

# JOINTAVBENCH: A BENCHMARK FOR JOINT AUDIO-VISUAL REASONING EVALUATION

Jianghan Chao<sup>1</sup>, Jianzhang Gao<sup>1</sup>, Wenhui Tan<sup>1</sup>, Yuchong Sun<sup>1</sup>, Ruihua Song<sup>1,\*</sup>, Liyun Ru<sup>2</sup>

<sup>1</sup>Gaoling School of Artificial Intelligence, Renmin University of China

<sup>2</sup>Baichuan Inc

{chaojh, rsong}@ruc.edu.cn

## ABSTRACT

Understanding videos inherently requires reasoning over both visual and auditory information. To properly evaluate Omni-Large Language Models (Omni-LLMs), which are capable of processing multi-modal information, including vision and audio, an effective benchmark must comprehensively cover three key aspects: (1) multi-modal dependency (i.e., questions that cannot be answered using vision or audio alone), (2) diverse audio information types (e.g., speech, sound events), and (3) varying scene spans. However, existing datasets fall short in one or more of these dimensions, limiting strict and comprehensive evaluation. To address this gap, we introduce JointAVBench, a novel benchmark with strict audio-video correlation, spanning five cognitive dimensions, four audio information types (speech, sound events, music, vocal traits), and three scene spans (single-, cross-, and full-scene). Given the high cost of manual annotation, we propose an automated pipeline that leverages state-of-the-art vision-LLMs, audio-LLMs, and general-purpose LLMs to synthesize questions and answers that strictly require joint audio-visual understanding. We evaluate leading vision-only, audio-only, and Omni-LLMs on our dataset. Results show that even the best-performing Omni-LLM achieves an average accuracy of only 62.6%, outperforming uni-modal baselines but revealing substantial room for improvement, especially in cross-scene reasoning.

**Project page:** <https://jointavbench.github.io>

## 1 INTRODUCTION

Humans can understand videos and the real world by seamlessly perceiving and integrating both visual and auditory information across different scenes, where diverse audio signals (e.g. speech, sound, music, or even vocal traits) are used to complement the visual scene in analyses. As illustrated in Figure 1(a), for a multi-scene video understanding task, such as determining the order across a visual object in scene 3 with the dialogue in scene 1 and scene 23, requires complex joint audio-visual reasoning. This process involves recognizing visual and auditory cues, correlating them across distinct temporal and spatial contexts, and reasoning with the acquired relations. Toward the goal of artificial general intelligence, equipping multimodal large language models (MLLMs) with such joint audio-visual reasoning ability is paramount.

While newly developed Omni-LLMs (Team et al., 2023; Bai et al., 2025; Han et al., 2024; Wu et al., 2024a; Su et al., 2023) aim to process both audio and visual inputs jointly, progress is hindered by the lack of a comprehensive benchmark dedicated to evaluating this crucial capability. Existing benchmarks exhibit several limitations: some lack strict audio-visual correlation controls (Hong et al., 2025; Geng et al., 2024), others primarily focus on static images or simple videos (Li et al., 2024c; Gong et al., 2024), and mostly cover only a limited range of audio types (Yang et al., 2025). Furthermore, nearly all existing benchmarks neglect the complexities of multi-scene reasoning, which is a core component of human cognition.

To address this critical gap, we introduce JointAVBench, the first comprehensive benchmark for evaluating Omni-LLMs’ joint audio-visual reasoning capabilities. Our benchmark features a systematic taxonomy covering five cognitive dimensions (e.g., temporal, plot, and long-form reasoning),

\*Corresponding author.



Figure 1: **Examples of JointAVBench.** (a) asks a cross-scene plot-related question that needs the visual information in Scene 3 and the speech information in Scene 1 and Scene 23 to reason the right answer. (b) asks a single-scene emotion-related question that needs the visual information of the speaker and his vocal traits to answer.

four audio signal types (vocal traits, music, speech, and sound event), and three distinct scene spans (single-, cross-, and full-scene). These features enable us to construct 15 challenging tasks with strict audio-visual correlations, providing a unified and rigorous evaluation framework. For example, the task in Figure 1 (b) tests Speaker Emotion Recognition (SER), i.e., a single-scene, vocal traits involved, and emotion-related task.

To overcome the immense cost of manual annotation, we propose a semi-automated pipeline to generate high-quality question-answer (QA) pairs. This three-stage process first generates detailed multimodal captions, then synthesizes questions that strictly require joint audio-visual reasoning, and finally performs a rigorous quality assurance step to ensure data fidelity. We then use human labor to filter out unqualified data. This approach enables us to construct a high-quality multi-choice benchmark of 2,853 samples that are designed to probe complex reasoning abilities.

We conduct extensive experiments on JointAVBench to evaluate three types of MLLMs: Omni-LLMs, Video-LLMs, and Audio-LLMs. Our results demonstrate that current Omni-LLMs, such as Qwen-omni and Gemini-2.5 flash, significantly outperform their single-modal counterparts. However, our analysis also reveals that these models exhibit uneven capabilities across different audio types and suffer from a substantial performance degradation with increasing scene complexity. Our comprehensive assessment highlights critical limitations in current models' audio-visual reasoning capacities, posing the potential for future improvement.

In summary, our contributions can be summarized as follows:

- We introduce JointAVBench, the first-ever comprehensive benchmark to evaluate joint audio-visual reasoning capability across five cognitive dimensions, four audio types, and three scene complexities.
- We propose a novel three-stage semi-automated pipeline for generating high-quality QA pairs with strict audio-visual correlations while reducing annotation difficulties and costs.

Table 1: Comparison between our benchmark and previous ones. **Anno.:** the construction method, where A for automatic pipeline, A+M for pipeline involving manual inspection, and M for manual pipeline. **Modality:** the modality involved. V for video, I for image, and A for audio. **Aud. Type:** number of different audio signal types included in the dataset or benchmark. **AV Corr. Ratio:** the ratio of true audio-visual correlated questions, discussed in detail in Appendix C.1. † This metric is evaluated only on the three MCQ audio-visual tasks of AVUT (with two additional audio-visual open-ended tasks yet to be released).

Benchmark/Dataset	Avg. Duration	#QA	Anno. Method	Modality	#Tasks	#Audio Types	AV Corr. Ratio
<i>Video Benchmarks/Datasets</i>							
EgoSchema (Mangalam et al., 2023)	180s	5,063	A+M	V	1	-	0
Video-MME (Fu et al., 2025a)	1,017.9s	2,700	M	V	12	-	0
MVBench (Li et al., 2024b)	16.0s	4,000	A	V	20	-	0
LVBench (Wang et al., 2024b)	4,101s	1,549	M	V	6	-	0
MMBench-Video (Fang et al., 2024)	165.4s	1,998	M	V	26	-	0
<i>Audiovisual Benchmarks/Datasets</i>							
Music-AVQA (Li et al., 2022a)	60s	45,867	M	V&A	9	1	56.7%
OmniBench (Li et al., 2024c)	-	1,142	M	I&A	8	3	100%
AV-Odyssey (Gong et al., 2024)	-	4,555	M	V/I&A	26	3	100%
LongVALE (Geng et al., 2024)	235s	-	A+M	V&A	3	3	76.2%
AVUT (Yang et al., 2025)	67.8s	13,774	A+M	V&A	8	2	77.8% <sup>†</sup>
WorldSense (Hong et al., 2025)	141.1s	3,172	M	V&A	26	3	62.9%
<b>JointAVBench (ours)</b>	<b>97.2s</b>	<b>2,853</b>	<b>A+M</b>	<b>V&amp;A</b>	<b>15</b>	<b>4</b>	<b>100%</b>

- We provide a comprehensive evaluation of current MLLMs on JointAVBench, demonstrating their limitations and highlighting the importance of developing truly integrated audio-visual reasoning Omni-LLMs.

## 2 RELATED WORKS

### 2.1 MULTIMODAL LARGE LANGUAGE MODELS

The rise of Large Language Models (LLMs) has spurred interest in extending their capabilities beyond text to multimodal inputs (Bi et al., 2024; Achiam et al., 2023; Radford et al., 2018). Early efforts, such as (Radford et al., 2021; Hurst et al., 2024; Li et al., 2022b; 2023b), demonstrate effective fusion of visual and textual modalities for cross-modal understanding. Subsequent studies (Chu et al., 2023; Radford et al., 2023) expand this paradigm to incorporate audio-text integration and achieve significant improvements. Later advances in hardware and memory optimization enable video-text modeling in MLLMs (Team et al., 2023; Gao et al., 2017; Geng et al., 2022; Huang et al., 2024), spanning across various domains and achieving progress such as long video understanding (Wang et al., 2024a; Yuan et al., 2025; Chen et al., 2024a) and movie understanding (He et al., 2024; Song et al., 2024). Recent works (Chowdhury et al., 2025; Shu et al., 2023; Cheng et al., 2024; Tang et al., 2025b; Fu et al., 2024; 2025b; Lu et al., 2022; 2024; Su et al., 2023; Wu et al., 2024a; Han et al., 2024) focus on achieving human-like audio-visual joint reasoning ability by interleaving audio, video, and text. This requires datasets to contain QAs with strict audio-visual correlations. To facilitate the development of Omni-LLMs, we present JointAVBench to evaluate the models’ audiovisual joint reasoning ability with questions that are fully audio-visual correlated.

### 2.2 AUDIO-VISUAL BENCHMARKS

With the development of MLLMs, various benchmarks have been constructed for the evaluation of MLLMs’ comprehensive abilities (Wu et al., 2024b; Li et al., 2024a; 2023a; Yue et al., 2024; Zhang et al., 2024a; Sakshi et al., 2024; Liu et al., 2024; Li et al., 2024b). Early datasets or benchmarks, such as AVQA (Yang et al., 2022), Music-AVQA (Li et al., 2022a), and AVInstruct (Ye et al., 2024), only focus on certain types of audio signals and lack strict audio-visual correlation. Subsequent works such as Omni-bench (Li et al., 2024c) and AV-Odyssey (Gong et al., 2024) consist primarily of only image and audio, lacking the evaluation of videos. The recent WorldSense (Hong et al., 2025) has delved into the problem. However, it lacks strict audio-visual correlation and emphasizes the evaluation of visual tasks. These datasets cannot capture the complex and interleaved auditory and scene details

Table 2: Task categories of our designed taxonomy. In audio signal type, we use SPE for speech, VOT for vocal traits, SEV for sound event, and MUS for music.

Scene Type	Cognitive Dimension	Audio Signal Type	Task Name	Task Code
Single	Temporal	SPE	Speech-based Timepoint Localization	STL
		SPE	Vision-Speech Sequence Recognition	VSSR
	Spatial	VOT	Speaker Spatial Localization	SPL
		SEV	Sounding Object Grounding	SOOG
		SEV	Sound Event Recognition	SOER
	Emotion	VOT	Speaker Emotion Recognition	SPER
MUS		Musical Tone Inference	MPTI	
Multiple	Long-form	SPE	Cross-scene Association	CSA
		SPE, VOT	Multi-plot Ordering	MPO
	Plot	SPE	Plot Development Prediction	PDP
		SEV, MUS	Audio Function Analysis	AFA
		SPE	Plot Temporal Grounding	PTG
Full	Long-form	SPE, VOT, SEV, MUS	Audio-Visual Detail Memory	AVDM
	Emotion	MUS	Musical Emotion Shift Inference	MESI
	Plot	SPE	Character Relationship Inference	CRI

in video (such as the details in Figure 1). In contrast, our proposed JointAVBench focuses on the evaluation of diverse audio signal types and multilevel scenes, aiming to conduct a comprehensive and systematic assessment of current MLLMs for joint audio-visual understanding.

### 3 JOINTAVBENCH

We propose JointAVBench, a benchmark for evaluating Omni-LLMs’ joint audio-visual reasoning ability. This section will first detail the benchmark’s core requirements, then the carefully designed data generation pipeline, including (i) omni-caption generation, (ii) QA pair creation, and (iii) rigorous quality control, with statistics provided at the end of this section.

#### 3.1 BENCHMARK REQUIREMENTS

The benchmark construction adheres to three fundamental requirements, ensuring comprehensive evaluation of joint audio-visual reasoning capabilities.

**Strict Audio-Visual Correlation.** We design a hierarchical taxonomy comprising 15 tasks (detailed in Table 2), ensuring that each task requires the integration of both visual and audio information to generate answers.

**High-quality Video Source.** Movie scenes are a natural source of extensive and diverse multimodal data. For our benchmark, we leverage the Short-Films 20K (SF20K) (Ghermi et al., 2024) dataset, which comprises 1,072 professionally produced movies rich in narrative and balanced audiovisual features. We then remove unavailable or grayscale videos, retaining 1,046 films that can be used to construct our benchmark.

**Multi-dimensional Task Taxonomy.** To ensure a comprehensive and fine-grained evaluation, we categorize our tasks along three key dimensions:

- **Cognitive Dimension:** Derived from a systematic analysis of previous studies (Fu et al., 2025a; Hong et al., 2025), this dimension assesses core cognitive abilities essential for video understanding. We define 5 types of cognitive dimensions: temporal, spatial, emotional, plot, and long-form.
- **Audio Types:** This dimension enables a comprehensive evaluation of audio understanding capabilities across all audio signal types. We divide audio into four types of signals: speech, vocal traits, sound event, and music.
- **Scene Complexity:** This dimension evaluates model performance across videos with varying temporal characteristics, using different scene types to quantify temporal information. We define three types of scene complexity: single-scene, multi-scene, and full-scene.

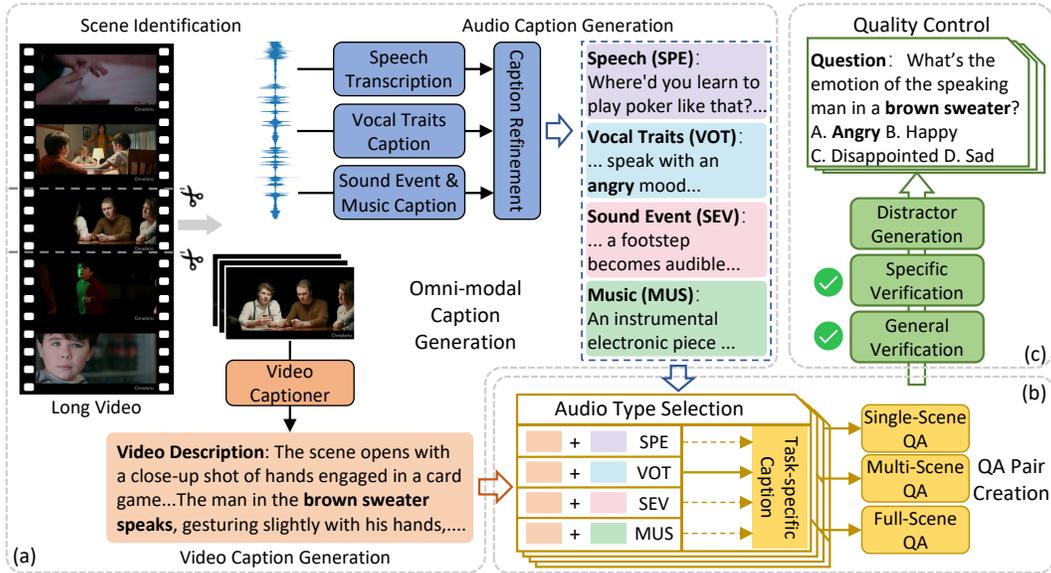


Figure 2: **Pipeline for JointAVBench.** Our construction pipeline is three-fold: (a) Omni-modal caption generation, (b) QA pair creation, and (c) Quality control.

### 3.2 BENCHMARK CONSTRUCTION

Our dataset construction pipeline is illustrated in Figure 2, which adopts a semi-automated process capable of handling diverse modality characteristics. More details can be found in the appendix.

#### 3.2.1 STAGE 1: OMNI-MODAL CAPTION GENERATION

**Scene Identification.** We first split the video into scenes with semantic consistency. Specifically, we follow the procedure in Panda-70m (Chen et al., 2024b) to divide long videos into distinct scenes with PySceneDetect<sup>1</sup> and then merge scenes with high semantic similarity, ensuring in-scene consistency. These segmented scenes retain considerable length, enabling us to capture richer contextual information within each scene.

**Video Caption Generation.** After scene identification, we directly generate visual descriptions for all segmented scenes, ensuring that static features (e.g. in-scene objects and characters) and dynamic features (e.g. transitions between shots and movements of characters) are well captured.

**Audio Caption Generation.** To ensure the diversity of audio types, we follow the requirements to generate captions for each audio type as shown in Figure 2. Notably, we observe that existing audio models have limitations in distinguishing between sound event and music, and therefore generate their captions simultaneously. Subsequently, we refine the audio captions by addressing the hallucination in the caption and separating the sound event caption and music caption using different LLM judges.

#### 3.2.2 STAGE 2: QA PAIR CREATION

To create QA pairs with strict audiovisual correlation, we design various question templates for tasks that LLM cannot easily understand (temporal, plot tasks that require complex audio-visual relation), while leaving other tasks to be curated by LLM to ensure question diversity (general tasks such as Character Relationship Inference). Additionally, when providing cross-modal descriptions, we strictly adhere to each task’s modality and scene requirements by inputting only the required modality descriptions from the designated scenes. For example, when generating data for the task Speaker Spatial Localization, we provide video captions along with vocal traits descriptions from only one scene. This procedure can eliminate possible interference from extraneous modalities and scenes.

<sup>1</sup><https://www.scenesdetect.com/>

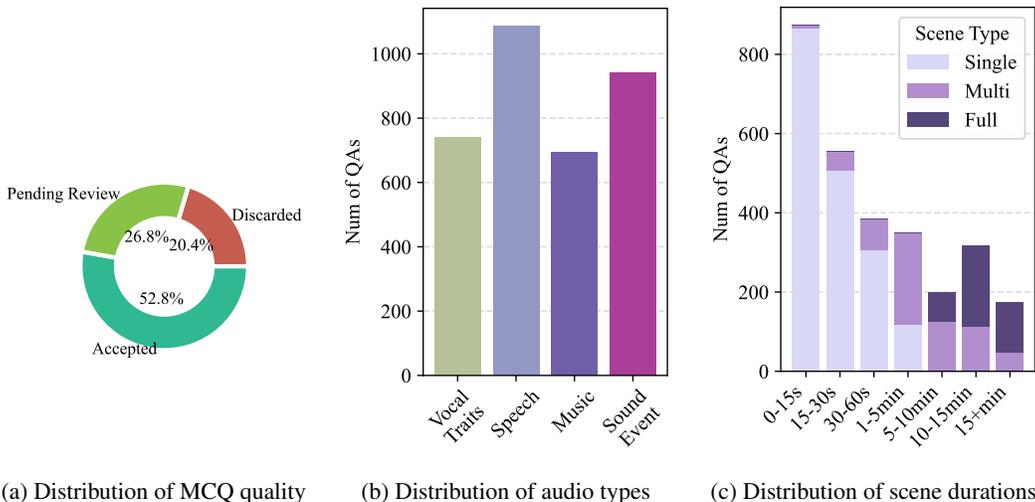


Figure 3: Statistics of JointAVBench.

### 3.2.3 STAGE 3: QUALITY CONTROL

We implement a multi-stage quality control process to address issues identified in the collected 9,109 QA pairs, such as mismatched question-answer pairs and redundant information. This process employs a general-to-specific verification strategy, where we guide models to use a chain-of-thought approach for step-by-step data filtering.

**General Verification.** We validate all QA pairs to ensure they meet fundamental standards. This includes a **Modality Check** to confirm that each QA pair necessitates both audio and video information, and a **Logic Check** to verify that answers are directly derivable from the question’s context. For instance, a question like “*What is the emotion of the adult male speaker?*” will be discarded if the audio contains only one male speaker, as the answer can be inferred from a single modality.

**Specific Verification.** This stage focuses on task-specific validations. We design the following three specific checks based on QA’s task: 1) **Sequence Check** to ensure the correct element order for sequence-based tasks; 2) **Ambiguity Check** to filter out overly generic QA pairs for complex reasoning tasks (e.g., “*What makes the door closing sound?*” with “*door closing*” as the answer); 3) **Audio Signal Type Check** to confirm that the required auditory information cannot be deduced from visual information for sound event and music.

**Distractor Generation.** For each verified QA pair, we craft three plausible but incorrect distractors to create challenging multiple-choice questions. These distractors incorporate diverse misdirections, such as replacing the sound source or confusing details.

## 3.3 HUMAN VERIFICATION

From the automated three-stage generation process, we obtain 3,974 MCQs. To ensure their quality and factual accuracy, we conducted a rigorous human verification process, and the results are illustrated in Figure 3a. Specifically, a team of human annotators rated QAs based on four key criteria: (i) answer correctness, (ii) information correctness, (iii) audio-visual dependency, and (iv) question difficulty. Based on these ratings, we categorize these data into three subsets: (1) Accepted: QAs that pass the answer correctness check and score highly on all other criteria, which are directly retained in the final dataset; (2) Pending Review: QAs that pass the answer correctness check but receive lower ratings on one or more additional criteria, which are subject to further selection according to their ratings; and (3) Discarded: QAs that fail the answer correctness check and are removed from the dataset. In total, we retained 2,853 QAs, achieving a data retention rate of 71.8%, which demonstrates that our automatic pipeline is highly effective at generating data of sufficient quality.

Table 3: Evaluation results of three types of mainstream MLLMs. We evaluate the performance of Omni-LLMs, Video-LLMs and Audio-LLMs on JointAVBench to provide a comprehensive analysis. † For neatness, we use short names of models, where v-SALMONN represents video-SALMONN series.

Model†	Size	STL	SPL	SOOG	SOER	SPER	MPTI	VSSR	CSA	MPO	PTG	AFA	PDP	AVDM	MESI	CRI	Avg
<i>Omni-LLMs</i>																	
Gemini2.5-Pro	-	<b>73.0</b>	<b>59.4</b>	60.8	68.9	35.2	68.1	76.5	43.8	<b>66.0</b>	<b>60.7</b>	65.5	45.7	<b>75.5</b>	66.1	81.9	<b>62.6</b>
Qwen3-Omni	30B	71.1	43.4	<b>73.8</b>	<b>78.4</b>	<b>35.7</b>	<b>80.3</b>	75.7	42.1	45.2	30.9	59.7	47.3	61.8	69.2	<b>84.0</b>	62.1
Qwen2.5-Omni	7B	71.3	35.3	59.8	72.3	30.6	63.4	<b>77.6</b>	51.2	40.4	20.8	<b>69.9</b>	47.3	47.3	69.9	70.3	56.2
Gemini2.5-Flash	-	62.6	43.8	55.7	68.6	23.3	59.1	42.6	39.2	46.2	25.3	65.3	<b>48.4</b>	64.8	67.7	81.3	52.6
v-SALMONN-2+	7B	40.9	23.2	58.4	59.1	17.6	72.8	52.6	30.8	40.4	22.1	55.7	45.2	39.0	70.7	61.6	47.3
VideoLLaMA2	7B	20.9	38.8	56.1	67.0	29.6	47.5	48.5	24.0	35.3	30.9	63.6	36.6	38.0	61.7	58.8	46.6
OneLLM	7B	33.0	44.2	45.6	37.8	29.9	29.7	33.9	<b>55.4</b>	31.7	32.9	46.6	44.1	34.5	34.6	50.3	38.5
v-SALMONN-o1	7B	32.2	30.0	35.1	43.5	14.0	44.7	32.0	25.6	20.0	36.2	55.4	30.1	35.5	66.9	58.8	37.3
v-SALMONN	7B	52.2	25.1	37.8	52.2	19.0	33.3	33.5	30.5	31.7	26.1	48.9	26.9	24.2	37.9	50.3	35.8
AVicuna	7B	31.9	29.3	35.1	38.8	16.6	31.2	21.4	25.0	21.2	30.9	43.7	30.1	27.6	29.5	44.3	30.6
<i>Video-LLMs</i>																	
InternVL-2.5	8B	28.7	37.9	59.8	71.1	23.6	64.1	52.2	42.5	44.2	27.5	63.6	41.9	50.0	68.4	68.3	51.3
VideoLLaMA3	7B	43.5	41.1	58.8	55.8	17.9	69.2	50.0	34.7	43.3	33.6	61.9	40.9	51.8	73.7	64.8	49.9
Qwen2.5-VL	7B	33.9	38.8	55.3	59.3	22.9	57.2	47.2	31.7	40.4	32.2	62.5	39.8	40.7	62.9	61.6	47.1
LLaVA-Video	7B	37.4	33.0	48.6	64.7	10.0	68.5	53.9	27.3	43.3	30.9	51.1	36.6	46.4	<b>76.7</b>	61.8	47.0
GPT-4o	-	30.4	34.8	55.7	69.7	11.6	53.6	24.8	40.5	13.5	14.1	51.7	47.3	50.9	56.4	70.9	43.3
<i>Audio-LLMs</i>																	
Kimi-Audio	7B	56.5	21.9	48.6	61.7	32.9	53.3	34.3	38.0	33.0	26.2	65.3	38.7	40.2	56.1	69.5	45.9
Qwen2-Audio	7B	54.1	24.3	39.5	54.3	34.6	40.0	34.3	33.0	32.7	27.6	55.0	29.8	32.9	46.6	58.1	40.0

### 3.4 BENCHMARK STATISTICS

JointAVBench consists of 2,853 high-quality, manually verified MCQs spanning all scene levels and audio types, with an average duration of 97.2s (Table 1). A detailed statistical analysis of the benchmark is presented in Figure 3. The number of QA pairs is balanced across diverse audio signal types (Figure 3b), showcasing our benchmark’s comprehensive coverage. Moreover, our dataset spans a wide range of video durations (Figure 3c), with single-scene, multi-scene, and full-scene tasks mainly comprising videos of less than 1 min, 1-10 min, and over 10 min, respectively.

## 4 EXPERIMENT

This section first demonstrates a comprehensive evaluation of mainstream MLLMs on our proposed benchmark, and then key factors that influence performance to provide valuable insights for future Omni-LLMs.

### 4.1 EXPERIMENT SETUP

**Models.** Our experiments are conducted on a diverse set of mainstream MLLMs. To comprehensively evaluate their joint audio-visual reasoning capability across different modalities, we categorize them into three groups: (i) Omni-modal LLMs: Qwen3-Omni (Xu et al., 2025b), Qwen2.5-Omni (Xu et al., 2025a), VideoLLaMA2 (Cheng et al., 2024), video-SALMONN-2+ (Tang et al., 2025a), video-SALMONN-o1 (Sun et al., 2025), video-SALMONN (Sun et al., 2024), OneLLM (Han et al., 2024), AVicuna (Tang et al., 2025b), Gemini2.5-Pro (Comanici et al., 2025), and Gemini2.5-Flash (Comanici et al., 2025); (ii) Video-LLMs: Qwen2.5-VL (Bai et al., 2025), LLaVA-Video (Zhang et al., 2024b), Video-LLaMA3 (Zhang et al., 2025), InternVL2.5 (Chen et al., 2024c), and GPT-4o (Hurst et al., 2024); and (iii) Audio-LLMs: Kimi-Audio (Ding et al., 2025) and Qwen2-Audio (Chu et al., 2024).

**Metrics and Experiment Settings.** To achieve evaluation consistency, we follow previous works (Hong et al., 2025; Fu et al., 2025a) and use accuracy as the evaluation metric. For a fair evaluation, we adopt the following protocols for all experiments. For open-source models, we use their official codebase with default configurations, while for closed-source models, we use their official APIs while keeping the configuration by default. To maintain comparability, we select open-source models with comparable 7B parameter sizes and enforce a unified sampling of 32 frames

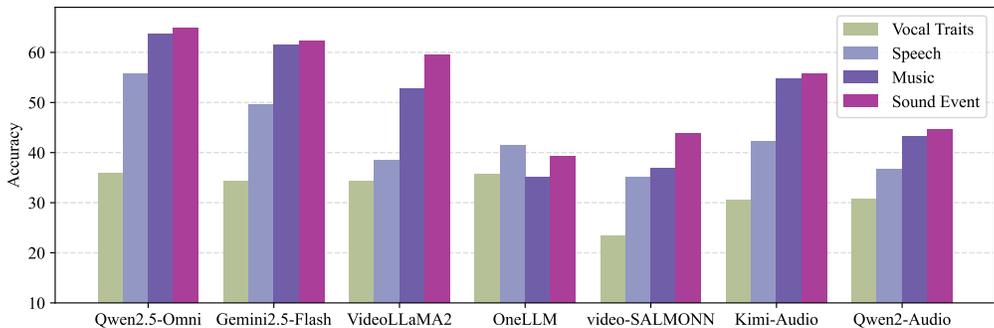


Figure 4: Results on JointAVBench across different audio types.

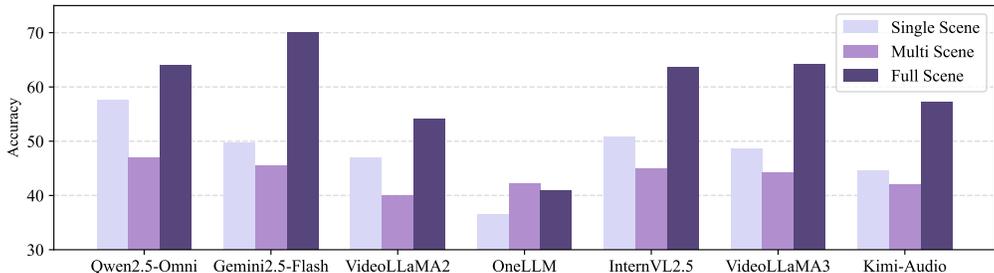


Figure 5: Results on JointAVBench across different scene types.

across all models. We also ensure that the text input to all models is limited to the question text, without any additional contextual information.

#### 4.2 RESULTS AND FINDINGS

**Overall Performance.** Our results, summarized in Table 3, reveal that current mainstream MLLMs perform sub-optimally on our benchmark, with the best performing model having an average accuracy of only 62.6%. This suggests a significant gap in their ability to process omni-modal information. Importantly, Omni-LLMs consistently outperform Video-LLMs and Audio-LLMs, highlighting the critical role of native modality integration. For instance, Qwen2.5-Omni significantly improves upon Intern-VL across most tasks.

**Breakdown Findings.** To gain a deeper understanding, we analyze model performance across various task categories and have the following observations.

**1) Models perform unevenly across different audio types, failing on tasks requiring vocal traits and speech.** We find a significant performance gap among different audio types (Figure 4). Models excel at tasks involving sound events and music, likely due to their stronger visual correspondence (where audio often matches visible objects or atmosphere). However, they struggle with more abstract audio, such as speech and vocal traits. This is likely due to a lack of training data focused on vocal traits, as most audio-visual datasets (Li et al., 2022a; Yang et al., 2022) overlook information like emotion and gender, leading to tasks like SPL, SPER, and MPO being the worst-performing overall.

**2) Multi-scene tasks usually yield worse results compared to single-scene tasks, while full-scene tasks often achieve better results.** Figure 5 shows the pronounced impact of scene complexity. Models perform well on single-scene tasks, particularly those requiring speech-based reasoning like STL and VSSR. This is likely because single scenes offer stable visual contexts and limited speech content, simplifying cross-modal correspondence. Conversely, multi-scene tasks requiring speech, such as MPO and PTG, yield worse performance, as they demand more complex processing of diverse scenes and cross-scene connections. Interestingly, while most models struggle with multi-scene tasks,

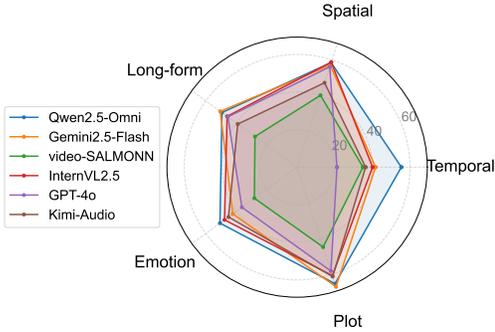


Figure 6: Results on JointAVBench across 5 cognitive dimensions.

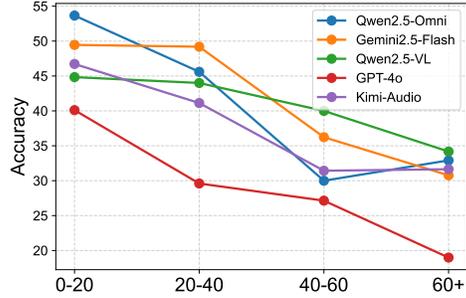


Figure 7: Results on JointAVBench with different number of scenes.

Table 4: Evaluation results of open-source omni-LLMs with different modality utilization. † For neatness, we use short names of models, where Qwen2.5 represents Qwen2.5-Omni, ViLLaMA2 represents VideoLLaMA2, and v-SALMONN represents video-SALMONN.

Model†	Modality	$N_o$	$N_u$	Avg	STL	SPL	SOOG	SOER	SPER	MPTI	VSSR	CSA	MPO	PTG	AFA	PDP	AVDM	MESI	CRI
Qwen2.5	A+V	8	1	56.2	71.3	35.3	59.8	72.3	30.6	63.4	77.6	51.2	40.4	20.8	69.9	47.3	47.3	69.9	70.3
	V			49.3	38.3	33.0	57.8	64.6	18.6	60.1	50.4	43.8	40.4	24.2	68.8	37.6	48.2	70.7	68.5
	A			52.3	71.3	23.9	64.4	66.5	34.2	56.2	57.0	41.3	45.2	18.1	69.3	41.9	49.1	64.7	68.3
VidLLaMA2	A+V	6	3	46.6	20.9	38.8	56.1	67.0	29.6	47.5	48.5	24.0	35.3	30.9	63.6	36.6	38.0	61.7	58.8
	V			46.6	27.8	37.9	56.8	67.4	27.9	49.3	43.9	28.9	46.0	26.2	62.5	38.7	38.0	62.4	52.7
	A			41.4	24.3	35.7	56.1	61.4	18.3	37.3	32.6	35.5	35.6	30.2	60.2	29.0	31.8	50.4	56.4
OneLLM	A+V	8	3	38.5	33.0	44.2	45.6	37.8	29.9	29.7	33.9	55.4	31.7	32.9	46.6	44.1	34.5	34.6	50.3
	V			32.7	27.8	35.7	37.8	28.6	15.3	28.6	31.7	31.4	33.7	32.2	38.6	50.5	34.5	35.3	53.3
	A			38.5	28.7	43.8	43.6	38.9	31.6	28.6	27.4	47.9	23.1	30.2	41.5	62.4	45.5	39.8	60.0
v-SALMONN	A+V	5	4	35.8	52.2	25.1	37.8	52.2	19.0	33.3	33.5	30.5	31.7	26.1	48.9	26.9	24.2	37.9	50.3
	V			34.8	23.5	24.3	44.4	49.0	16.8	35.5	29.1	28.3	36.5	26.2	47.2	40.9	31.3	38.3	44.1
	A			35.7	53.9	22.6	39.0	50.6	21.4	34.9	32.6	36.4	28.7	25.7	44.3	31.2	23.8	39.7	44.2

they perform better on full-scene tasks, which focus on global narratives rather than fine-grained details. This highlights that improving models’ cross-scene reasoning capabilities is needed.

**3) Omni-models perform worse on emotional and spatial tasks than single-modal models.** As demonstrated in Figure 6, while Omni-LLMs generally perform best in 11 of our 15 tasks, they surprisingly fall behind single-modality models on emotion-based tasks. This suggests that in some cases, single-modality models can better focus on emotion cues without the distraction of additional modalities. Furthermore, Omni-LLMs perform poorly on spatial tasks like SOOG and SOER, even falling behind Video-LLMs. This is likely because models primarily rely on spatial information from video and fail to effectively integrate complementary audio cues. This finding highlights a critical limitation in current models’ ability to perform true audio-visual spatial reasoning.

**4) Increased scene number leads to models’ performance degradation on multi-scene tasks.** Our analysis reveals that increasing the number of scenes adversely affects multi-scene task performance (Figure 7), with accuracy dropping sharply by approximately 20% from 0-20 to over 60 scenes. This underscores the uneven performance of current MLLMs across diverse scenes, highlighting a critical area for improvement.

**5) Omni-modal Models Demonstrate Effective Modality Fusion.** We quantify the effectiveness of joint reasoning by defining  $N_o$  and  $N_u$  as the number of tasks where a model’s audio-visual (A+V) performance is better and worse than its single-modality baseline, respectively, in Table 4. For all models,  $N_o$  significantly outweighs  $N_u$ , confirming that integrating audio and video fundamentally enhances overall reasoning capability. Furthermore, as a model’s overall performance increases, its  $N_o$  count rises while its  $N_u$  count falls. This pattern indicates that more advanced models are better adept at modality fusion. For instance, while early Omni-LLMs (Cheng et al., 2024; Han et al., 2024; Sun et al., 2024) show only marginal gains over single-modality baselines, Qwen2.5-Omni’s

dual-modality performance significantly surpasses its single-modality results. This confirms that true joint reasoning is a hallmark of mature omni-modal architectures.

## 5 CONCLUSION

In this paper, we propose JointAVBench, a comprehensive benchmark for evaluating joint audio-visual reasoning, distinguished by a hierarchical taxonomy and a high-quality, automated generation pipeline. Each question in JointAVBench is meticulously designed to necessitate the integrated understanding of both visual and a specific type of audio input. We further ensure benchmark quality through human verification. Our extensive experiments reveal that even the best-performing models achieve an accuracy of only 62.6%, underscoring the substantial need for more powerful omni-modal models with enhanced audio-visual fusion capabilities.

## ACKNOWLEDGMENTS

This work is supported by the National Natural Science Foundation of China (No. 62276268) and Baichuan Inc. We also gratefully acknowledge the insightful comments and suggestions provided by the anonymous reviewers.

## REFERENCES

- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*, 2023.
- Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibao Song, Kai Dang, Peng Wang, Shijie Wang, Jun Tang, et al. Qwen2. 5-vl technical report. *arXiv preprint arXiv:2502.13923*, 2025.
- Xiao Bi, Deli Chen, Guanting Chen, Shanhuang Chen, Damai Dai, Chengqi Deng, Honghui Ding, Kai Dong, Qiu Shi Du, Zhe Fu, et al. Deepseek llm: Scaling open-source language models with longtermism. *arXiv preprint arXiv:2401.02954*, 2024.
- Lin Chen, Jinsong Li, Xiaoyi Dong, Pan Zhang, Conghui He, Jiaqi Wang, Feng Zhao, and Dahua Lin. Sharegpt4v: Improving large multi-modal models with better captions. In *European Conference on Computer Vision*, pp. 370–387. Springer, 2024a.
- Tsai-Shien Chen, Aliaksandr Siarohin, Willi Menapace, Ekaterina Deyneka, Hsiang-wei Chao, Byung Eun Jeon, Yuwei Fang, Hsin-Ying Lee, Jian Ren, Ming-Hsuan Yang, et al. Panda-70m: Captioning 70m videos with multiple cross-modality teachers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 13320–13331, 2024b.
- Zhe Chen, Weiyun Wang, Yue Cao, Yangzhou Liu, Zhangwei Gao, Erfei Cui, Jinguo Zhu, Shenglong Ye, Hao Tian, Zhaoyang Liu, et al. Expanding performance boundaries of open-source multimodal models with model, data, and test-time scaling. *arXiv preprint arXiv:2412.05271*, 2024c.
- Zesen Cheng, Sicong Leng, Hang Zhang, Yifei Xin, Xin Li, Guanzheng Chen, Yongxin Zhu, Wenqi Zhang, Ziyang Luo, Deli Zhao, et al. Videollama 2: Advancing spatial-temporal modeling and audio understanding in video-llms. *arXiv preprint arXiv:2406.07476*, 2024.
- Sanjoy Chowdhury, Hanan Gani, Nishit Anand, Sayan Nag, Ruohan Gao, Mohamed Elhoseiny, Salman Khan, and Dinesh Manocha. Aurelia: Test-time reasoning distillation in audio-visual llms. *arXiv preprint arXiv:2503.23219*, 2025.
- Yunfei Chu, Jin Xu, Xiaohuan Zhou, Qian Yang, Shiliang Zhang, Zhijie Yan, Chang Zhou, and Jingren Zhou. Qwen-audio: Advancing universal audio understanding via unified large-scale audio-language models. *arXiv preprint arXiv:2311.07919*, 2023.
- Yunfei Chu, Jin Xu, Qian Yang, Haojie Wei, Xipin Wei, Zhifang Guo, Yichong Leng, Yuanjun Lv, Jinzheng He, Junyang Lin, et al. Qwen2-audio technical report. *arXiv preprint arXiv:2407.10759*, 2024.
- Gheorghe Comanici, Eric Bieber, Mike Schaekermann, Ice Pasupat, Noveen Sachdeva, Inderjit Dhillon, Marcel Blistein, Ori Ram, Dan Zhang, Evan Rosen, et al. Gemini 2.5: Pushing the frontier with advanced reasoning, multimodality, long context, and next generation agentic capabilities. *arXiv preprint arXiv:2507.06261*, 2025.
- Ding Ding, Zeqian Ju, Yichong Leng, Songxiang Liu, Tong Liu, Zeyu Shang, Kai Shen, Wei Song, Xu Tan, Heyi Tang, et al. Kimi-audio technical report. *arXiv preprint arXiv:2504.18425*, 2025.
- Xinyu Fang, Kangrui Mao, Haodong Duan, Xiangyu Zhao, Yining Li, Dahua Lin, and Kai Chen. Mmbench-video: A long-form multi-shot benchmark for holistic video understanding. *Advances in Neural Information Processing Systems*, 37:89098–89124, 2024.
- Chaoyou Fu, Haojia Lin, Zuwei Long, Yunhang Shen, Meng Zhao, Yifan Zhang, Shaoqi Dong, Xiong Wang, Di Yin, Long Ma, et al. Vita: Towards open-source interactive omni multimodal llm. *arXiv preprint arXiv:2408.05211*, 2024.
- Chaoyou Fu, Yuhang Dai, Yongdong Luo, Lei Li, Shuhuai Ren, Renrui Zhang, Zihan Wang, Chenyu Zhou, Yunhang Shen, Mengdan Zhang, et al. Video-mme: The first-ever comprehensive evaluation benchmark of multi-modal llms in video analysis. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pp. 24108–24118, 2025a.

- Chaoyou Fu, Haojia Lin, Xiong Wang, Yi-Fan Zhang, Yunhang Shen, Xiaoyu Liu, Yangze Li, Zuwei Long, Heting Gao, Ke Li, et al. Vita-1.5: Towards gpt-4o level real-time vision and speech interaction. *arXiv preprint arXiv:2501.01957*, 2025b.
- Jiyang Gao, Chen Sun, Zhenheng Yang, and Ram Nevatia. Tall: Temporal activity localization via language query. In *Proceedings of the IEEE international conference on computer vision*, pp. 5267–5275, 2017.
- Tiantian Geng, Feng Zheng, Xiaorong Hou, Ke Lu, Guo-Jun Qi, and Ling Shao. Spatial-temporal pyramid graph reasoning for action recognition. *IEEE Transactions on Image Processing*, 31: 5484–5497, 2022.
- Tiantian Geng, Jinrui Zhang, Qingni Wang, Teng Wang, Jinming Duan, and Feng Zheng. Longvale: Vision-audio-language-event benchmark towards time-aware omni-modal perception of long videos. *arXiv preprint arXiv:2411.19772*, 2024.
- Ridouane Ghermi, Xi Wang, Vicky Kalogeiton, and Ivan Laptev. Short film dataset (sfd): A benchmark for story-level video understanding. *arXiv preprint arXiv:2406.10221*, 2024.
- Kaixiong Gong, Kaituo Feng, Bohao Li, Yibing Wang, Mofan Cheng, Shijia Yang, Jiaming Han, Benyou Wang, Yutong Bai, Zhuoran Yang, et al. Av-odyssey bench: Can your multimodal llms really understand audio-visual information? *arXiv preprint arXiv:2412.02611*, 2024.
- Jiaming Han, Kaixiong Gong, Yiyuan Zhang, Jiaqi Wang, Kaipeng Zhang, Dahua Lin, Yu Qiao, Peng Gao, and Xiangyu Yue. Onellm: One framework to align all modalities with language. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 26584–26595, 2024.
- Yichen He, Yuan Lin, Jianchao Wu, Hanchong Zhang, Yuchen Zhang, and Ruicheng Le. Storyteller: Improving long video description through global audio-visual character identification. *arXiv preprint arXiv:2411.07076*, 2024.
- Jack Hong, Shilin Yan, Jiayin Cai, Xiaolong Jiang, Yao Hu, and Weidi Xie. Worldsense: Evaluating real-world omnimodal understanding for multimodal llms. *arXiv preprint arXiv:2502.04326*, 2025.
- Bin Huang, Xin Wang, Hong Chen, Zihan Song, and Wenwu Zhu. Vtimellm: Empower llm to grasp video moments. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 14271–14280, 2024.
- Aaron Hurst, Adam Lerer, Adam P Goucher, Adam Perelman, Aditya Ramesh, Aidan Clark, AJ Ostrow, Akila Welihinda, Alan Hayes, Alec Radford, et al. Gpt-4o system card. *arXiv preprint arXiv:2410.21276*, 2024.
- J Richard Landis and Gary G Koch. The measurement of observer agreement for categorical data. *biometrics*, pp. 159–174, 1977.
- 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*, 2023a.
- Bohao Li, Yuying Ge, Yixiao Ge, Guangzhi Wang, Rui Wang, Ruimao Zhang, and Ying Shan. Seed-bench: Benchmarking multimodal large language models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 13299–13308, 2024a.
- Guangyao Li, Yake Wei, Yapeng Tian, Chenliang Xu, Ji-Rong Wen, and Di Hu. Learning to answer questions in dynamic audio-visual scenarios. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 19108–19118, 2022a.
- Junnan Li, Dongxu Li, Caiming Xiong, and Steven Hoi. Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In *International conference on machine learning*, pp. 12888–12900. PMLR, 2022b.

- Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. In *International conference on machine learning*, pp. 19730–19742. PMLR, 2023b.
- Kunchang Li, Yali Wang, Yinan He, Yizhuo Li, Yi Wang, Yi Liu, Zun Wang, Jilan Xu, Guo Chen, Ping Luo, et al. Mvbench: A comprehensive multi-modal video understanding benchmark. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 22195–22206, 2024b.
- Yizhi Li, Ge Zhang, Yinghao Ma, Ruibin Yuan, Kang Zhu, Hangyu Guo, Yiming Liang, Jiaheng Liu, Zekun Wang, Jian Yang, et al. Omnibench: Towards the future of universal omni-language models. *arXiv preprint arXiv:2409.15272*, 2024c.
- Yuanxin Liu, Shicheng Li, Yi Liu, Yuxiang Wang, Shuhuai Ren, Lei Li, Sishuo Chen, Xu Sun, and Lu Hou. Tempcompass: Do video llms really understand videos? *arXiv preprint arXiv:2403.00476*, 2024.
- Jiasen Lu, Christopher Clark, Rowan Zellers, Roozbeh Mottaghi, and Aniruddha Kembhavi. Unified-io: A unified model for vision, language, and multi-modal tasks. *arXiv preprint arXiv:2206.08916*, 2022.
- Jiasen Lu, Christopher Clark, Sangho Lee, Zichen Zhang, Savya Khosla, Ryan Marten, Derek Hoiem, and Aniruddha Kembhavi. Unified-io 2: Scaling autoregressive multimodal models with vision language audio and action. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 26439–26455, 2024.
- Kartikeya Mangalam, Raiymbek Akshulakov, and Jitendra Malik. Egoschema: A diagnostic benchmark for very long-form video language understanding. *Advances in Neural Information Processing Systems*, 36:46212–46244, 2023.
- Alec Radford, Karthik Narasimhan, Tim Salimans, Ilya Sutskever, et al. Improving language understanding by generative pre-training. 2018.
- 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 *International conference on machine learning*, pp. 8748–8763. PMLR, 2021.
- Alec Radford, Jong Wook Kim, Tao Xu, Greg Brockman, Christine McLeavey, and Ilya Sutskever. Robust speech recognition via large-scale weak supervision. In *International conference on machine learning*, pp. 28492–28518. PMLR, 2023.
- S Sakshi, Utkarsh Tyagi, Sonal Kumar, Ashish Seth, Ramaneswaran Selvakumar, Oriol Nieto, Ramani Duraiswami, Sreyan Ghosh, and Dinesh Manocha. Mmau: A massive multi-task audio understanding and reasoning benchmark. *arXiv preprint arXiv:2410.19168*, 2024.
- Fangxun Shu, Lei Zhang, Hao Jiang, and Cihang Xie. Audio-visual llm for video understanding. *arXiv preprint arXiv:2312.06720*, 2023.
- Enxin Song, Wenhao Chai, Guan hong Wang, Yucheng Zhang, Haoyang Zhou, Feiyang Wu, Haozhe Chi, Xun Guo, Tian Ye, Yanting Zhang, et al. Moviechat: From dense token to sparse memory for long video understanding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 18221–18232, 2024.
- 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.
- Guangzhi Sun, Wenyi Yu, Changli Tang, Xianzhao Chen, Tian Tan, Wei Li, Lu Lu, Zejun Ma, Yuxuan Wang, and Chao Zhang. video-salmonn: Speech-enhanced audio-visual large language models. *arXiv preprint arXiv:2406.15704*, 2024.
- Guangzhi Sun, Yudong Yang, Jimin Zhuang, Changli Tang, Yixuan Li, Wei Li, Zejun Ma, and Chao Zhang. video-salmonn-ol: Reasoning-enhanced audio-visual large language model. *arXiv preprint arXiv:2502.11775*, 2025.

- Changli Tang, Yixuan Li, Yudong Yang, Jimin Zhuang, Guangzhi Sun, Wei Li, Zejun Ma, and Chao Zhang. video-salmonn 2: Captioning-enhanced audio-visual large language models. *arXiv preprint arXiv:2506.15220*, 2025a.
- Yunlong Tang, Daiki Shimada, Jing Bi, Mingqian Feng, Hang Hua, and Chenliang Xu. Empowering llms with pseudo-untrimmed videos for audio-visual temporal understanding. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pp. 7293–7301, 2025b.
- Gemini Team, Rohan Anil, Sebastian Borgeaud, Jean-Baptiste Alayrac, Jiahui Yu, Radu Soricut, Johan Schalkwyk, Andrew M Dai, Anja Hauth, Katie Millican, et al. Gemini: a family of highly capable multimodal models. *arXiv preprint arXiv:2312.11805*, 2023.
- Jiawei Wang, Liping Yuan, Yuchen Zhang, and Haomiao Sun. Tarsier: Recipes for training and evaluating large video description models. *arXiv preprint arXiv:2407.00634*, 2024a.
- Weihan Wang, Zehai He, Wenyi Hong, Yean Cheng, Xiaohan Zhang, Ji Qi, Xiaotao Gu, Shiyu Huang, Bin Xu, Yuxiao Dong, et al. Lvbench: An extreme long video understanding benchmark. *arXiv preprint arXiv:2406.08035*, 2024b.
- Shengqiong Wu, Hao Fei, Leigang Qu, Wei Ji, and Tat-Seng Chua. Next-gpt: Any-to-any multimodal llm. In *Forty-first International Conference on Machine Learning*, 2024a.
- Siwei Wu, Yizhi Li, Kang Zhu, Ge Zhang, Yiming Liang, Kaijing Ma, Chenghao Xiao, Haoran Zhang, Bohao Yang, Wenhui Chen, et al. Scimmir: Benchmarking scientific multi-modal information retrieval. *arXiv preprint arXiv:2401.13478*, 2024b.
- Jin Xu, Zhifang Guo, Jinzheng He, Hangrui Hu, Ting He, Shuai Bai, Keqin Chen, Jialin Wang, Yang Fan, Kai Dang, et al. Qwen2. 5-omni technical report. *arXiv preprint arXiv:2503.20215*, 2025a.
- Jin Xu, Zhifang Guo, Hangrui Hu, Yunfei Chu, Xiong Wang, Jinzheng He, Yuxuan Wang, Xian Shi, Ting He, Xinfu Zhu, et al. Qwen3-omni technical report. *arXiv preprint arXiv:2509.17765*, 2025b.
- An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Haoran Wei, et al. Qwen2. 5 technical report. *arXiv preprint arXiv:2412.15115*, 2024.
- Pinci Yang, Xin Wang, Xuguang Duan, Hong Chen, Runze Hou, Cong Jin, and Wenwu Zhu. Avqa: A dataset for audio-visual question answering on videos. In *Proceedings of the 30th ACM international conference on multimedia*, pp. 3480–3491, 2022.
- Yudong Yang, Jimin Zhuang, Guangzhi Sun, Changli Tang, Yixuan Li, Peihan Li, Yifan Jiang, Wei Li, Zejun Ma, and Chao Zhang. Acvubench: Audio-centric video understanding benchmark. *arXiv preprint arXiv:2503.19951*, 2025.
- Qilang Ye, Zitong Yu, Rui Shao, Xinyu Xie, Philip Torr, and Xiaochun Cao. Cat: Enhancing multimodal large language model to answer questions in dynamic audio-visual scenarios. In *European Conference on Computer Vision*, pp. 146–164. Springer, 2024.
- Liping Yuan, Jiawei Wang, Haomiao Sun, Yuchen Zhang, and Yuan Lin. Tarsier2: Advancing large vision-language models from detailed video description to comprehensive video understanding. *arXiv preprint arXiv:2501.07888*, 2025.
- Xiang Yue, Yuansheng Ni, Kai Zhang, Tianyu Zheng, Ruoqi Liu, Ge Zhang, Samuel Stevens, Dongfu Jiang, Weiming Ren, Yuxuan Sun, et al. Mmmu: A massive multi-discipline multimodal understanding and reasoning benchmark for expert agi. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 9556–9567, 2024.
- Boqiang Zhang, Kehan Li, Zesen Cheng, Zhiqiang Hu, Yuqian Yuan, Guanzheng Chen, Sicong Leng, Yuming Jiang, Hang Zhang, Xin Li, et al. Videollama 3: Frontier multimodal foundation models for image and video understanding. *arXiv preprint arXiv:2501.13106*, 2025.
- Yi-Fan Zhang, Huanyu Zhang, Haochen Tian, Chaoyou Fu, Shuangqing Zhang, Junfei Wu, Feng Li, Kun Wang, Qingsong Wen, Zhang Zhang, et al. Mme-realworld: Could your multimodal llm challenge high-resolution real-world scenarios that are difficult for humans? *arXiv preprint arXiv:2408.13257*, 2024a.

Yuanhan Zhang, Jinming Wu, Wei Li, Bo Li, Zejun Ma, Ziwei Liu, and Chunyuan Li. Video instruction tuning with synthetic data. *arXiv preprint arXiv:2410.02713*, 2024b.

## A MORE DETAILS ON JOINTAVBENCH

### A.1 TASK DEFINITION

We construct a taxonomy of 15 tasks requiring audio-visual joint reasoning ability based on the benchmark’s requirements. The detailed task descriptions with examples are presented in Table 5. Since single-scene tasks and multi-scene tasks require focus on movie details (e.g., visual cues, temporal relationships), question templates are designed to ensure the generated questions both capture critical information and comply with task specifications.

Table 5: Problem Categories and Definitions of JointAVBench’s Taxonomy

Task Name	Code	Category Description	Question Example
Speech-based Timepoint Localization	STL	Locate the temporal position of an object in the dialogue	Which objects are mentioned only in the dialogue but not clearly shown in the video, and when does the first object appear in the dialogue?
Speaker Spatial Localization	SPL	Locate the spatial position of a character in the video	Where’s the character that says ‘I’m gonna give you a compliment now’ with a contemptuous tone located in the video?
Sounding Object Grounding	SOOG	Locate the spatial position of a sound-emitting object in the movie	What is the spatial position of the object that produced the loud bang in the scene?
Sound Event Recognition	SOER	Infer what action occurred that caused the sound	What makes the high-pitched, bright sound?
Speaker Emotion Recognition	SPER	Identify the emotions of characters	What’s the emotion of the speaker that wears a brown jacket over a blue shirt?
Musical Tone Inference	MPTI	Determine the overall tone of a scene	What is the overall atmosphere of the scene?
Vision-Speech Sequence Recognition	VSSR	Determine the chronological order of element appearances	In what order were the following items mentioned in the video? (a) ‘Did you see this guy smoking pot?’ (b) The police officer holding an object. (c) ‘Excuse me, sir.’
Cross-scene Association	CSA	Identify associations between elements across different scenes	Which dialogue in the remaining parts is most relevant to what the man does in 18.56s-35.16s?
Multi-plot Ordering	MPO	Order different audio-visual details across different movie segments	In what order were the following items mentioned in the video? (a) The girl is seen adjusting an oxygen mask on a child (b) The woman is seen trimming flower stems. (c) The man says, ‘‘This isn’t funny anymore.’’
Plot Temporal Grounding	PTG	Identify the approximate temporal position of plot segments	When did the woman in the green tulle outfit reveal her success in the music industry?
Audio Function Analysis	AFA	Analyze the purpose of sound effects and background music in movie segments	How does the video depict the movement of the man in the beige suit?
Plot Development Prediction	PDP	Predict future plot developments based on existing plot elements	Which of the following options is most likely to occur after this video ends?
Audio-Visual Detail Memory	AVDM	Test the model’s long-term memory capability	What was the man doing when the police officer asked if they were smoking?
Musical Emotion Shift Inference	MESI	Identify emotional changes and trends throughout the movie	How does the emotional tone evolve from the beginning to the middle of the movie?
Character Relationship Inference	CRI	Infer complex relationships between characters based on the overall plot	What is the relationship between the man in the brown jacket and the police officer?

### A.2 MORE STATISTICS

To further analyze the descriptive characteristics of audiovisual events in our dataset, we performed lexical frequency analysis on both captions and QA pairs, visualized through the word cloud in Figure 8c. The results highlight frequent visual descriptors (e.g., ‘‘position’’, ‘‘object’’, ‘‘location’’) and auditory terms (e.g., ‘‘sound’’, ‘‘speaker’’, ‘‘speech’’). Furthermore, we examined the distribution of question lengths (in words), as illustrated in Figure 8d, which confirms the conciseness of our formulated questions with minimal redundancy. Figure 8a presents the quantitative distribution of QA pairs across different tasks, revealing that single-scene tasks exhibit particularly rich audiovisual

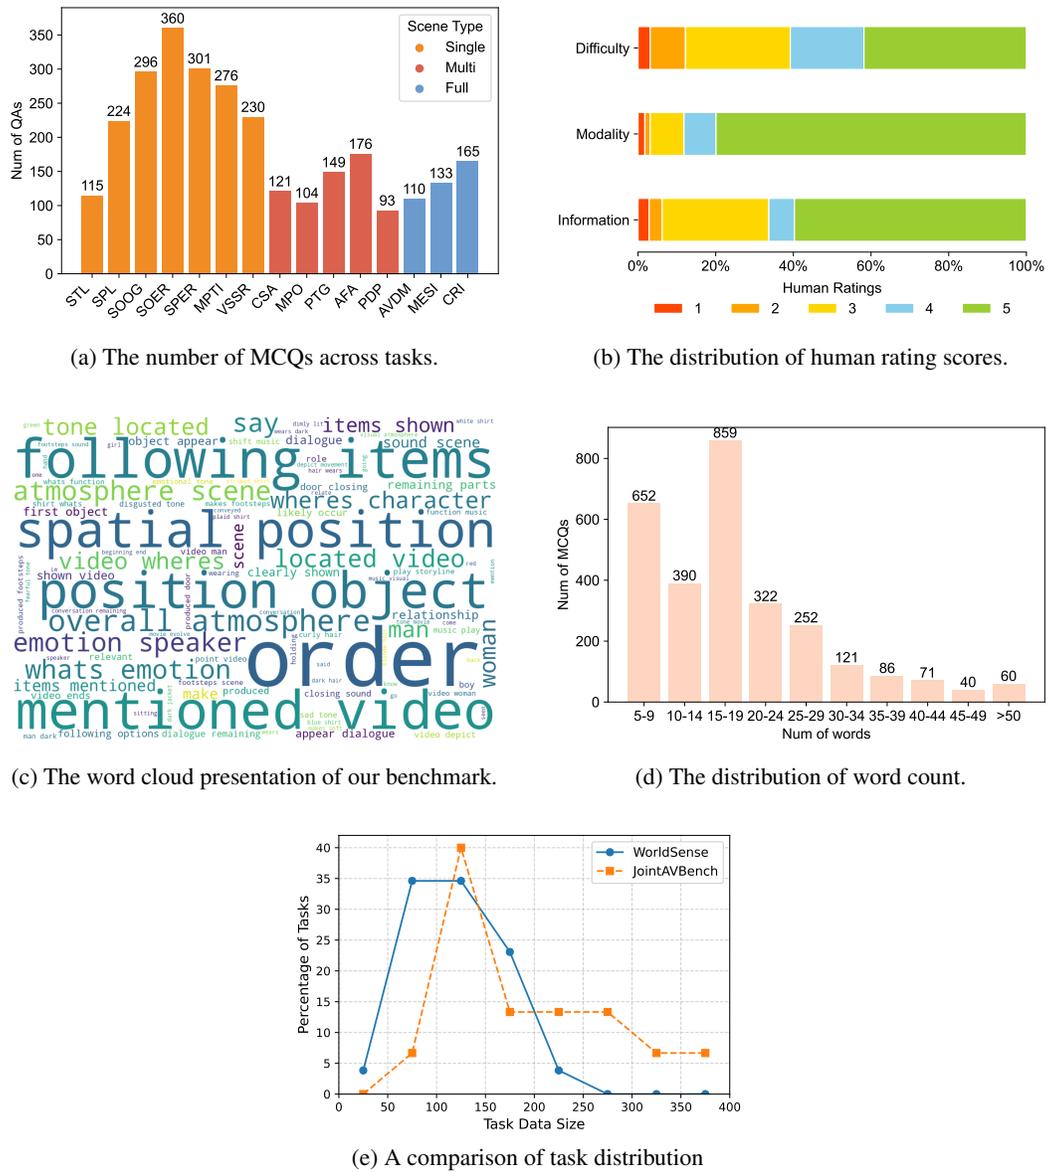


Figure 8: More details of JointAVBench. The number 1-5 in Figure (b) are human ratings, 1 represents the lowest score and 5 represents the highest score.

correlations and simplicity in question designs. We also present the comparison of our task sample size with that of WorldSense (Hong et al., 2025) in Figure 8e. This comparison indicates that our benchmark contains more data samples per task compared with WorldSense, which also proves our benchmark’s credibility.

### A.3 MANUAL CHECK

To ensure the quality of JointAVBench, we engage a group of annotators to evaluate and correct the questions according to the criteria specified in Section 3.3. The annotation process consists of two main steps: First, annotators rate the generated MCQs; then, for incorrect MCQs where the correct answer appears in the distractors, they replace the designated correct answer with the appropriate distractor. Through this process, we maintain high-quality MCQs while discarding only those that fail

to meet our quality standards. After human verification, we discard 17.4% of the data for incorrect answers. Other data are filtered based on the other 3 judging criteria, resulting in 2,853 MCQs.

Figure 8b presents the human evaluation scores of the automatically generated MCQs. The results demonstrate that our pipeline generates high-quality questions, with the majority of MCQs receiving the highest rating across all evaluation criteria. Specifically, we find that our QAs are relatively difficult, with the difficulty criteria having high ratings.

#### A.4 ANNOTATION INFORMATION

To ensure that our manual data check process is rigorous, we have selected professional annotators. Our annotation team is provided by a professional crowdsourcing labeling company. Each of our annotators has a bachelor’s degree from a highly ranked university, is proficient in English, and has passed a qualification test on sample tasks before participating.

To assess the reliability of our annotations, we randomly selected 400 samples and had them independently annotated by two annotators. We calculated Cohen’s Kappa coefficient to measure the agreement. The resulting Cohen’s Kappa was 0.713, which is interpreted as “Substantial Agreement” according to established benchmarks (Landis & Koch, 1977). This strong level of agreement validates the clarity of our annotation guidelines and the reliability of our final benchmark.

#### A.5 EXPERIMENT DETAILS

To ensure comparability between MLLMs, we selected similar model parameters and video frames. For open-source models, we selected 7B as the parameter size (we selected InternVL2.5-8B (Chen et al., 2024c) since it does not provide a 7B model). We also selected 32 frames as the maximum number of frames per video to ensure that the frame number did not affect performance. For closed-source models, we adhered to the official model settings. Notably, for GPT-4o (Hurst et al., 2024), we only input video frames alongside the question text, and therefore, the model is unable to access timestamps. During the experiment, we randomly shuffled the options to ensure that the distribution of correct answer prefixes was uniform and free from biases. For Gemini2.5-Flash, we use the setting: temperature=1.0, max\_temperature=2.0, top\_p=0.95, top\_k=64, thinking=True. For Gemini2.5-Pro, we use the setting: temperature=1.0, max\_temperature=2.0, top\_p=0.95, top\_k=64, thinking=True.

All experiments were conducted on NVIDIA H-100 GPUs and can be reproduced using a single H-100 (80G) GPU. During the automatic generation process, we used the official API (model name: ‘qwen2.5-72b-instruct’) from Qwen to ensure long-context capability and generation stability.

#### A.6 CASES

Figure 9 presents detailed examples from our benchmark. The questions and options in JointAVBench are designed to incorporate both audio and video modalities and to evaluate audio-visual joint reasoning capabilities. We present the task names, a few frames of the video, questions, options, and correct answers in the cases.

## B DETAILS ON GENERATION PIPELINE

### B.1 VIDEO CAPTION GENERATION

We generate visual captions using Qwen2.5-VL (Bai et al., 2025) at a rate of 1 frame per second (fps) for each identified scene, regardless of its quality level (low or high). To ensure the captions capture both static and dynamic scene information, we carefully design a video captioning prompt as shown in Figure 10.

### B.2 AUDIO CAPTION GENERATION

To ensure the audio captions contain rich information about diverse audio signal types, we generate separate captions for each type. We then utilize an LLM with carefully designed Chain-of-Thought (CoT) prompts to reduce caption hallucination.

**Vocal Traits, Sound Event, and Music Description.** We find that directly using general-purpose audio-language models (ALMs) to generate overall audio captions often overlooks important details and may produce inaccurate results. Therefore, we employ Qwen2.5-Omni (Xu et al., 2025a), an open-source multimodal model, to separately generate descriptions of vocal traits, sound events, and music components. Notably, current ALMs cannot reliably distinguish between sound events and music. We address this limitation by generating both sound event and music descriptions initially, then separating them during post-processing. The detailed prompt templates are provided in Figure 10.

**Subtitle Transcription.** For dialogue transcription, our primary objectives are accurate speech recognition and precise timestamp generation. Since general ALMs underperform in timestamp estimation, we utilize Whisper-v3 (Radford et al., 2023), an advanced and widely utilized automatic speech recognition (ASR) system, to ensure transcription quality.

**Audio Caption Refinement.** The initial audio descriptions contain hallucinations (*i.e.*, factually incorrect or repetitive outputs). We employ Qwen-2.5 (Yang et al., 2024) to perform three refinement steps: (1) distinguishing between background music and sound events, (2) aligning vocal characteristics with dialogue transcripts, and (3) removing redundant content. The detailed prompt engineering for this process is illustrated in Figure 11.

### B.3 QA PAIR GENERATION

We utilize Qwen2.5 (Yang et al., 2024) to generate QA pairs following predefined question templates. For single-scene and multi-scene tasks, we only provide descriptions for high-quality scenes. For full-scene tasks, we include all scene descriptions regardless of quality to ensure no details are omitted. After generating multi-scene QA pairs, we verify the interval between questions to ensure they require information from multiple scenes, using the prompt shown in Figure 16 to identify their information intervals. Additionally, we require the model to generate a brief justification for each answer while generating QA pairs.

### B.4 QUALITY CONTROL

In the quality control stage, we employ extensive CoT techniques to ensure that Qwen2.5 (Yang et al., 2024) achieves optimal performance in identifying potential hallucinations.

During the general check, we utilize only the QA pair and its explanation to filter out unqualified QA pairs. This stage includes four checks: modality, format, content, and speculation checks. The details of each check are as follows: (i) modality check assesses whether the modality clues used in the QA pair are derived from dual modalities; (ii) format check evaluates whether the answer corresponds to the question in format (*e.g.*, the answer explains two items, but the question asks about only one item); (iii) content check verifies whether the answer can be logically inferred from the question based on the explanation; (iiii) speculation check examines whether the answer relies excessively on speculation rather than concrete evidence. The prompts used in this stage are shown in Figure 13.

For the specific check, we design task-specific prompts based on the definition of each task. The prompts for each specific check are presented in Figure 14 and Figure 15. Note that for better adaptation to different tasks, the prompt for the sequence check varies slightly between VSSR and MPO, and the ambiguity check varies slightly between SPL and SOER.

### B.5 DISTRACTOR GENERATION

Since generating distractors requires additional information from the video, we generated them after filtering the QA pairs. In this process, we designed a generation prompt incorporating various error types to ensure option diversity and complexity, as illustrated in Figure 16. This includes common error categories such as incorrect details and temporal/spatial misplacement.

Table 6: AV Correlation Scores for Different Datasets

Dataset	Automatic Score	Human Score
Music-AVQA	56.7%	54.5%
OmniBench	100%	94.0%
AV-Odyssey	100%	99.5%
LongVALE	76.2%	75.0%
WorldSense	62.9%	60.5%
JointAVBench (ours)	93.5%	94.5%

Table 7: Examples of failure cases.

Task	Question	Groundtruth Answer	Model Answer
SPL	Where’s the character that says ‘Oh, my God’ with a vibrant shocked tone located in the video? A. In the left B. In the center C. Standing behind the loveseat D. On the right	A	D. On the right.
SPER	What’s the emotion of the speaker that wears a vibrant green tulle outfit with hair rollers and dramatic makeup? A. Surprised B. Amused C. Confused D. Excited	D	A. Surprised.
MPO	In what order were the following items mentioned in the video? (a) ‘When do you need to have the van back’. (b) The man talked about cars with a joyful tone (c) The man is driving a car along a road lined with bare trees A. (c) (b) (a) B. (a) (b) (c) C. (b) (c) (a) D. (c) (a) (b)	D	The correct answer is A.
PTG	When did the woman in the green tulle outfit reveal her success in the music industry? A. 51.05s-80.79s B. 126.38s-174.51s C. 86.46s-126.38s D. 45.84s-51.05s	A	The correct answer is B.

## C MORE EXPERIMENTS

### C.1 EVALUATION OF AV CORRELATION FOR PREVIOUS WORKS

We evaluate the AV correlation ratio (the proportion of questions that truly require auditory and visual information to answer) for previous works in Table 1. However, since some benchmarks do not evaluate the ratio themselves, we utilize Qwen-2.5 (Yang et al., 2024) to judge the correlation. We utilize our modality check prompt (shown in Figure 13) and sample 1,000 QA pairs from each dataset to judge. We also recruit human volunteers to verify the AV correlation ratio by asking them to assess whether the randomly sampled 100 data items from each benchmark require both auditory and visual information to answer. The overall comparison is described in Table 6, indicating that our automatically calculated score has high correspondence to human scores. This score indicates that our benchmark can assess joint audio-visual reasoning ability.

### C.2 ERROR STUDY

We’d like to provide some failure cases and deep analysis in this section. The failure cases are selected from the worst-performing tasks.

Based on the failure cases, we find that:

- **Models fail to understand vocal traits.** In the first example, the model fails to find the spatial information based on speech and vocal traits. In the second example, the model can’t find vocal traits based on visual information. These two examples indicate that the ability to align visual information with vocal traits remains low in current models and needs to be improved.
- **Models fail to understand temporal information.** In the third example, the question tests the model’s ability to understand the storyline and arrange the detailed information. And the fourth example tests the model’s temporal grounding ability from a long video. These two examples showcase that future works should focus on increasing the model’s ability to understand temporal relationships in audio-visual scenarios.

## D LIMITATIONS

We acknowledge that the JointAVBench has several limitations in its generation and experimental evaluation. First, the dataset is exclusively derived from SF20K, which may introduce biases in data distribution. Second, our designed taxonomy, while comprehensive, may not encompass all dimensions of audio-visual joint reasoning capabilities. Nevertheless, we have rigorously ensured that the included tasks cover critical aspects and effectively assess the target abilities. Third, due to computational constraints, our experiments were limited to selected representative MLLMs rather than an exhaustive evaluation. We intend to address these limitations in future work through dataset expansion and more extensive benchmarking.

## E BROADER IMPACTS

We constructed JointAVBench to facilitate research and development in omni-LLMs and video understanding. We anticipate that this dataset may yield both positive and negative societal impacts. JointAVBench offers several potential benefits, including: (1) enabling development of human-like agent systems, (2) advancing video understanding tools, and (3) creating assistive software for people with disabilities. However, the dataset also presents certain risks, such as privacy concerns and copyright issues. We believe a thorough discussion of these benefits and challenges will lead to a more comprehensive understanding of the dataset’s societal implications.

## F DECLARATION OF LLM USAGE

During our research, we use LLMs as a major dataset construction tool, including dataset generation, quality control, and experiments. During our paper writing, we use LLMs to polish our paper and correct the defects in our paper.

## G SAFEGUARDS

To ensure our benchmark excludes unsafe content, we adopted two key measures. First, we used the publicly released SF20K dataset (Ghermi et al., 2024) as the foundation, which provides pre-filtered safe content. Second, we employed the official Qwen API during benchmark generation, whose built-in safety mechanisms automatically screen both input prompts and output responses for potentially unsafe video recommendations.

## H LICENSE

The JointAVBench dataset is released under the CC BY-NC-SA 4.0 license. Subsequent research using this dataset must comply with the license terms.



Figure 9: Additional cases of JointAVBench. The first and second row represents single-scene tasks, the third row represents multi-scene tasks, and the last fourth row represents full-scene tasks.

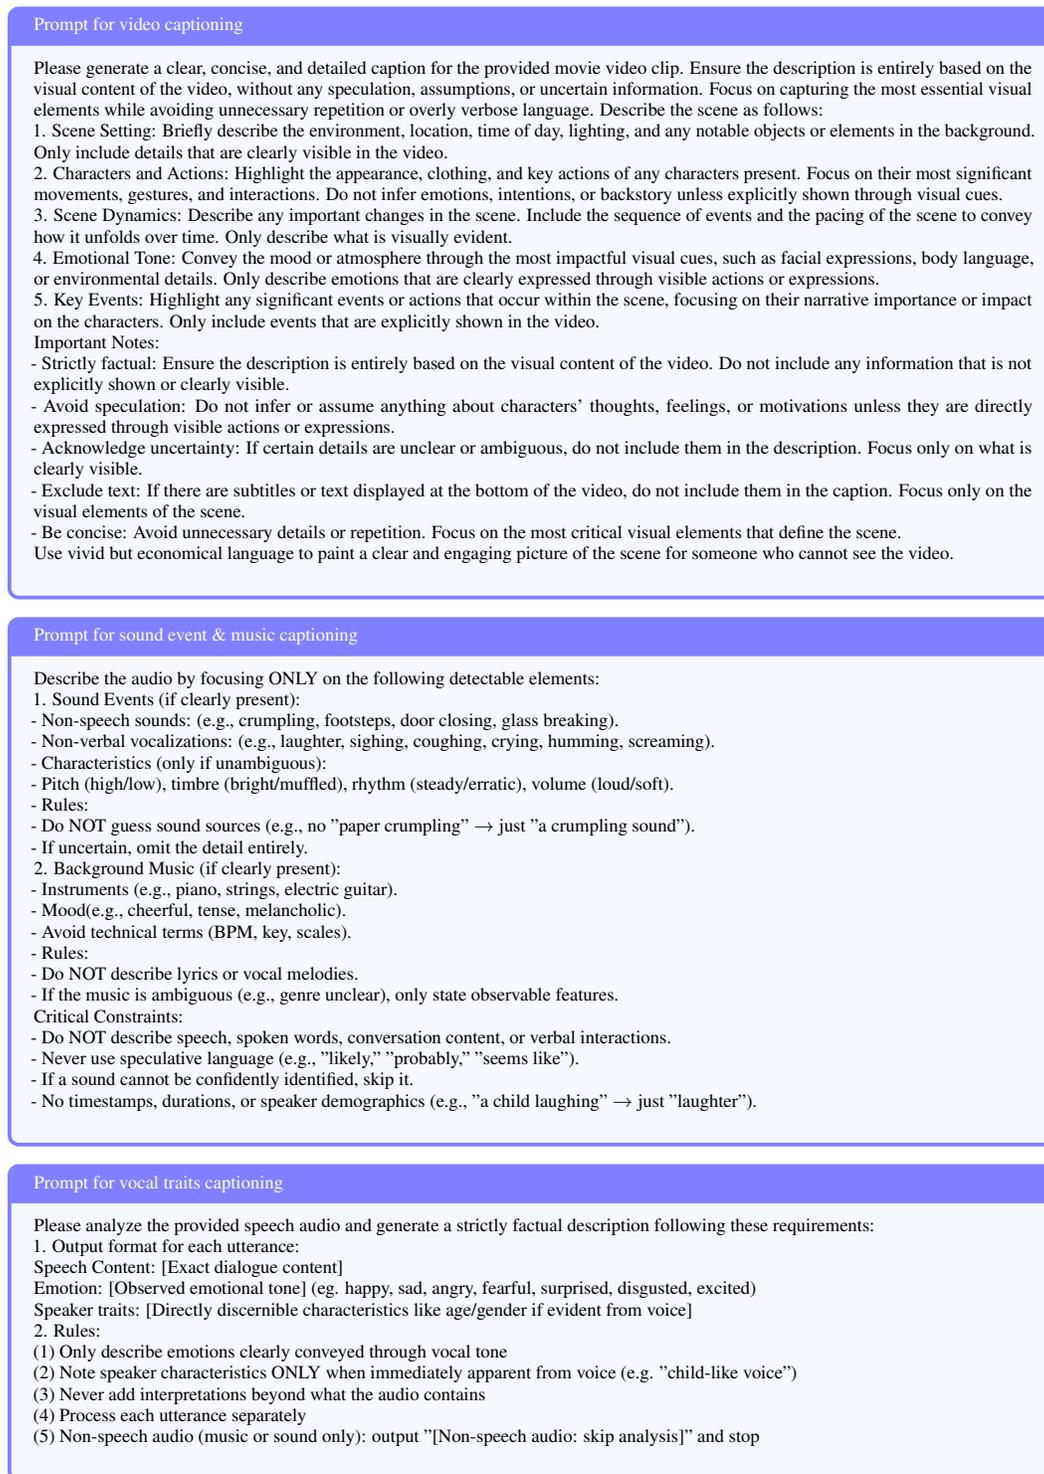


Figure 10: Prompts for generating omni-modal caption.

## Prompt for cleaning and separating sound event caption and music caption

You are an expert in audio data cleaning. Your task is to clean and organize audio captions by separating them into two distinct categories: [music] and [sound event].

Please analyze the provided audio caption and separate it into two parts:

[music]: Include descriptions related to music, such as background music, musical instruments, and any emotional or atmospheric content (e.g., tense, sad). If no music is described, output "None".

[sound event]: Include descriptions of Non-musical sounds (e.g., actions, engine noises, wind) and Non-verbal vocal sounds (e.g., laughter, sighing, coughing) (EXCEPT for any mention of "Crumpling sound" - these should be completely omitted). If no sound events are described (after filtering out Crumpling sounds), output "None".

Do not include any content related to speech, dialogue, or verbal communication in either section. Remember not to generate any explanations or notations after the separated part output.

Output Format:

[music] ;description of music and emotional characteristics if any; or None

[sound event] ;description of sound events; or None

Example: ;example;

## Prompt for cleaning vocal traits caption

You are an expert in speech and subtitle data analysis. You have access to the following:

Vocal traits: The text that describes the emotion and other characteristics of speech sentence by sentence, each utterance contains speech content, emotion, and speaker traits.

Subtitle: The transcription of speech. Each line in the subtitle represents a single phrase.

Your task is to evaluate the provided vocal traits by comparing them with the corresponding subtitle. The goal is to determine whether the vocal traits contain usable information based on their alignment with the subtitle. Your analysis should be rigorous and follow a structured approach.

Please follow the guidelines below:

1. Comparison with Subtitle:

- Directly compare the Speech Content text with the provided subtitle.
- Discard any speech content marked with neutral emotion (e.g., "neutral tone", "neutral mood") immediately, regardless of subtitle alignment.
- For non-neutral speech content:
- If the speech content contains phrases that overlap or align with the subtitle (e.g., shared keywords or contextual similarity), proceed to the next stage.
- If no alignment is found, the vocal traits are unusable.

2. Output of Emotional Information:

For each utterance:

- Only if the vocal traits align with the subtitle and are non-neutral:
- Replace any speech content in the vocal traits (i.e., quoted or referenced dialogue) with the exact matching phrase from the subtitle.
- Preserve all other emotional/tonal descriptors (e.g., "sad mood," "English accent").
- Format the output as a coherent description combining the subtitle content and emotional features.
- If no alignment exists, output '[Unavailable]':
- Always prefix the final output with '[Output]' or '[Unavailable]':

Provide the final cleaned utterances in the following format:

[Output] ;utterances with emotional/tonal information; if subtitle roughly matches or [Unavailable] if NO utterance is available

Example: ;example;

Input:

Vocal Traits: {vocal\_traits}

Subtitle: {subtitle}

Please follow the guidelines to clean the vocal traits text, and provide the final output in the specified format. Remember to output '[Unavailable]' only when no utterance is available.

Figure 11: Prompt for audio caption refinement.

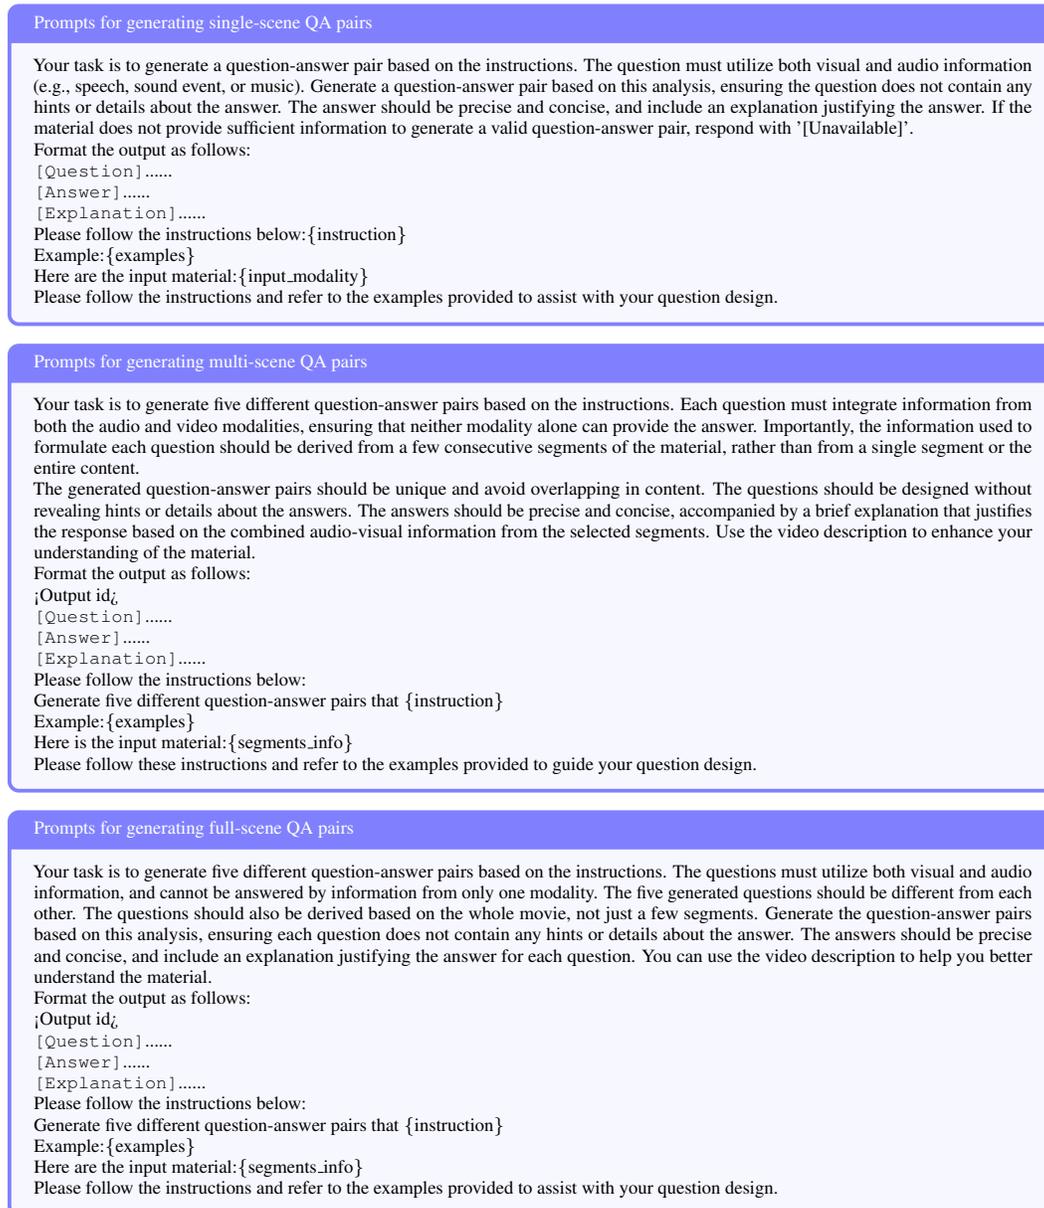


Figure 12: Prompts for generating QA pairs

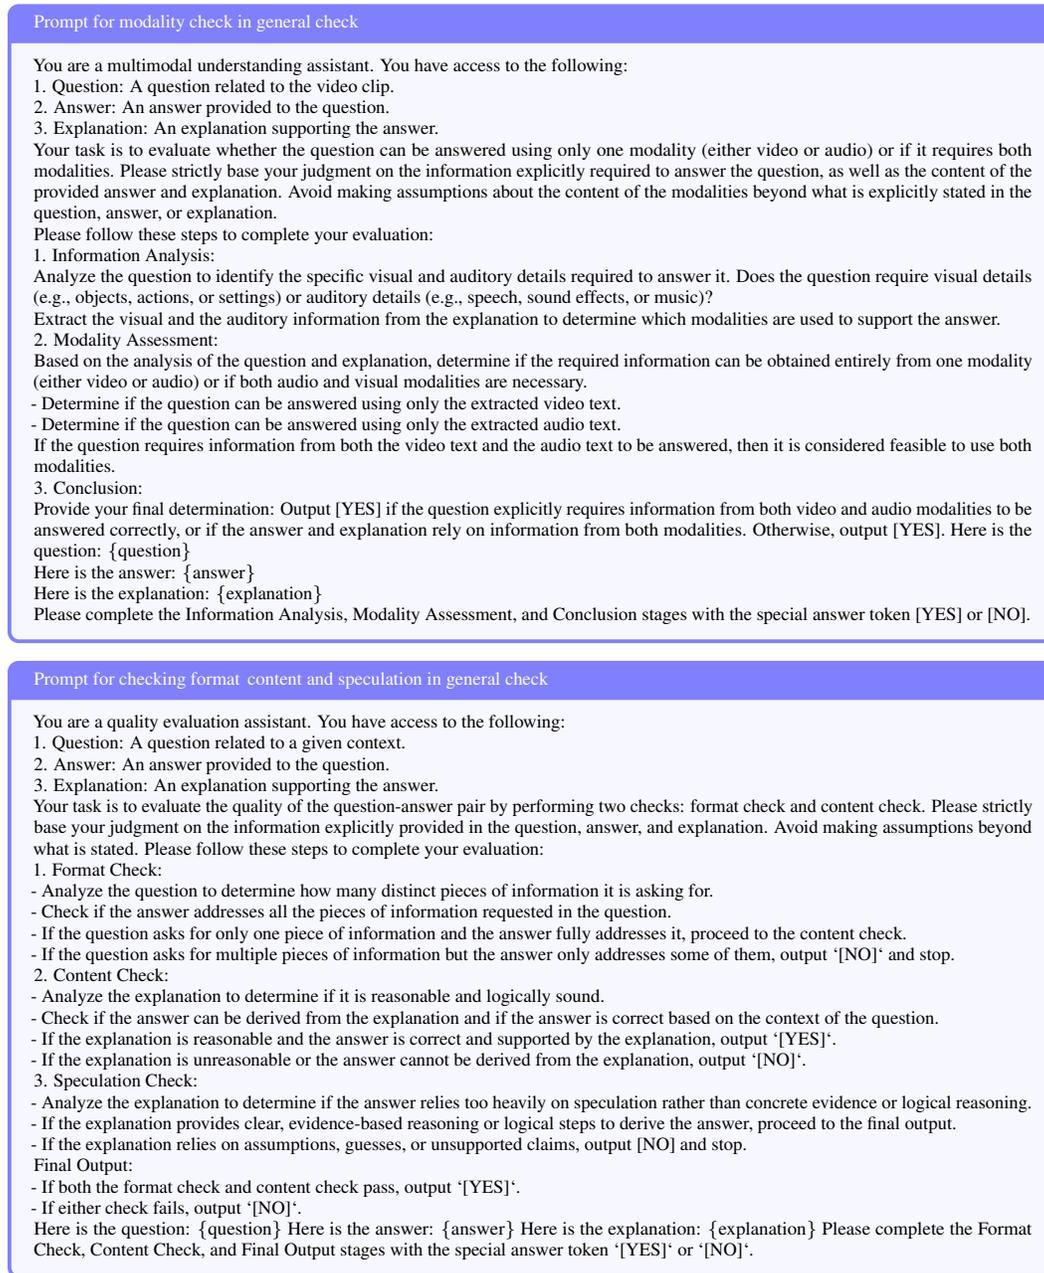


Figure 13: Prompts for general checks



Figure 14: Prompts for sequence check and ambiguity check.

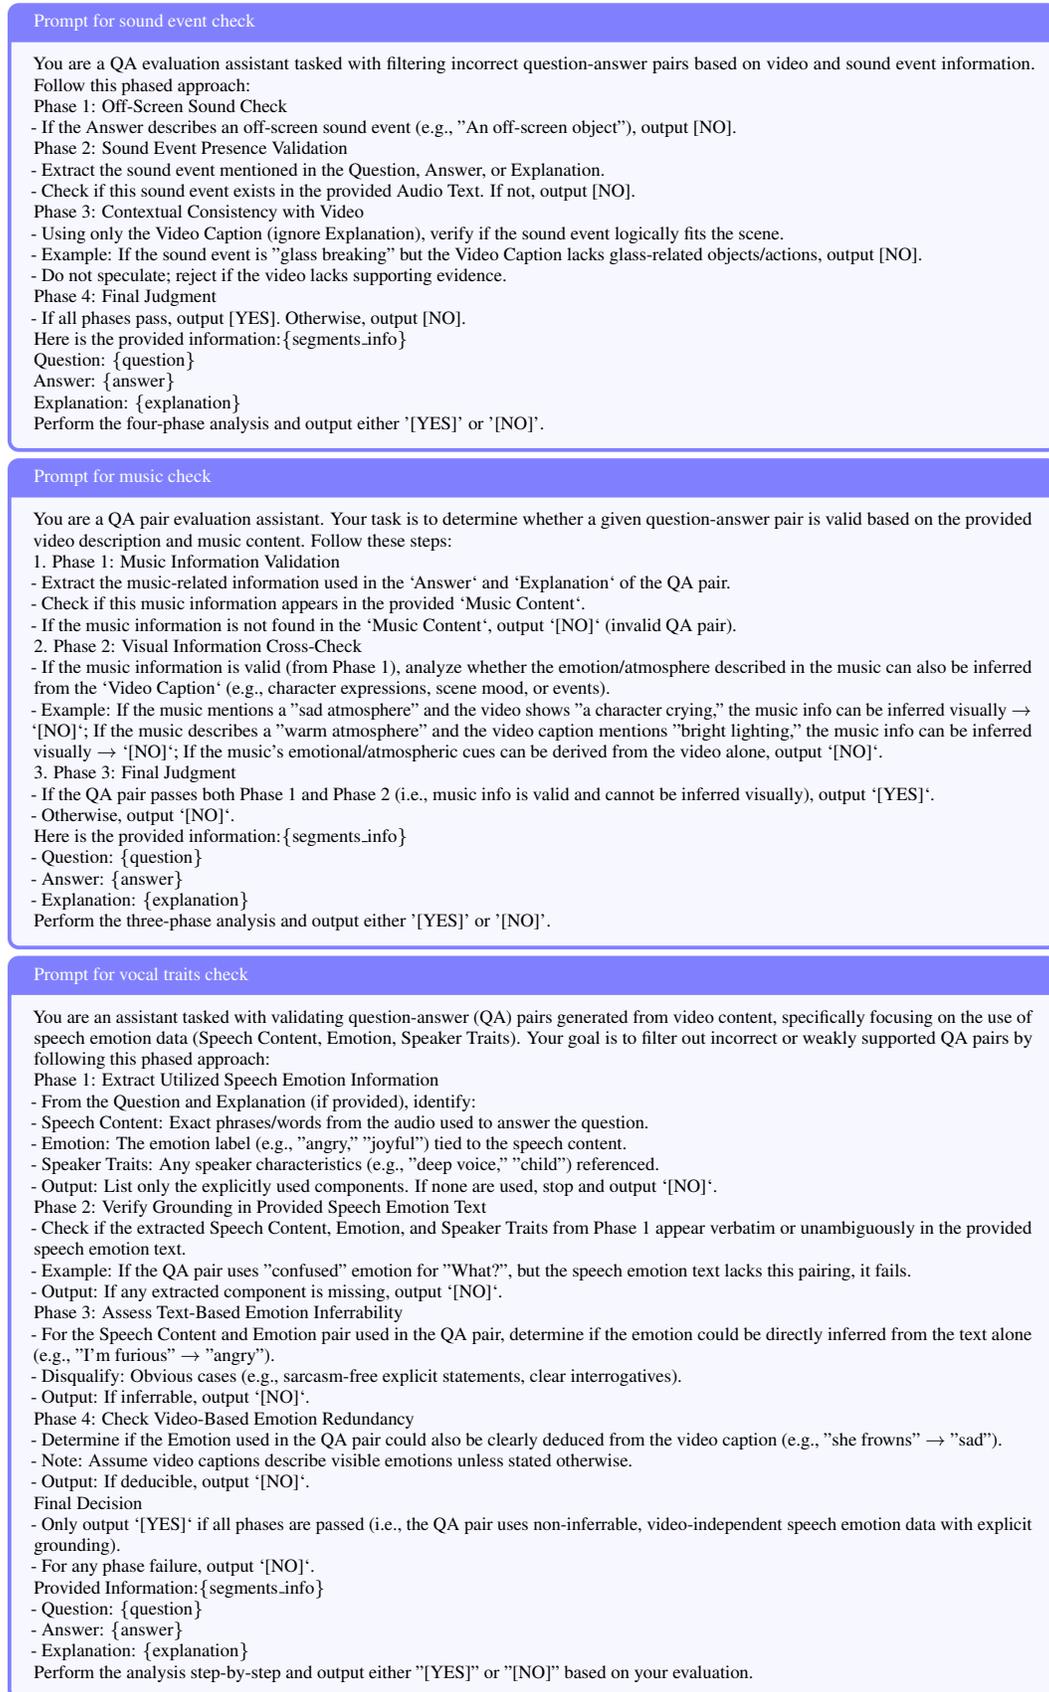


Figure 15: Prompts for audio check.

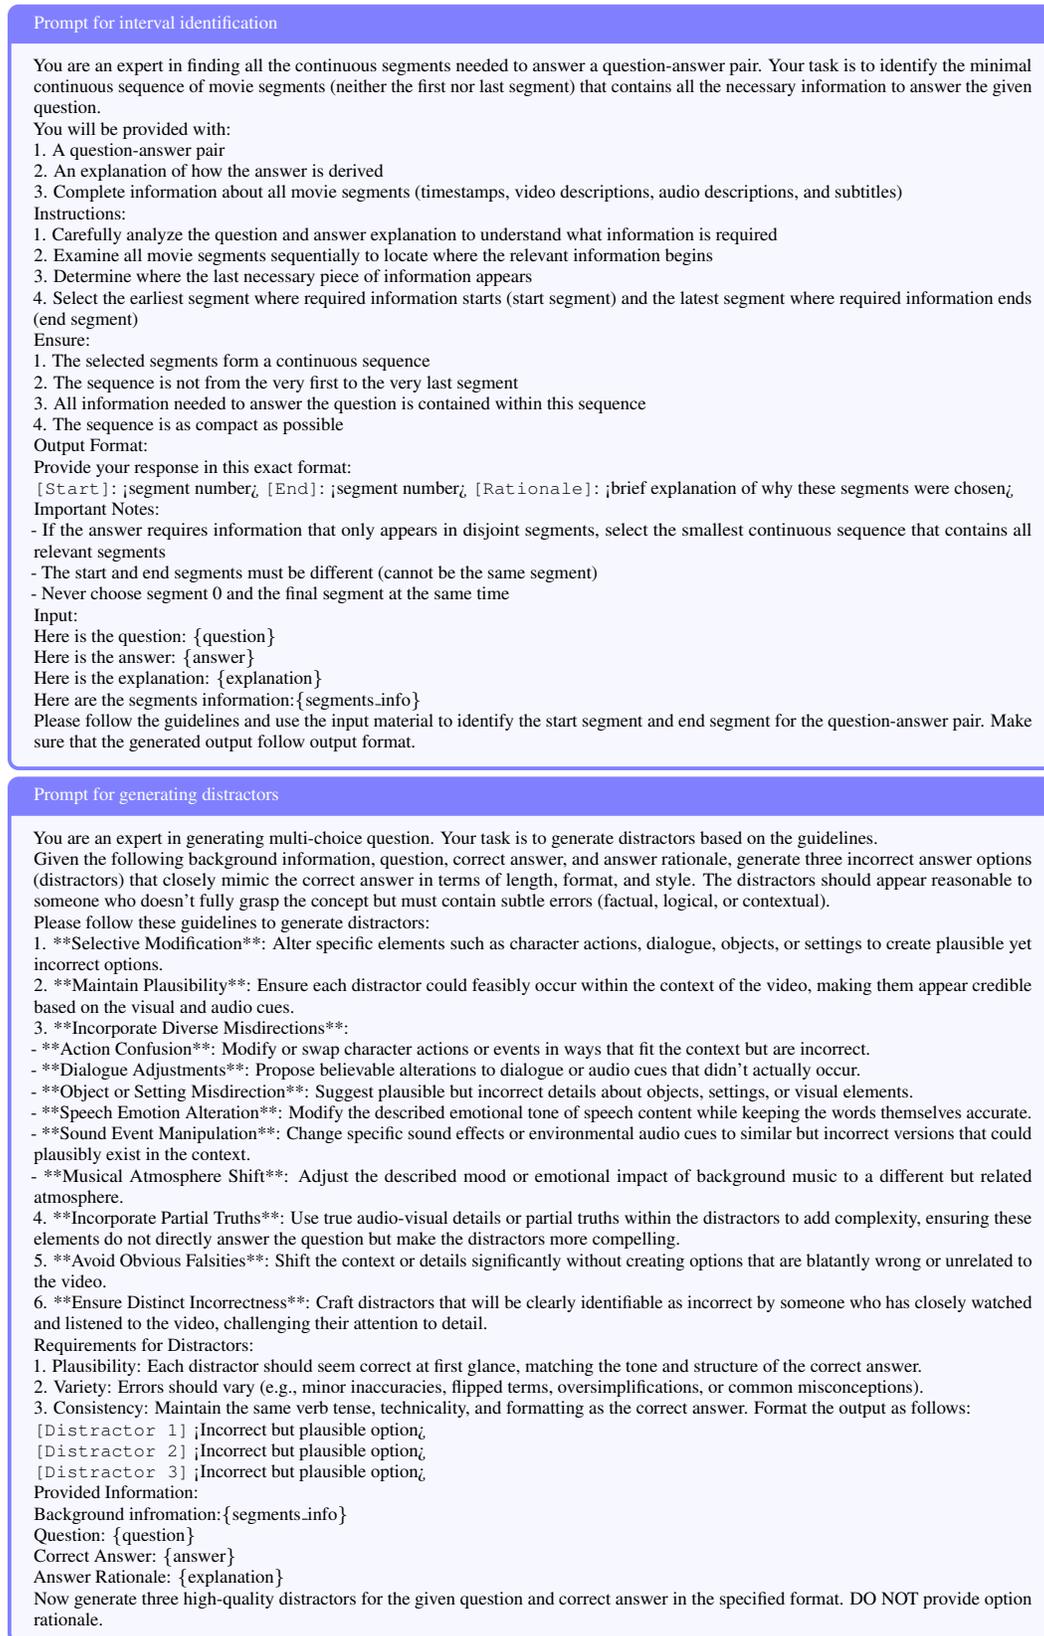


Figure 16: Prompts for interval check and distractor generation.