resource, it facilitates cross-cultural emotion analysis. Additionally, it introduces the MDI framework for context-aware emotion recognition, achieving strong performance with state-of-the-art methods.

ShapeNet: ShapeNet [7] is a large-scale repository of 3D CAD models, organized using the WordNet taxonomy and enriched with semantic annotations like part structures, symmetry planes, rigid alignments, and physical dimensions. Featuring over 3,000,000 indexed models, with 220,000 categorized into 3,135 WordNet synsets, it serves as a key resource for object visualization, geometric analysis, and benchmarking in computer graphics and vision research.

SQA3D: Situated Question Answering in 3D Scenes (SQA3D) [30] introduces a task to assess embodied agents' understanding of 3D environments through reasoning about spatial relations, commonsense, navigation, and multi-hop logic. Using 650 ScanNet scenes, the dataset includes 6.8k situations, 20.4k descriptions, and 33.4k questions. Results reveal a significant performance gap, with state-of-the-art models achieving 47.20% accuracy compared to 90.06% by humans, underscoring the challenge of 3D multi-modal reasoning.

TVQA: TVQA [23] is a multimodal dataset for video question answering, comprising 925 episodes from six TV shows across sitcoms (*The Big Bang Theory, How I Met Your Mother, Friends*), medical dramas (*Grey's Anatomy, House*), and a crime drama (*Castle*), totaling 461 hours. The dataset includes 21,793 clips, segmented into 60-second or 90-second intervals, with aligned subtitles and character-tagged transcripts to facilitate question-answer generation.

MMSum: Multimodal summarization with multimodal output (MSMO) has potential but is hindered by dataset limitations such as poor maintenance and small size. To overcome this, the MMSum dataset [40] was developed, featuring human-validated summaries across 17 categories and 170 subcategories, ensuring diverse, high-quality multimodal learning. Benchmarking on MMSum supports video, text, and multimodal summarization tasks.

LibriSpeech: LibriSpeech [37] is a 1000-hour read English speech dataset derived from LibriVox audiobooks, sampled at 16 kHz. Designed for training and

Table 3. Comprehensive summary of the datasets under task specific application with modality, description, size and year of release.

Dataset	Modality	Description	Size	Time
OmniACT	Visuals + Text	A dataset and benchmark for assessing agents' abilities to generate scripts.	9802 data points	2024
InternVid	Video + Text	A large-scale video-centric dataset for video-text representation learning.	7M+ videos	2024
SlideVQA	Image + Text	A multi-image document VQA dataset requiring reasoning in slide decks.	2.6k+ slide decks, 52k+ images	2024
Peacock dataset	Text + Image	Dataset created for training the Peacock models.	916K image-text pairs	2024
ImageNet	Image	Large-scale object recognition dataset with millions of images.	150GB	2021
CMU MultimodalSD	Text + Image + Audio	Synchronized dataset of video, audio, and text data.	12GB	2023
M3ED	Text + Video	Emotion recognition dataset from videos.	35GB	2022
MELD	Text + Audio	Emotion recognition in dialogues from TV shows.	5GB	2019
TVQA	Text + Video	Question-answering dataset based on TV show videos.	200K QA pairs	2023
CommonVoice	Audio	Multilingual voice recordings for speech recognition.	1,000 hours	2019
AudioCaps	Audio	Audio captioning dataset for sound events.	30K audio clips	2019
HowTo100M	Video + Text	Instructional videos with textual descriptions.	100M videos	2019
EmotionCaps	Text + Audio	Audio clips annotated with emotional descriptions.	100K audio clips	2023
ShapeNet	3D	3D models dataset for shape understanding tasks.	51K models	2015
SQA3D	Text + 3D	Spatial question answering in 3D environments.	40K scenes	2022
MMSum	Text	Dataset for multimodal summarization.	200K documents	2023
LibriSpeech	Audio	Corpus of read English speech for ASR.	1,000 hours	2019
VoxCeleb	Audio	Speaker recognition dataset with celebrity voice samples.	1,200 hours	2020
CATER	Video	Action recognition dataset in video content.	130K clips	2020

evaluating speech recognition systems, it includes language-model training data and pre-built models. Models trained on LibriSpeech show superior performance with lower error rates on WSJ test sets compared to those trained directly on WSJ data.

CATER: CATER [16] tackles the challenges in video understanding by providing a synthetic dataset designed to evaluate spatiotemporal reasoning with 3D objects and intricate motions. It offers diagnostic tools to assess such models.

VoxCeleb: VoxCeleb [34] is a dataset of over 100,000 audio-visual utterances from 1,251 celebrities, collected via an automated pipeline using YouTube videos, synchronization CNNs for speaker activity, and facial recognition for identity. It includes diverse demographics, professions, and accents, with real-world audio conditions like noise and degradation. Benchmarking shows CNN-based models excel in speaker identification. Unlike SITW, it contains both audio and video, with VoxCeleb2 offering an expanded version.

4   Multimodal Datasets for Domain-Specific Applications

Recent advancements in large language models have spurred the creation of multimodal datasets tailored to specific domains, complementing existing task-focused datasets. These domain-specific datasets leverage the integration of multiple modalities to address unique challenges across industries. For a detailed overview, refer to Fig. 5 and Table 4.

Medical Imaging: Domain-specific multimodal datasets, particularly in medical imaging, are instrumental for tasks like diagnosis, treatment planning, and patient monitoring. For instance, the MIMIC-CXR dataset [20] pairs radiological images with clinical reports, enabling models to understand correlations between visual data and medical language. Similarly, PathGen-1.6M [55], an open-source

dataset with 1.6 million image-caption pairs, aids in training models for pathological image analysis, supporting disease detection and prognosis.

Autonomous Driving: Multimodal datasets are essential in autonomous driving, integrating inputs like camera footage, LiDAR, and GPS for robust self-driving systems. Widely used datasets, such as the KITTI Vision Benchmark Suite [14] and nuScenes [6], support training models to navigate urban environments, detect objects, and predict road user behavior.

Geospatial Intelligence: Multimodal datasets, integrating satellite images, topographic maps, and textual data, are pivotal in geospatial intelligence for applications like land-use classification, change detection, and disaster response[42]. Examples such as SpaceNet [12] and EuroSAT [19] enable LLMs to derive significant insights by leveraging diverse geospatial information sources.


Fig. 5. Datasets grouped by various domains

Scientific Domain: Two prominent datasets enhance scientific figure retrieval tasks. SciOL [47], an extensive open-access corpus, supports multimodal models in scientific research, while MuLMS-Img [47] focuses on high-quality image-text pairings in materials science. Both demonstrate significant improvements in figure classification, captioning, and retrieval tasks. Another resource, MMSci [25], compiles data from Nature Communications across 72 disciplines, offering benchmarks for figure captioning and multiple-choice tasks. This dataset reveals substantial model performance gaps, underscoring its potential as a valuable training tool.

Egocentric and Interactive environment: Recent advances in egocentric video datasets have significantly enhanced research in computer vision and natural language processing, particularly for analyzing everyday activities and interactions. Notable contributions include Ego4D [17], featuring 3,670 hours of first-person videos, and ALFRED [44], which links natural language instructions to household action sequences.

MINEDOJO [13] builds on this by offering a framework for developing versatile agents through simulations, a multimodal knowledge base, and flexible agent architectures.

Miscellaneous: The Beijing Text-Traffic (BjTT) dataset [58] offers a comprehensive multimodal resource for Intelligent Transportation Systems (ITS) traffic prediction. By integrating textual context with time-series traffic data, it overcomes the constraints of traditional datasets with over 32,000 traffic records paired with descriptive text.

5   Dataset Characteristics and Limitations

The development of advanced multimodal language models relies heavily on large, diverse, and well-annotated datasets. For example, datasets like MS-COCO [9] offer extensive real-world data distributions, enabling models to learn intricate cross-modal relationships. Broad modality coverage, as seen in Flickr30k [54], and high-quality annotations enhance training and evaluation effectiveness. Practical applicability is supported by datasets like SpaceNet [12], which align models with real-world tasks.

Table 4. Comprehensive summary of the datasets under domain specific needs with the dataset type, description, size and year of release.

Dataset	Type	Description	Size	Time
MIMIC-CXR	Image + Text	Chest X-ray images with radiology reports.	370,000 images	2019
PathGen	Image + Text	A synthetic dataset used for training and evaluating models in generating and understanding complex path trajectories	1.6M images	2023
SciOL	Text	Science-oriented question-answering dataset.	1M QA pairs	2024
MMSci	Text + Image	Multimodal science dataset with images and text.	200K instances	2024
MINEDOJO	Text + Video	Interactive learning dataset from gameplay videos.	10K clips	2022
ALFRED	Text + 3D	Instruction following dataset in 3D environments.	100K tasks	2020
Ego4D	Video	Egocentric video understanding dataset.	7K hours	2022
KITTI	Image + Lidar	Autonomous driving dataset with images and Lidar data.	40GB	2012
nuScenes	Image + Lidar	Self-driving car dataset with urban data.	1.4TB	2020
EuroSat	Image	Satellite image classification dataset.	1.5GB	2019
SpaceNet	Image	Satellite imagery for urban mapping tasks.	600GB	2018
BjTT	Text + Time Series	A dataset for traffic prediction, combining time-series traffic data with descriptive text.	32,000+ time-series records	2024

However, challenges persist, including data biases, imbalances , and limited diversity in task representation. Moreover, privacy and security concerns underscore the need for responsible dataset use[28]. Overcoming these challenges through thoughtful design and curation is vital for advancing robust and ethical multimodal learning.

6   Emerging Trends and Future Dataset Needs

Multimodal learning has seen significant progress, propelled by advancements in large language models (LLMs) and expansive multimodal datasets. A key trend is the creation of increasingly diverse and complex datasets that mirror real-world scenarios, enhancing the field's potential for addressing practical challenges.

- **Diverse and Geographically Representative Datasets:** Researchers are advancing multimodal datasets by integrating tactile, olfactory, and physiological signals, while prioritizing geographical and linguistic diversity to enhance generalization and mitigate biases in large-scale models [5].
- **Complex Interactions and Real-World Applications:** Future multi-modal datasets should encompass intricate intermodal interactions while tackling practical applications like healthcare, autonomous systems, and environmental monitoring [26]. To ensure transparency, reproducibility, and ethical use, standardized documentation and benchmarking practices are essential [27].

7 Conclusion

Advancement of multimodal learning relies on development of specialized datasets, which we have grouped into three main categories: training needs, task-specific needs, and domain-specific needs. Training datasets, which are vital for pretraining and instruction tuning methods such as SFT and RLHF, act as the backbone for building and refining multimodal systems.The second category includes needs that are task-specific: these are datasets built for specific tasks that make a model perform well for narrow applications. Finally, domain-specific datasets are essential for solving the peculiarities of specific industries and making the model adaptable and context-sensitive. Although challenges are still there, such as data diversity and effective cross-modal integration, emerging trends in data augmentation and cross-modal learning are mellowing these challenges. Therefore, the development and diversification of these categories of datasets will be highly essential to unlocking the potentials of multimodal systems and further innovations in AI.

References

1. Alwajih, F., Nagoudi, E., Bhatia, G., Mohamed, A., Abdul-Mageed, M.: Peacock: A family of arabic multimodal large language models and benchmarks. arXiv preprint arXiv:2403.01031 (2024)
2. Ardila, R., Branson, M., Davis, K., Henretty, M., Kohler, M., Meyer, J., Morais, R., Saunders, L., Tyers, F., Weber, G.: Common voice: A massively-multilingual speech corpus. arXiv preprint arXiv:1912.06670 (2019)
3. Aytar, Y., Vondrick, C., Torralba, A.: See, hear, and read: Deep aligned representations. arXiv preprint arXiv:1706.00932 (2017)
4. Bai, T., Liang, H., Wan, B., Xu, Y., Li, X., Li, S., Yang, L., Li, B., Wang, Y., Cui, B., Huang, P., Shan, J., He, C., Yuan, B., Zhang, W.: A survey of multimodal large language model from a data-centric perspective (2024), `https://arxiv.org/abs/2405.16640`
5. Baltrušaitis, T., Ahuja, C., Morency, L.: Multimodal machine learning: A survey and taxonomy. IEEE transactions on pattern analysis and machine intelligence **41**(2), 423–443 (2018)

6. Caesar, H., Bankiti, V., Lang, A.H., Vora, S., Liong, V.E., Xu, Q., Krishnan, A., Pan, Y., Baldan, G., Beijbom, O.: nuscenes: A multimodal dataset for autonomous driving (2020), https://arxiv.org/abs/1903.11027

7. Chang, A., Funkhouser, T., Guibas, L., Hanrahan, P., Huang, Q., Li, Z., Savarese, S., Savva, M., Song, S., Su, H., Xiao, J.: Shapenet: An information-rich 3d model repository. arXiv preprint arXiv:1512.03012 (2015)

8. Changpinyo, S., Sharma, P., Ding, N., Soricut, R.: Conceptual 12m: Pushing web-scale image-text pre-training to recognize long-tail visual concepts (2021), https://arxiv.org/abs/2102.08981

9. Chen, X., Fang, H., Lin, T.Y., Vedantam, R., Gupta, S., Dollar, P., Zitnick, C.L.: Microsoft coco captions: Data collection and evaluation server (2015), https://arxiv.org/abs/1504.00325

10. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Li, F.: Imagenet: A large-scale hierarchical image database. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 248–255 (2009)

11. Du, J., Na, X., Liu, X., Bu, H.: Aishell-2: Transforming mandarin asr research into industrial scale (2018), https://arxiv.org/abs/1808.10583

12. Etten, A.V., Lindenbaum, D., Bacastow, T.M.: Spacenet: A remote sensing dataset and challenge series (2019), https://arxiv.org/abs/1807.01232

13. Fan, L., Wang, G., Jiang, Y., Mandlekar, A., Yang, Y., Zhu, H., Tang, A., Huang, D., Zhu, Y., Anandkumar, A.: Minedojo: Building open-ended embodied agents with internet-scale knowledge. Advances in Neural Information Processing Systems **35**, 18343–18362 (2022)

14. Geiger, A., Lenz, P., Urtasun, R.: Are we ready for autonomous driving? the kitti vision benchmark suite. In: Conference on Computer Vision and Pattern Recognition (CVPR) (2012)

15. Gemmeke, J., Ellis, D., Freedman, D., Jansen, A., Lawrence, W., Moore, R., Plakal, M., Ritter, M.: Audio set: An ontology and human-labeled dataset for audio events. In: 2017 IEEE international conference on acoustics, speech and signal processing (ICASSP). pp. 776–780 (2017)

16. Girdhar, R., Ramanan, D.: Cater: A diagnostic dataset for compositional actions and temporal reasoning. CoRR (2019), https://arxiv.org/abs/1910.04744

17. Grauman, K., Westbury, A., Byrne, E., Chavis, Z., Furnari, A., Girdhar, R., Hamburger, J., Jiang, H., Liu, M., Liu, X., Martin, M.: Ego4d: Around the world in 3,000 hours of egocentric video. Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition pp. 18995–19012 (2022)

18. Hakimov, S., Schlangen, D.: Images in language space: Exploring the suitability of large language models for vision & language tasks. arXiv preprint arXiv:2305.13782 (2023)

19. Helber, P., Bischke, B., Dengel, A., Borth, D.: Eurosat: A novel dataset and deep learning benchmark for land use and land cover classification (2019)

20. Johnson, A.E.W., Pollard, T.J., Greenbaum, N.R., Lungren, M.P., ying Deng, C., Peng, Y., Lu, Z., Mark, R.G., Berkowitz, S.J., Horng, S.: Mimic-cxr-jpg, a large publicly available database of labeled chest radiographs (2019)

21. Kapoor, R., Butala, Y., Russak, M., Koh, J., Kamble, K., Alshikh, W., Salakhutdinov, R.: Omniact: A dataset and benchmark for enabling multimodal generalist autonomous agents for desktop and web. arXiv preprint arXiv:2402.17553 (2024)

22. Kim, C., Kim, B., Lee, H., Kim, G.: 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. pp. 119–132 (2019)

23. Lei, J., Yu, L., Bansal, M., Berg, T.: Tvqa: Localized, compositional video question answering. arXiv preprint arXiv:1809.01696 (2018)
24. Li, B., Zhang, Y., Chen, L., Wang, J., Pu, F., Yang, J., Li, C., Liu, Z.: Mimic-it: Multi-modal in-context instruction tuning (2023), https://arxiv.org/abs/2306.05425
25. Li, Z., Yang, X., Choi, K., Zhu, W., Hsieh, R., Kim, H., Lim, J., Ji, S., Lee, B., Yan, X., Petzold, L.: Mmsci: A multimodal multi-discipline dataset for phd-level scientific comprehension. arXiv preprint arXiv:2407.04903 (2024)
26. Liang, P., Cheng, Y., Fan, X., Ling, C., Nie, S., Chen, R., Deng, Z., Allen, N., Auerbach, R., Mahmood, F., Salakhutdinov, R.: Quantifying & modeling multi-modal interactions: An information decomposition framework. Advances in Neural Information Processing Systems **36** (2024)
27. Liang, P., Lyu, Y., Fan, X., Agarwal, A., Cheng, Y., Morency, L., Salakhutdinov, R.: Multizoo & multibench: a standardized toolkit for multimodal deep learning. The Journal of Machine Learning Research **24**(1), 11056–11062 (2023)
28. Liang, P., Zadeh, A., Morency, L.: Foundations and trends in multimodal machine learning: Principles, challenges, and open questions. arXiv preprint arXiv:2209.03430 (2022)
29. Liu, H., Li, C., Wu, Q., Lee, Y.J.: Visual instruction tuning (2023), https://arxiv.org/abs/2304.08485
30. Ma, X., Yong, S., Zheng, Z., Li, Q., Liang, Y., Zhu, S.C., Huang, S.: Sqa3d: Situated question answering in 3d scenes (2023), https://arxiv.org/abs/2210.07474
31. Manivannan, M., Nethrapalli, V., Cartwright, M.: Emotioncaps: Enhancing audio captioning through emotion-augmented data generation. arXiv preprint arXiv:2410.12028 (2024)
32. Manzoor, M.A., Albarri, S., Xian, Z., Meng, Z., Nakov, P., Liang, S.: Multimodality representation learning: A survey on evolution, pretraining and its applications (2024), https://arxiv.org/abs/2302.00389
33. Miech, A., Zhukov, D., Alayrac, J.B., Tapaswi, M., Laptev, I., Sivic, J.: Howto100m: Learning a text-video embedding by watching hundred million narrated video clips (2019), https://arxiv.org/abs/1906.03327
34. Nagrani, A., Chung, J., Zisserman, A.: Voxceleb: a large-scale speaker identification dataset. arXiv preprint arXiv:1706.08612 (2017)
35. Nan, K., Xie, R., Zhou, P., Fan, T., Yang, Z., Chen, Z., Li, X., Yang, J., Tai, Y.: Openvid-1m: A large-scale high-quality dataset for text-to-video generation (2024), https://arxiv.org/abs/2407.02371
36. Naveed, H., Khan, A., Qiu, S., Saqib, M., Anwar, S., Usman, M., Akhtar, N., Barnes, N., Mian, A.: A comprehensive overview of large language models. arXiv preprint arXiv:2307.06435 (2023)
37. Panayotov, V., Chen, G., Povey, D., Khudanpur, S.: Librispeech: an asr corpus based on public domain audio books. In: 2015 IEEE international conference on acoustics, speech and signal processing (ICASSP). pp. 5206–5210 (2015)
38. Pawłowski, M., Wróblewska, A., Sysko-Romańczuk, S.: Does a technique for building multimodal representation matter?–comparative analysis. arXiv preprint arXiv:2206.06367 (2022)
39. Poria, S., Hazarika, D., Majumder, N., Naik, G., Cambria, E., Mihalcea, R.: Meld: A multimodal multi-party dataset for emotion recognition in conversations. arXiv preprint arXiv:1810.02508 (2018)
40. Qiu, J., Zhu, J., Han, W., Kumar, A., Mittal, K., Jin, C., Yang, Z., Li, L., Wang, J., Zhao, D., Li, B., Wang, L.: Mmsum: A dataset for multimodal summarization and thumbnail generation of videos (2023), https://arxiv.org/abs/2306.04216

41. Rahate, A., Walambe, R., Ramanna, S., Kotecha, K.: Multimodal co-learning: Challenges, applications with datasets, recent advances and future directions. Information Fusion **81**, 203–239 (2022)

42. Roberts, J., Lüddecke, T., Sheikh, R., Han, K., Albanie, S.: Charting new territories: Exploring the geographic and geospatial capabilities of multimodal llms (2024), `https://arxiv.org/abs/2311.14656`

43. Schuhmann, C., Beaumont, R., Vencu, R., Gordon, C., Wightman, R., Cherti, M., Coombes, T., Katta, A., Mullis, C., Wortsman, M., Schramowski, P., Kundurthy, S., Crowson, K., Schmidt, L., Kaczmarczyk, R., Jitsev, J.: Laion-5b: An open large-scale dataset for training next generation image-text models (2022), `https://arxiv.org/abs/2210.08402`

44. Shridhar, M., Thomason, J., Gordon, D., Bisk, Y., Han, W., Mottaghi, R., Zettlemoyer, L., Fox, D.: Alfred: A benchmark for interpreting grounded instructions for everyday tasks (2020), `https://arxiv.org/abs/1912.01734`

45. Stan, G., Rohekar, R., Gurwicz, Y., Olson, M., Bhiwandiwalla, A., Aflalo, E., Wu, C., Duan, N., Tseng, S., Lal, V.: Lvlm-intrepret: An interpretability tool for large vision-language models. arXiv preprint arXiv:2404.03118 (2024)

46. Tanaka, R., Nishida, K., Nishida, K., Hasegawa, T., Saito, I., Saito, K.: Slidevqa: A dataset for document visual question answering on multiple images. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 37, pp. 13636–13645 (2023)

47. Tarsi, T., Adel, H., Metzen, J., Zhang, D., Finco, M., Friedrich, A.: Sciol and mulms-img: Introducing a large-scale multimodal scientific dataset and models for image-text tasks in the scientific domain. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. pp. 4560–4571 (2024)

48. Wang, Y., He, Y., Li, Y., Li, K., Yu, J., Ma, X., Li, X., Chen, G., Chen, X., Wang, Y., He, C.: Internvid: A large-scale video-text dataset for multimodal understanding and generation. arXiv preprint arXiv:2307.06942 (2023)

49. Wu, J., Gan, W., Chen, Z., Wan, S., Philip, S.: Multimodal large language models: A survey. In: 2023 IEEE International Conference on Big Data (BigData). pp. 2247–2256. IEEE (December 2023)

50. Yan, A., Yang, Z., Wu, J., Zhu, W., Yang, J., Li, L., Lin, K., Wang, J., McAuley, J., Gao, J., Wang, L.: List items one by one: A new data source and learning paradigm for multimodal llms (2024), `https://arxiv.org/abs/2404.16375`

51. Yang, D., Huang, S., Lu, C., Han, X., Zhang, H., Gao, Y., Hu, Y., Zhao, H.: Vript: A video is worth thousands of words (2024), `https://arxiv.org/abs/2406.06040`

52. Yin, S., Fu, C., Zhao, S., Li, K., Sun, X., Xu, T., Chen, E.: A survey on multimodal large language models. National Science Review **11**(12) (Nov 2024). `https://doi.org/10.1093/nsr/nwae403`, `http://dx.doi.org/10.1093/nsr/nwae403`

53. Yin, Z., Wang, J., Cao, J., Shi, Z., Liu, D., Li, M., Sheng, L., Bai, L., Huang, X., Wang, Z., Shao, J., Ouyang, W.: Lamm: Language-assisted multi-modal instruction-tuning dataset, framework, and benchmark (2023), `https://arxiv.org/abs/2306.06687`

54. Young, P., Lai, A., Hodosh, M., Hockenmaier, J.: From image descriptions to visual denotations: New similarity metrics for semantic inference over event descriptions. Transactions of the Association for Computational Linguistics **2**, 67–78 (2014)

55. Yuxuan, S., Yunlong, Z., Yixuan, S., Chenglu, Z., Zhongyi, S., Kai, Z., Jingxiong, L., Xingheng, L., Tao, L., Lin, Y.: Pathgen-1.6m: 1.6 million pathology image-text pairs generation through multi-agent collaboration. arXiv preprint arXiv:2407.00203 (2024)

56. Zadeh, A., Liang, P., Poria, S., Cambria, E., Morency, L.P.: Multimodal language analysis in the wild: Cmu-mosei dataset and interpretable dynamic fusion graph. pp. 2236–2246 (01 2018). `https://doi.org/10.18653/v1/P18-1208`

57. Zhang, C., Yang, Z., He, X., Deng, L.: Multimodal intelligence: Representation learning, information fusion, and applications. IEEE Journal of Selected Topics in Signal Processing **14**(3), 478–493 (2020)

58. Zhang, C., Zhang, Y., Shao, Q., Feng, J., Li, B., Lv, Y., Piao, X., Yin, B.: Bjtt: A large-scale multimodal dataset for traffic prediction. IEEE Transactions on Intelligent Transportation Systems (2024)

59. Zhang, D., Yu, Y., Li, C., Dong, J., Su, D., Chu, C., Yu, D.: Mm-llms: Recent advances in multimodal large language models. arXiv preprint arXiv:2401.13601 (2024)

60. Zhang, J., Huang, J., Jin, S., Lu, S.: Vision-language models for vision tasks: A survey (2024), `https://arxiv.org/abs/2304.00685`

61. Zhao, B., Wu, B., He, M., Huang, T.: Svit: Scaling up visual instruction tuning (2023), `https://arxiv.org/abs/2307.04087`

62. Zhao, J., Zhang, T., Hu, J., Liu, Y., Jin, Q., Wang, X., Li, H.: M3ed: Multi-modal multi-scene multi-label emotional dialogue database. arXiv preprint arXiv:2205.10237 (2022)

63. Zhao, R., Chen, H., Wang, W., Jiao, F., Do, X., Qin, C., Ding, B., Guo, X., Li, M., Li, X., Joty, S.: Retrieving multimodal information for augmented generation: A survey. arXiv preprint arXiv:2303.10868 (2023)