HarmonySet: A Comprehensive Dataset for Understanding Video-Music Semantic Alignment and Temporal Synchronization

Zitang Zhou[*,1,2], **Ke Mei**[*,1], **Yu Lu**[1,3,✉], **Tianyi Wang**[1], **Fengyun Rao**[1]
[1] Wechat, Tencent Inc. [2] Beijing University of Posts and Telecommunications
[3] Zhejiang University

Figure 1. We introduce HarmonySet, the first instruction tuning dataset for MLLMs to understand the alignment between video and music. While existing MLLMs typically offer surface-level interpretations of video-music relationships, HarmonySet includes 48,328 video-music pairs, each annotated with rich information on rhythmic synchronization, emotional alignment, thematic coherence, and cultural relevance.

Abstract

This paper introduces HarmonySet, a comprehensive dataset designed to advance video-music understanding. HarmonySet consists of 48,328 diverse video-music pairs, annotated with detailed information on rhythmic synchronization, emotional alignment, thematic coherence, and cultural relevance. We propose a multi-step human-machine collaborative framework for efficient annotation, combining human insights with machine-generated descriptions to identify key transitions and assess alignment across multiple dimensions. Addition-

[*] Equal conrtibution [✉] Corresponding author: Yu Lu.

ally, we introduce a novel evaluation framework with tasks and metrics to assess the multi-dimensional alignment of video and music, including rhythm, emotion, theme, and cultural context. Our extensive experiments demonstrate that HarmonySet, along with the proposed evaluation framework, significantly improves the ability of multimodal models to capture and analyze the intricate relationships between video and music. Project page: https://harmonyset.github.io/.

1. Introduction

The rapid growth of online video platforms has led to a rising demand for multimodal content analysis across video, text, and music. This demand is fueled by advancements in large-scale multimodal datasets [7–14] and models, such as Video Multimodal Large Language Models (MLLMs) [15–19], which show strong potential for understanding video semantics and performing cross-modal reasoning tasks.

Despite these advancements, a key challenge remains in the domain of video-music understanding, where capturing the complex semantic and temporal relationships between video content and music proves difficult [20]. Effective video-music understanding requires the ability to recognize nuanced elements, such as emotional tone, narrative progression, and symbolic imagery—critical aspects that underlie the synchronization between video and music. Current models [21–24], however, often provide surface-level interpretations of video-music relationships, failing to capture deeper, context-specific insights, such as rhythm synchronization, emotional alignment, and thematic coherence (as illustrated in the left panel of Figure 1).

A significant limitation in addressing these challenges is the lack of effective datasets that provide comprehensive annotations for video-music understanding. Existing datasets offer paired video and music content [1, 2, 5, 25–28], but their textual annotations typically consist of basic descriptions [3, 6, 29] that fail to capture the detailed semantic alignment and temporal synchronization necessary for effective training of MLLMs. This results in a limited understanding of how music influences the narrative rhythm and emotional tone of video content.

Creating datasets that capture these complex video-music relationships with detailed annotations is a labor-intensive process. Annotators must watch videos while listening to the accompanying music, carefully identifying key transitions to ensure precise temporal alignment. Furthermore, evaluating video-music pairs is inherently subjective [30], as personal taste and cultural context can significantly influence interpretations, making it difficult to standardize annotations.

To address these challenges, we introduce HarmonySet, a novel dataset designed to facilitate a deeper understanding of video-music alignment. HarmonySet consists of 48,328 diverse video-music pairs, curated from a broad range of genres to ensure comprehensive representation. Each pair is annotated with rich information on key aspects of temporal synchronization and semantic alignment, enabling more robust training of multimodal models. As illustrated in Figure 1 , HarmonySet provides annotations that go beyond simple descriptions, offering detailed insights into how video and music align both temporally and semantically.

To efficiently generate these annotations, we propose a multi-step human-machine collaborative labeling framework. Initially, human annotators identify key timestamps that mark synchronized transitions between video and music, forming the foundation for deeper analysis. These timestamps serve as anchors for categorizing the video-music alignment into dimensions such as rhythm synchronization, emotional alignment, thematic coherence, and cultural relevance. Annotators then assess each dimension on a scale, ensuring that the annotations capture the full complexity of the video-music relationship. Machine-generated descriptions are subsequently produced by an MLLM [19], which utilizes the identified timestamps and video metadata to provide detailed, context-aware descriptions of the video-music alignment. This combined human-machine approach significantly reduces annotation workload while maintaining high-quality, multi-dimensional insights.

In addition to the dataset, we introduce a novel evaluation framework for benchmarking video-music understanding models. Our framework includes a series of tasks and metrics designed to evaluate critical aspects of video-music alignment, such as temporal synchronization, emotional congruence, and thematic integration. By providing standardized benchmarks, we aim to establish a more rigorous approach to evaluating the performance of models in understanding the complex interplay between video and music.

Comprehensive experiments demonstrate that both HarmonySet and our evaluation framework significantly enhance the ability of multimodal models to capture and analyze the intricate relationships between video and music.

Our key contributions are threefold:

- We introduce HarmonySet, a diverse collection of video-music pairs with rich annotations on rhythmic synchronization, emotional alignment, and thematic coherence, addressing the gap in existing datasets for video-music understanding.
- We propose an efficient, multi-step human-machine framework for annotating video-music relationships. This approach combines human insights with machine-generated descriptions to label key transitions and assess alignment across multiple dimensions.
- We introduce a new evaluation framework with tasks and metrics for assessing temporal alignment, emotional congruence, and thematic integration, providing a standardized benchmark for video-music understanding tasks.

Table 1. **Overview of Video-Music Datasets.** HarmonySet provides comprehensive video-music content, and stands out among existing video-music datasets by offering both semantic matching and temporal synchronization annotations.

Dataset	Year	Music Style	#Hours	#Videos	#Annotations	Semantic Matching	Temporal Synchronization
TT-150K [1]	2021	diverse	-	146,351	-	✗	✗
MovieClips [2]	2022	diverse	230	20,000	-	✗	✗
MuseChat [3]	2024	songs	-	98,206	98,206	✓	✗
BGM909 [4]	2024	piano	-	909	9,090	✗	✗
SVM-10K [5]	2024	diverse	-	10,000	-	✗	✗
MMTrail [6]	2024	diverse	27,100	290,000	290,000	✗	✗
HarmonySet (Ours)	2024	diverse	458.8	48,328	48,328	✓	✓

2. Related Work

2.1. Video-Audio Datasets

Existing datasets used for training MLLMs emphasize general audio features, rather than the specific musical elements that are central to modern video multimodal contents. For instance, AudioSet [31] and VGGSound [32] are large-scale datasets primarily designed for audio event recognition. Other datasets like FSD50K [33] and ESC-50 [34] are also commonly employed in pre-training multimodal models that accept audio inputs.

2.2. Video-Music Datasets

As shown in Table 1, recent benchmarks [25, 27, 28] incorporate video-music content, exploring video-level visual-music semantic alignment. For instance, TT-150K [1] collected 150,000 short videos with music tracks for video-music recommendation. SVM-10K [5] collected short videos with high likes for filtering high-quality music. MovieClips [2] comprises 20,000 videos sourced from the MovieClips YouTube channel. However, these datasets merely offer paired music and video data without detailed annotations, limiting their utility in enhancing MLLM capabilities. Some datasets [4, 26, 35–39] provide annotations for video-music rhythm matching, such as BGM909 [4] providing short music descriptions, music chords, and beats, but they lack analysis of emotional alignment and semantic transitions. MMTrail [6] provides trailer videos and includes descriptions for MLLM instruction tuning, but it does not thoroughly investigate the video-music relationship. Musechat [3] and YT8M-MusicTextClips [29] automatically formulated music recommendation dialogues. None of these datasets deliver cohesive and multi-dimensional reasoning on the intricate video-music relationships. The temporal synchronization that enhances the harmony between music and visual narratives remains largely unexplored.

2.3. Video Datasets and Benchmarks

Traditional Vision-Language (VL) benchmarks [40–44] have primarily focused on specific capabilities such as multimodal retrieval and vision question answering (QA). The advent of multimodal large language models (MLLMs) has spurred the development of benchmarks designed to assess more integrated VL tasks [7–9, 45–51]. For instance, VideoMME [13], MM-Vet [12], Q-Bench [14], EgoSchema [10], and MMBench [11] emphasize comprehensive VL skills. These benchmarks introduce evaluation metrics that go beyond simple model hierarchies, providing a more nuanced assessment of model performance across a range of vision-language tasks.

2.4. Multimodal Large Language Models

Video large language models have evolved significantly from captioning tools like BLIP2 [52] to more advanced systems such as VideoChat [53] and Video-LLaVA [17], which demonstrate capabilities in dialogue generation and question-answering [15, 16, 18, 53–55]. Increasingly, models are also incorporating audio modalities [56, 57]. Examples include VideoLLaMA2 [21], video-SALMONN [22], Macaw-LLM [24], and VALOR [23], which can analyze both video and audio content and provide open-ended text outputs. These methods leverage powerful language models and can provide a deeper understanding of the relationship between video, audio, and text content, going beyond mere video-audio matching.

3. The HarmonySet Dataset & Benchmark

HarmonySet is designed to advance the understanding of video-music relationships by examining how background music aligns with and enhances visual narratives. This dataset emphasizes key aspects of synchronization and semantic alignment, focusing on temporal dynamics, rhythm, theme, emotion, and cultural relevance. In this section, we describe the data collection and annotation process, present dataset statistics, and discuss quality control measures. We demonstrate that **HarmonySet** is a pioneering resource for studying video-music alignment, offering rich insights into the synchronization between music and visual storytelling.

3.1. Video Collection

To ensure a diverse and high-quality collection of video-music pairs, we implemented a hierarchical tagging structure to facilitate the identification of videos that feature well-aligned background music. This structure includes primary categories such as *Life & Emotions*, *Arts & Performance*, *Travel & Events*, *Sports & Fitness*, *Knowledge*, and *Technology & Fashion*, each of which represents a broad genre, format, and cultural expression. These categories are further subdivided into 43 specific subcategories (see Figure 2, left). In addition, we generated 293 relevant keywords derived from these subcategories to guide our video search process. Using these keywords, we crawled videos from YouTube Shorts, ensuring a variety of music genres and visual content. The dataset exclusively includes videos with user-added background music that complements the visual content. To ensure data consistency, annotators manually reviewed the collected videos to remove those lacking music. To verify the presence of music, we employed the PANNs [58] model, which confirmed that 83% of the videos from our search contained music. Videos without music were excluded to maintain dataset integrity.

3.2. Annotation Construction

To make **HarmonySet** a valuable resource for research on video-music relationships, we implemented a multi-phase annotation process that captures various aspects of the audiovisual content. The annotation process consists of two primary phases: manual annotation by trained annotators and automated refinement using machine-generated annotations.

3.2.1. Manual Annotation

Manual annotation includes two main components: synchronization with timestamps and multi-dimensional label assignment.

Synchronization Annotation: Annotators identify key moments in the video, such as transitions or shifts in the visual narrative (e.g., scene changes or plot twists). They assess whether the music changes at these points and whether these changes align with the visual transitions, marking the timestamps for temporal synchronization.

Labeling: A structured labeling system is used to evaluate the relationship between the video and music across four dimensions: rhythm and synchronization, theme and content, emotion, and cultural relevance. Each label reflects the extent of alignment between the music and video. For example, in the *content alignment* dimension, possible labels include "strongly related," "indirectly related," "unrelated," and "conflicting." In the *narrative enhancement* dimension, labels include "enhancing," "suggesting," "reversing," "independent narrative," and "no supplement." Annotators select the most appropriate label for each dimension, providing a nuanced and multi-faceted understanding of the video-music relationship.

Each video is annotated by three independent annotators to ensure objectivity and reliability. The final annotations are derived from the consensus among these annotators, minimizing individual biases and enhancing the robustness of the dataset.

3.2.2. Quality Control

A rigorous quality review process was implemented to ensure accurate and reliable annotations. A dedicated reviewer cross-checked each annotation for key timestamp accuracy consistency and factual grounding of the labels. This process mitigates potential biases and ensures data quality.

3.2.3. Automated Annotation Curation

Following manual annotation, we employed *Gemini 1.5 Pro* [19] to generate enhanced annotations. The inputs for this process included the video and audio content, the manually verified annotations (used as ground truth), and the video metadata (e.g., titles, and descriptions). The system was tasked with generating detailed descriptions of the video-music relationship, focusing on the four key dimensions: rhythm and synchronization, theme and content, emotion, and cultural relevance. The final output provides temporally aligned, multi-dimensional annotations that offer deeper insights into the alignment between video and music. Specific generation prompts and further details on the automated annotation process can be found in the appendix.

3.3. Instruction Tuning Dataset Statistics

To further enhance the utility of the HarmonySet dataset for training multimodal models, we created an instruction-tuning dataset. This dataset includes structured annotations that provide detailed explanations of the video-music relationships, enabling the fine-tuning of models for MLLMs to better understand video-music content.

The HarmonySet instruction tuning dataset consists of **44,470** video-music pairs, each with an annotation that provides a structured explanation of the video-music connections. Figure 2 (middle & right) illustrates the distribution of annotation length and video durations within the HarmonySet dataset. The videos are between 2.96 and 63.38 seconds in length, with an average duration of 31.5 seconds, contributing to a total of **458.8** hours of video and music content. The average number of words in HarmonySet annotations is **352.65** words for each video-music pair.

3.4. HarmonySet Benchmark

HarmonySet-OE HarmonySet-OE is comprised of 3,858 video-music pairs along with their accompanying annotations. MLLMs are required to address the video-music alignment relationships, including both temporal synchronization and semantic matching. The expected responses are designed to be open-ended to cover the diverse angles of video-music

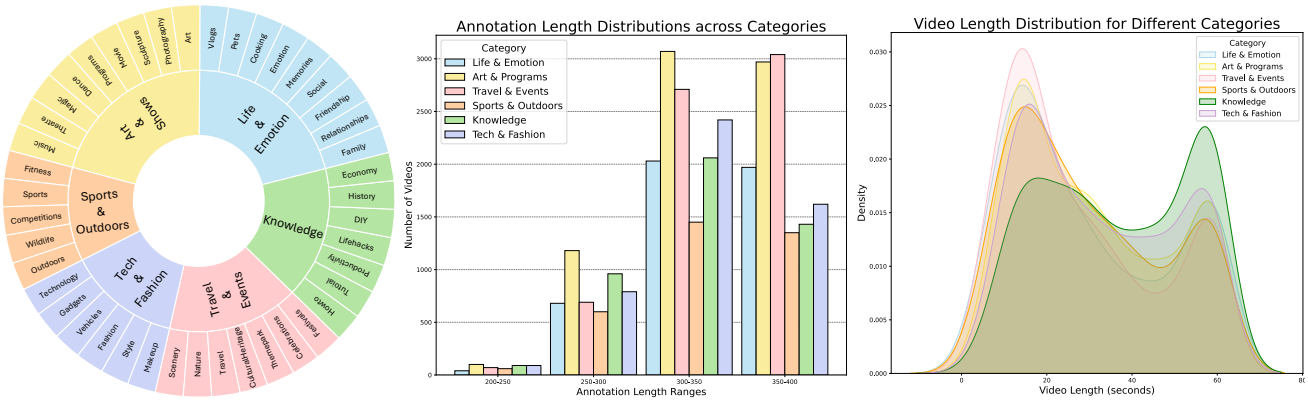

Figure 2. **HarmonySet Statistics.** (Left) HarmonySet covers 6 main categories and is divided into 43 subclasses with a full spectrum of content types. (Middle) Distributions of the number of words across categories in HarmonySet annotations. HarmonySet has a balanced annotation length across 6 main categories. (Right) Video duration distributions for different categories. The video durations are concentrated between 10 seconds and 60 seconds, with a rich number of videos in each time segment.

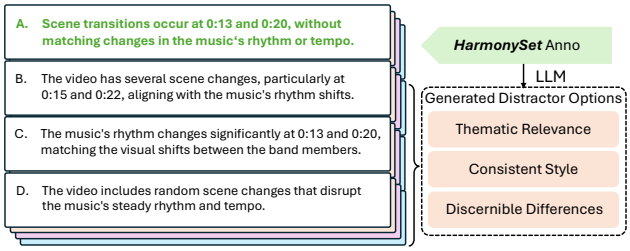

Figure 3. An example of HarmonySet-MC curation. We used LLM to convert open-ended annotations into multiple-choice options, with HarmonySet annotations serving as the correct options. Wrong options are constructed to be challenging yet distinguishable from the correct option.

relationships. Traditional language metrics like BLEU-4 [59] and ROUGE-L [60] are only sensitive to lexical variations and cannot identify changes in sentence semantics. Recent study [61] has proved LLM [62] to be a reliable evaluation tool for open-ended responses. Therefore, MLLM scores are obtained by comparing the MLLM outputs with the ground truth provided by HarmonySet-OE using LLM. The specific prompt of LLM for evaluation can be found in Appendix.

HarmonySet-MC We further developed HarmonySet-MC, a multiple-choice extension of HarmonySet-OE, to facilitate a more structured and objective evaluation process. Specifically, we instructed *GPT-4o* to use the annotated answer as the correct option and create three wrong options. These distracting options were carefully crafted to meet the following criteria: 1) Maintain thematic relevance to the correct answer, avoiding overly obvious discrepancies, 2) Resemble the correct answer in length and sentence structure, avoiding superficial distinctions, and 3) Present discernible semantic differences compared to the correct answer, see Figure 3 for an example. HarmonySet-MC includes the

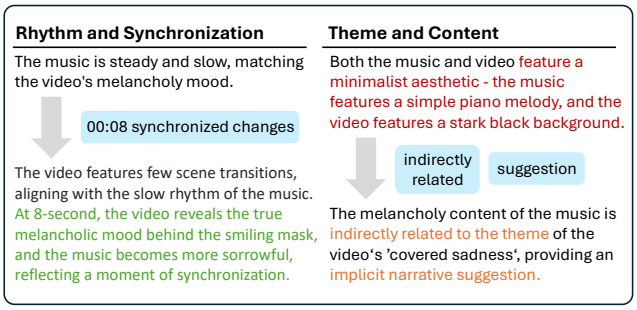

Figure 4. An example of annotation before and after the introduction of manual labels. The red text highlights an unreasonable explanation that arises in the absence of human guidance.

same 3,858 video-music pairs with HarmonySet-OE, each with four multiple-choice questions corresponding to four distinct aspects of rhythm, emotion, theme, and cultural context. HarmonySet-MC offers a convenient evaluation tool, allowing direct assessment of model performance using multiple-choice accuracy.

3.5. Dataset Assessment

Due to the potential influence of individual biases on the annotations, we conducted a consensus evaluation after completing all annotations to ensure agreement and objectivity. The results, detailed in Table 3, show that HarmonySet gets 92% high consensus.

Figure 4 illustrates the comparison of a video's two aspects of annotation before and after injecting manual labels. Manual annotations provide more accurate semantic alignment understanding and incorporate specific temporal synchronization information. Without human knowledge, models tend to generate spurious video-music connections or meaningless responses. Experiment in Table 4 also confirms the

Table 2. **Main Results on HarmonySet-OE.** We tested Gemini-1.5 Pro and open source MLLMs including VideoLLaMA2 and video-SALMONN. The bottom part presents results of VideoLLaMA2 finetuned on our instruction tuning dataset. While Gemini-1.5 Pro leads among untrained models, VideoLLaMA2 finetuned with HarmonySet demonstrates significant improvement and a strong understanding of video-music alignment. Results on synchronization can be found in R & S (Rhythm & Synchronization), and semantic matching results consist of scores in T (Theme), E (Emotion), and C (Culture).

Models	LLM	Metrics	Life & Emotion	Art & Performance	Travel & Events	Sports & Outdoors	Knowledge	Tech & Fashion	Overall
- Close-source MLLM									
Gemini-1.5 Pro [63]	-	R & S	5.30	6.05	5.69	4.94	4.91	4.98	5.43
		T	5.25	5.76	5.75	4.41	4.46	4.49	**5.18**
		E	5.28	5.75	5.60	4.59	4.45	4.43	5.15
		C	4.64	4.91	4.77	3.85	4.27	4.03	4.51
- Open-source MLLMs									
VideoLLaMA2 [21]	Qwen2-7B	R & S	3.89	4.80	4.56	4.01	3.39	3.54	4.15
		T	4.09	4.83	4.93	3.89	3.44	3.71	4.29
		E	4.36	5.01	5.02	4.08	3.44	3.49	4.38
		C	2.95	3.46	3.69	2.56	2.32	2.52	3.05
video-SALMONN [22]	Vicuna-13B-v1.5	R & S	2.43	3.53	2.98	2.68	2.32	2.51	2.83
		T	3.24	4.18	3.97	3.23	2.96	3.00	3.55
		E	3.11	4.12	3.84	3.13	2.56	2.70	3.38
		C	1.85	2.51	2.51	1.77	1.68	1.84	2.12
- With HarmonySet									
VideoLLaMA2 (HarmonySet)	Qwen2-7B	R & S	5.43	6.35	6.03	4.94	5.33	4.83	**5.55**
		T	5.12	5.21	5.03	4.84	5.21	4.85	5.06
		E	5.25	6.41	5.84	4.00	4.88	4.47	**5.26**
		C	4.87	4.98	4.72	3.31	5.23	4.09	**4.62**

Table 3. Consensus evaluation result. We randomly sampled 10% of the data and had human reviewers select "Low", "Medium", or "High" as their level of agreement with the annotations. HarmonySet gets 92% high consensus, showing a high level of agreement.

	Low	Medium	High
Consensus	3%	5%	92%

significant positive impact of human annotation on improving video-music understanding.

3.6. Dataset Property

HarmonySet emphasis on temporal synchronization. One crucial factor of video-music connection lies in temporal synchronization. For example, the music becomes more intense just as an athlete makes a final sprint. Such synchronized changes contribute significantly to an immersive visual-music experience. Static images paired with music can only achieve content or mood matching, lacking the ability to express such dynamic transitions. HarmonySet not only includes detailed annotations on the overall pace suitability and beat matching but also provides timestamped explanations of video-music transitions. 58% of the data in HarmonySet contains key timestamp annotations, providing valuable support for understanding temporal relationships between video and music.

HarmonySet provides deep semantic alignment understanding. The semantic resonance between video and music often manifests as a subtle connection that is difficult to articulate. HarmonySet categorizes this semantic alignment into four dimensions, providing a comprehensive framework for understanding these complex relationships.

4. Experiments

4.1. Baselines

We conduct the evaluation on Gemini 1.5 Pro [19] and state-of-the-art open-source video-audio MLLMs, including VideoLLaMA2 [21] and video-SALMONN [22]. For a fair comparison, we adopt the zero-shot setting to infer HarmonySet-OE questions with all MLLMs based on the same prompt. In the experiments presented in Table 2, we used a consistent 16 frames for the video input of open-source models for both inference and fine-tuning. A special case is Gemini 1.5 Pro, which supports relatively long multimodal contexts, and videos are sampled at 1 frame per second for the input. In the Appendix, we provide detailed information regarding the architecture and the parameter size for all open-source MLLMs evaluated in this paper, as well as additional results for more MLLMs under various settings.

4.2. Main Results

Table 2 shows the main results on HarmonySet-OE. Gemini-1.5 Pro generally outperforms untrained open-source MLLMs across all categories and metrics, significantly exceeding the second-best model in Rhythm Synchronization (by 1.28) and Culture (by 1.46). This might be due to its

Table 4. Comparison between performance of VideoLLaMA2 trained on fully automated data and HarmonySet data. Models trained with our instruction tuning data demonstrates a clear advantage, validating the value of human expertise in providing rich information on synchronization and semantic alignment and the effectiveness of our human-machine collaborative framework.

Models	Metrics	Life & Emotion	Art & Performance	Travel & Events	Sports & Outdoors	Knowledge	Tech & Fashion	Overall
VideoLLaMA2 (F.A., 10k)	R & S	4.56	5.05	4.97	4.28	4.20	4.17	4.59
	T	4.20	5.01	4.85	3.49	3.36	3.41	4.16
	E	4.29	4.76	4.67	4.03	3.76	3.79	4.28
	C	3.53	3.98	3.67	2.93	3.13	3.02	3.44
VideoLLaMA2 (HarmonySet, 10k)	R & S	4.69	5.58	5.30	4.49	4.36	4.25	**4.86**
	T	4.66	5.02	4.98	4.40	4.39	4.29	**4.70**
	E	4.64	5.43	5.26	4.06	3.85	3.78	**4.66**
	C	3.99	4.30	4.25	2.97	3.79	3.34	**3.89**

Table 5. Human and model performance on HarmonySet-MC. While VideoLLaMA2 tuned on HarmonySet surpasses Gemini-1.5 Pro in certain aspects, it still falls short of human performance, highlighting both the challenging nature of our task and the limitations of current models.

	R & S (Acc.)	T (Acc.)	E (Acc.)	C (Acc.)
Gemini-1.5 Pro	41.84%	45.45%	44.43%	50.40%
Video-LLaMA2	21.76%	48.95%	52.76%	24.29%
Video-LLaMA2 (HarmonySet)	10.63%	54.16%	47.32%	36.66%
Human	**85.26%**	**88.19%**	**84.49%**	**93.81%**

Table 6. Results of VideoLLaMA2 trained on HarmonySet with 16, 32, and 64 frames. Using 64 frames yields the lowest scores, indicating potential redundancy or even negative effects from excessive visual input within short (<1 minute) videos.

	R & S	T	E	C
16 Frames	5.55	5.06	5.26	4.62
32 Frames	**5.59**	**5.08**	**5.29**	**4.65**
64 Frames	5.49	4.94	5.21	4.53

capacity for long context inputs and timestamped outputs, allowing for better alignment with human annotations that often consider temporal relationships. The performance in culturally nuanced pairings might stem from Gemini-1.5 Pro's extensive training data.

Untrained VideoLLaMA2 and video-SALMONN both underperform Gemini-1.5 Pro. While VideoLLaMA2 shows moderate rhythmic synchronization (4.15), its semantic matching capabilities are weaker, particularly in cultural understanding (3.05). This suggests a deficiency in comprehending nuanced cultural differences and contextual information. Video-SALMONN consistently scores lowest across all metrics and categories, struggling with understanding temporal synchronization, emotional congruence, thematic integration, and cultural relevance. The open-source models' weaker performance likely stems from having less training data (in both quantity and quality, especially regarding cultural nuances) and limited input capacity, hindering analysis of complex relationships and rhythmic synchronization requiring longer temporal contexts.

Training VideoLLaMA2 on HarmonySet yields substantial improvements, boosting its Rhythm Synchronization score by 1.40 and Culture score by 1.57. This surpasses previous state-of-the-art results in most domains, demonstrating HarmonySet's effectiveness in addressing limitations of models trained on lower-quality data and fostering deeper multimodal understanding.

4.3. HarmonySet-MC Results

We evaluated different models on HarmonySet-MC and invited individuals unfamiliar with the dataset to provide human performance. From Table 5, it is evident that human performance significantly outperforms the best model results across all evaluation perspectives. This highlights the challenging nature of our dataset, which effectively measures the gap between model performance and human-level understanding. These results indicate that current models still struggle to effectively understand the complex interplay between video and music, underscoring the need for further advancements in model training and dataset annotation to bridge this performance gap. The complete experimental results on HarmonySet-MC can be found in the Appendix.

4.4. Ablations

Ablation on frame number. We designed experiments to investigate the impact of the number of frames used during training on the model's performance, as shown in Table 6. Increasing the frames from 16 to 32 improves the model's performance across all aspects. However, using 64 frames yields the worst performance, even lower than with 16 frames. This might be attributed to the focus of our videos on short-form content under one minute, where 64 frames could introduce increased computational complexity, potential overfitting, or information redundancy.

Can fully automated data provide sufficient capability? HarmonySet provides valuable video-music annotations

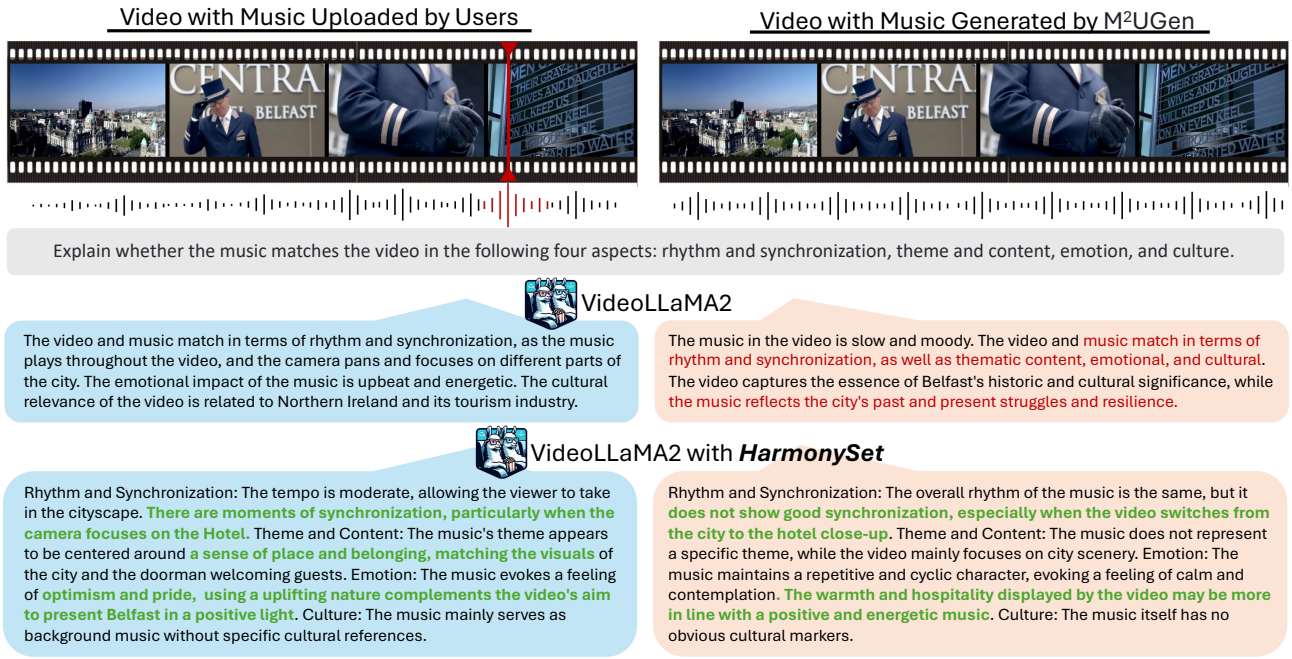

Figure 5. VideoLLaMA2's response before and after training with our instruction tuning dataset. The left video features human-composed soundtracks, while the right video is with AI-generated soundtracks. Without HarmonySet, the model often provides the wrong justification for the generated music for its harmony with the visual content (highlighted in red text). The trained model offers more insightful analysis and can effectively assess both human-composed and AI-generated music. Our dataset facilitates a deeper understanding of both synchronization and semantic alignment.

that enhance the multimodal understanding of MLLMs. To evaluate the dataset's effectiveness, we compared two types of training data: annotations fully generated by Gemini 1.5 Pro and those derived from our HarmonySet. Both training processes utilized 10,000 samples. Results show in Table 4 that training on the fully automated data yields minimal performance gains. Models trained with HarmonySet annotations consistently surpass those trained with auto-generated data, especially in capturing synchronization and semantic alignment between video and music. These results highlight the importance of a human-machine collaborative framework in enhancing multimodal models' video-music understanding.

4.5. Effectiveness on Assessing AI-Generated Music

Generating music for videos [64–68] is a highly challenging task that requires harmony between music and visual narratives in terms of synchronization and semantic alignment. Figure 5 shows VideoLLaMA2's improved ability to differentiate between human-composed and AI-generated soundtracks after training with HarmonySet. The vanilla VideoLLaMA2 struggled to justify the generated music's harmony with the video, likely due to a lack of understanding of what constitutes harmonious alignment. With HarmonySet, the model provides more detailed temporal and semantic analyses, enabling more objective evaluation. These results offer valuable insights for future video soundtrack creation.

5. Conclusion

We introduce HarmonySet, the first dataset focused on facilitating the ability of MLLMs in comprehensive video-music understanding. HarmonySet comprises a diverse domain of videos with high-quality music, each annotated with structured explanations detailing the semantic matching and temporal synchronization between video and music. Our extensive evaluation of state-of-the-art MLLMs, encompassing commercial and open-source models, reveals the limitations of MLLMs' reasoning about the interplay between visual and musical elements. This highlights the challenge of achieving in-depth video-music understanding. There are many exciting directions to build upon this work, including developing novel MLLM architectures specifically tailored for video-music analysis and investigating the potential for cross-modal knowledge transfer between video and music content. We hope HarmonySet will inspire future research and development in improving the capabilities of MLLMs.

References

[1] Jing Yi, Yaochen Zhu, Jiayi Xie, and Zhenzhong Chen. Cross-modal variational auto-encoder for content-based micro-video

background music recommendation. IEEE Transactions on Multimedia, 25:515–528, 2021. 2, 3

[2] Dídac Surís, Carl Vondrick, Bryan Russell, and Justin Salamon. It's time for artistic correspondence in music and video. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 10564–10574, 2022. 2, 3

[3] Zhikang Dong, Xiulong Liu, Bin Chen, Pawel Polak, and Peng Zhang. Musechat: A conversational music recommendation system for videos. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 12775–12785, 2024. 2, 3

[4] Sizhe Li, Yiming Qin, Minghang Zheng, Xin Jin, and Yang Liu. Diff-bgm: A diffusion model for video background music generation. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 27348–27357, 2024. 3

[5] Jiajie Teng, Huiyu Duan, Yucheng Zhu, Sijing Wu, and Guangtao Zhai. Mvbind: Self-supervised music recommendation for videos via embedding space binding. arXiv preprint arXiv:2405.09286, 2024. 2, 3

[6] Xiaowei Chi, Yatian Wang, Aosong Cheng, Pengjun Fang, Zeyue Tian, Yingqing He, Zhaoyang Liu, Xingqun Qi, Jiahao Pan, Rongyu Zhang, et al. Mmtrail: A multimodal trailer video dataset with language and music descriptions. arXiv preprint arXiv:2407.20962, 2024. 2, 3

[7] Shuai Bai, Shusheng Yang, Jinze Bai, Peng Wang, Xingxuan Zhang, Junyang Lin, Xinggang Wang, Chang Zhou, and Jingren Zhou. Touchstone: Evaluating vision-language models by language models. arXiv preprint arXiv:2308.16890, 2023. 2, 3

[8] 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, pages 22195–22206, 2024.

[9] 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, pages 13299–13308, 2024. 3

[10] Karttikeya 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. 3

[11] Yuan Liu, Haodong Duan, Yuanhan Zhang, Bo Li, Songyang Zhang, Wangbo Zhao, Yike Yuan, Jiaqi Wang, Conghui He, Ziwei Liu, et al. Mmbench: Is your multi-modal model an all-around player? In European Conference on Computer Vision, pages 216–233. Springer, 2025. 3

[12] Weihao Yu, Zhengyuan Yang, Linjie Li, Jianfeng Wang, Kevin Lin, Zicheng Liu, Xinchao Wang, and Lijuan Wang. Mm-vet: Evaluating large multimodal models for integrated capabilities. arXiv preprint arXiv:2308.02490, 2023. 3

[13] Chaoyou Fu, Yuhan 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. arXiv preprint arXiv:2405.21075, 2024. 3

[14] Haoning Wu, Zicheng Zhang, Erli Zhang, Chaofeng Chen, Liang Liao, Annan Wang, Chunyi Li, Wenxiu Sun, Qiong Yan, Guangtao Zhai, et al. Q-bench: A benchmark for general-purpose foundation models on low-level vision. arXiv preprint arXiv:2309.14181, 2023. 2, 3

[15] Muhammad Maaz, Hanoona Rasheed, Salman Khan, and Fahad Shahbaz Khan. Video-chatgpt: Towards detailed video understanding via large vision and language models. arXiv preprint arXiv:2306.05424, 2023. 2, 3

[16] Renrui Zhang, Jiaming Han, Chris Liu, Peng Gao, Aojun Zhou, Xiangfei Hu, Shilin Yan, Pan Lu, Hongsheng Li, and Yu Qiao. Llama-adapter: Efficient fine-tuning of language models with zero-init attention. arXiv preprint arXiv:2303.16199, 2023. 3

[17] Bin Lin, Yang Ye, Bin Zhu, Jiaxi Cui, Munan Ning, Peng Jin, and Li Yuan. Video-llava: Learning united visual representation by alignment before projection. arXiv preprint arXiv:2311.10122, 2023. 3

[18] Bo Li, Yuanhan Zhang, Dong Guo, Renrui Zhang, Feng Li, Hao Zhang, Kaichen Zhang, Yanwei Li, Ziwei Liu, and Chunyuan Li. Llava-onevision: Easy visual task transfer. arXiv preprint arXiv:2408.03326, 2024. 3

[19] Gemini Team, Petko Georgiev, Ving Ian Lei, Ryan Burnell, Libin Bai, Anmol Gulati, Garrett Tanzer, Damien Vincent, Zhufeng Pan, Shibo Wang, et al. Gemini 1.5: Unlocking multimodal understanding across millions of tokens of context. arXiv preprint arXiv:2403.05530, 2024. 2, 4, 6

[20] RHYTHMIC SYNCHRONIZATION. Muvi: Video-to-music generation with seman-tic alignment and rhythmic synchronization. 2

[21] 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. 2, 3, 6

[22] 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. 3, 6

[23] Sihan Chen, Xingjian He, Longteng Guo, Xinxin Zhu, Weining Wang, Jinhui Tang, and Jing Liu. Valor: Vision-audio-language omni-perception pretraining model and dataset. arXiv preprint arXiv:2304.08345, 2023. 3

[24] Chenyang Lyu, Minghao Wu, Longyue Wang, Xinting Huang, Bingshuai Liu, Zefeng Du, Shuming Shi, and Zhaopeng Tu. Macaw-llm: Multi-modal language modeling with image, audio, video, and text integration. arXiv preprint arXiv:2306.09093, 2023. 2, 3

[25] Sungeun Hong, Woobin Im, Jongbin Ryu, and Hyun S Yang. Sspp-dan: Deep domain adaptation network for face recognition with single sample per person. In 2017 IEEE International Conference on Image Processing (ICIP), pages 825–829. IEEE, 2017. 2, 3

[26] Bochen Li, Xinzhao Liu, Karthik Dinesh, Zhiyao Duan, and Gaurav Sharma. Creating a multitrack classical music performance dataset for multimodal music analysis: Challenges, insights, and applications. IEEE Transactions on Multimedia, 21(2):522–535, 2018. 3

[27] Zeyue Tian, Zhaoyang Liu, Ruibin Yuan, Jiahao Pan, Qifeng Liu, Xu Tan, Qifeng Chen, Wei Xue, and Yike Guo. Vidmuse: A simple video-to-music generation framework with long-short-term modeling. arXiv preprint arXiv:2406.04321, 2024. 3

[28] 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, pages 19108–19118, 2022. 2, 3

[29] Daniel McKee, Justin Salamon, Josef Sivic, and Bryan Russell. Language-guided music recommendation for video via prompt analogies. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 14784–14793, 2023. 2, 3

[30] Andres Ferraro, Xavier Serra, and Christine Bauer. Break the loop: Gender imbalance in music recommenders. In Proceedings of the 2021 conference on human information interaction and retrieval, pages 249–254, 2021. 2

[31] Jort F Gemmeke, Daniel PW Ellis, Dylan Freedman, Aren Jansen, Wade Lawrence, R Channing Moore, Manoj Plakal, and Marvin Ritter. Audio set: An ontology and human-labeled dataset for audio events. In 2017 IEEE international conference on acoustics, speech and signal processing (ICASSP), pages 776–780. IEEE, 2017. 3

[32] Honglie Chen, Weidi Xie, Andrea Vedaldi, and Andrew Zisserman. Vggsound: A large-scale audio-visual dataset. In ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), pages 721–725. IEEE, 2020. 3

[33] Eduardo Fonseca, Xavier Favory, Jordi Pons, Frederic Font, and Xavier Serra. Fsd50k: an open dataset of human-labeled sound events. IEEE/ACM Transactions on Audio, Speech, and Language Processing, 30:829–852, 2021. 3

[34] Karol J Piczak. Esc: Dataset for environmental sound classification. In Proceedings of the 23rd ACM international conference on Multimedia, pages 1015–1018, 2015. 3

[35] Ruilong Li, Shan Yang, David A Ross, and Angjoo Kanazawa. Ai choreographer: Music conditioned 3d dance generation with aist++. In Proceedings of the IEEE/CVF International Conference on Computer Vision, pages 13401–13412, 2021. 3

[36] Le Zhuo, Zhaokai Wang, Baisen Wang, Yue Liao, Chenxi Bao, Stanley Peng, Songhao Han, Aixi Zhang, Fei Fang, and Si Liu. Video background music generation: Dataset, method and evaluation. In Proceedings of the IEEE/CVF International Conference on Computer Vision, pages 15637–15647, 2023.

[37] Jiashuo Yu, Yaohui Wang, Xinyuan Chen, Xiao Sun, and Yu Qiao. Long-term rhythmic video soundtracker. In International Conference on Machine Learning, pages 40339–40353. PMLR, 2023.

[38] Kaixing Yang, Xukun Zhou, Xulong Tang, Ran Diao, Hongyan Liu, Jun He, and Zhaoxin Fan. Beatdance: A beat-based model-agnostic contrastive learning framework for music-dance retrieval. In Proceedings of the 2024 International Conference on Multimedia Retrieval, pages 11–19, 2024.

[39] Ye Zhu, Kyle Olszewski, Yu Wu, Panos Achlioptas, Menglei Chai, Yan Yan, and Sergey Tulyakov. Quantized gan for complex music generation from dance videos. In European Conference on Computer Vision, pages 182–199. Springer, 2022. 3

[40] Jun Xu, Tao Mei, Ting Yao, and Yong Rui. Msr-vtt: A large video description dataset for bridging video and language. In Proceedings of the IEEE conference on computer vision and pattern recognition, pages 5288–5296, 2016. 3

[41] Dejing Xu, Zhou Zhao, Jun Xiao, Fei Wu, Hanwang Zhang, Xiangnan He, and Yueting Zhuang. Video question answering via gradually refined attention over appearance and motion. In Proceedings of the 25th ACM international conference on Multimedia, pages 1645–1653, 2017.

[42] Raghav Goyal, Samira Ebrahimi Kahou, Vincent Michalski, Joanna Materzynska, Susanne Westphal, Heuna Kim, Valentin Haenel, Ingo Fruend, Peter Yianilos, Moritz Mueller-Freitag, et al. The" something something" video database for learning and evaluating visual common sense. In Proceedings of the IEEE international conference on computer vision, pages 5842–5850, 2017.

[43] Will Kay, Joao Carreira, Karen Simonyan, Brian Zhang, Chloe Hillier, Sudheendra Vijayanarasimhan, Fabio Viola, Tim Green, Trevor Back, Paul Natsev, et al. The kinetics human action video dataset. arXiv preprint arXiv:1705.06950, 2017.

[44] Junbin Xiao, Xindi Shang, Angela Yao, and Tat-Seng Chua. Next-qa: Next phase of question-answering to explaining temporal actions. In Proceedings of the IEEE/CVF conference on computer vision and pattern recognition, pages 9777–9786, 2021. 3

[45] Kaining Ying, Fanqing Meng, Jin Wang, Zhiqian Li, Han Lin, Yue Yang, Hao Zhang, Wenbo Zhang, Yuqi Lin, Shuo Liu, et al. Mmt-bench: A comprehensive multimodal benchmark for evaluating large vision-language models towards multitask agi. arXiv preprint arXiv:2404.16006, 2024. 3

[46] Junjie Zhou, Yan Shu, Bo Zhao, Boya Wu, Shitao Xiao, Xi Yang, Yongping Xiong, Bo Zhang, Tiejun Huang, and Zheng Liu. Mlvu: A comprehensive benchmark for multi-task long video understanding. arXiv preprint arXiv:2406.04264, 2024.

[47] Haoning Wu, Dongxu Li, Bei Chen, and Junnan Li. Longvideobench: A benchmark for long-context interleaved video-language understanding. arXiv preprint arXiv:2407.15754, 2024.

[48] Yu Lu, Ruijie Quan, Linchao Zhu, and Yi Yang. Exploiting unlabeled videos for video-text retrieval via pseudo-supervised learning. IEEE Transactions on Image Processing, 2024.

[49] Yu Lu, Feiyue Ni, Haofan Wang, Xiaofeng Guo, Linchao Zhu, Zongxin Yang, Ruihua Song, Lele Cheng, and Yi Yang. Show me a video: A large-scale narrated video dataset for coherent story illustration. IEEE Transactions on Multimedia, 2023.

[50] Yu Lu, Ruijie Quan, Linchao Zhu, and Yi Yang. Zero-shot video grounding with pseudo query lookup and verifica-

tion. IEEE Transactions on Image Processing, 33:1643–1654, 2024.

[51] Binzhu Xie, Sicheng Zhang, Zitang Zhou, Bo Li, Yuanhan Zhang, Jack Hessel, Jingkang Yang, and Ziwei Liu. Funqa: Towards surprising video comprehension. In European Conference on Computer Vision, pages 39–57. Springer, 2024. 3

[52] 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, pages 19730–19742. PMLR, 2023. 3

[53] KunChang Li, Yinan He, Yi Wang, Yizhuo Li, Wenhai Wang, Ping Luo, Yali Wang, Limin Wang, and Yu Qiao. Videochat: Chat-centric video understanding. arXiv preprint arXiv:2305.06355, 2023. 3

[54] Jingkang Yang, Yuhao Dong, Shuai Liu, Bo Li, Ziyue Wang, Haoran Tan, Chencheng Jiang, Jiamu Kang, Yuanhan Zhang, Kaiyang Zhou, et al. Octopus: Embodied vision-language programmer from environmental feedback. In European Conference on Computer Vision, pages 20–38. Springer, 2024.

[55] Hang Zhang, Xin Li, and Lidong Bing. Video-LLaMA: An instruction-tuned audio-visual language model for video understanding. In Yansong Feng and Els Lefever, editors, Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing: System Demonstrations, pages 543–553, Singapore, December 2023. Association for Computational Linguistics. 3

[56] Sihan Chen, Handong Li, Qunbo Wang, Zijia Zhao, Mingzhen Sun, Xinxin Zhu, and Jing Liu. Vast: A vision-audio-subtitle-text omni-modality foundation model and dataset. Advances in Neural Information Processing Systems, 36:72842–72866, 2023. 3

[57] Bin Zhu, Bin Lin, Munan Ning, Yang Yan, Jiaxi Cui, HongFa Wang, Yatian Pang, Wenhao Jiang, Junwu Zhang, Zongwei Li, et al. Languagebind: Extending video-language pretraining to n-modality by language-based semantic alignment. arXiv preprint arXiv:2310.01852, 2023. 3

[58] Qiuqiang Kong, Yin Cao, Turab Iqbal, Yuxuan Wang, Wenwu Wang, and Mark D Plumbley. Panns: Large-scale pre-trained audio neural networks for audio pattern recognition. IEEE/ACM Transactions on Audio, Speech, and Language Processing, 28:2880–2894, 2020. 4, 1

[59] Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. Bleu: a method for automatic evaluation of machine translation. In Proceedings of the 40th annual meeting of the Association for Computational Linguistics, pages 311–318, 2002. 5, 2

[60] Chin-Yew Lin. Rouge: A package for automatic evaluation of summaries. In Text summarization branches out, pages 74–81, 2004. 5, 3

[61] Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric Xing, et al. Judging llm-as-a-judge with mt-bench and chatbot arena. Advances in Neural Information Processing Systems, 36:46595–46623, 2023. 5

[62] OpenAI, Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, et al. Gpt-4 technical report, 2024. 5

[63] 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. 6

[64] Shansong Liu, Atin Sakkeer Hussain, Chenshuo Sun, and Ying Shan. M^{2} ugen: Multi-modal music understanding and generation with the power of large language models. arXiv preprint arXiv:2311.11255, 2023. 8

[65] Kun Su, Xiulong Liu, and Eli Shlizerman. How does it sound? generation of rhythmic soundtracks for human movement videos. In Conf. Neural Inf. Process. Syst, volume 35, pages 0–10, 2021.

[66] Wei Guo, Heng Wang, Weidong Cai, and Jianbo Ma. Gotta hear them all: Sound source aware vision to audio generation. arXiv preprint arXiv:2411.15447, 2024.

[67] Xinhao Mei, Varun Nagaraja, Gael Le Lan, Zhaoheng Ni, Ernie Chang, Yangyang Shi, and Vikas Chandra. Foley-gen: Visually-guided audio generation. In 2024 IEEE 34th International Workshop on Machine Learning for Signal Processing (MLSP), pages 1–6. IEEE, 2024.

[68] Xiulong Liu, Kun Su, and Eli Shlizerman. Tell what you hear from what you see–video to audio generation through text. arXiv preprint arXiv:2411.05679, 2024. 8

[69] Thibault Sellam, Dipanjan Das, and Ankur P Parikh. Bleurt: Learning robust metrics for text generation. arXiv preprint arXiv:2004.04696, 2020. 3

[70] Ramakrishna Vedantam, C Lawrence Zitnick, and Devi Parikh. Cider: Consensus-based image description evaluation. In Proceedings of the IEEE conference on computer vision and pattern recognition, pages 4566–4575, 2015. 3

[71] Zhibiao Wu and Martha Palmer. Verb semantics and lexical selection. arXiv preprint cmp-lg/9406033, 1994. 3

HarmonySet: A Comprehensive Dataset for Understanding Video-Music Semantic Alignment and Temporal Synchronization

Supplementary Material

A. More details of HarmonySet

A.1. Dataset Construction pipeline

Videos in HarmonySet are sourced from the YouTube Shorts platform and were crawled using 293 keywords we designed. The complete list of keywords is shown in Figure 6.

Figure 7 illustrates the multi-phase annotation process. The raw data we crawled includes videos along with their audio and metadata. The metadata includes the title, author, duration, and video width and height, as shown in Figure 8. We used PANNs [58] for music tagging and video filtering. The music tags generated by PANNs were added to the metadata, and videos without music were filtered out based on the tagging results. The filtering criterion was: if the labels with top 2 probabilities do not include 'music', the video is deleted.

The filtered videos were then assigned to human annotators for detailed screening to **ensure video-music pair quality** and to **exclude non-ethical and sensitive content**. The instructions for the annotators were:

Music Check: *If there is no background music (e.g., pure human voice, pure environmental sound, pure noise, no sound, etc. Music mixed with human voice counts as having music), please flag the video. Listen to the entire video before making this determination, as music may only be present in a portion of the video.*

Content Suitability Check: *Carefully review the entire video for any content that is: Non-Ethical: This includes, but is not limited to, content that promotes or depicts illegal activities, harmful behavior, or discrimination. Sensitive: This includes, but is not limited to, content that is sexually suggestive, graphically violent, or exploits, abuses, or endangers children.*

Video Quality Check: *Please also assess the overall video quality. Flag any videos with technical issues, such as severe distortion, extremely poor resolution, or corrupted files.*

Human Annotation We conducted a rigorous annotator selection process. We recruited 120 annotators to pre-annotate 500 videos. The 120 annotators are all experts who have previous formal experience in video annotation work. After the pre-annotation, we retained 25 individuals who demonstrated both accuracy, diversity, and speed in their annotations. For videos with music, human annotators were to mark key time points and label tags. In the key time point annotation, annotators first identified moments representing visual narrative turning points or key points, then determined whether the music synchronously changed with the video at those moments. The instructions were:

Please mark up to three important time points in the video. If there are no changes throughout the video, fill in 0. Then determine: A) The music changes precisely in sync with the video at the turning point; B) The music changes near the turning point but is not strongly synchronized; C) The music does not change when the video turns. The answer format should be: video timestamp + comma + uppercase letter option, separated by semicolons between time points. Example: 00:10,A; 00:20,B; 00:30,C (Non-synchronization means the visual changes but the music remains the same. Examples of synchronized music changes include [outfit change on beat], [music changes to a victorious tune after a basketball shot], [music reaches a climax as the video reaches its most exciting moment], etc.)

For label tagging, the structured label system is shown in Figure 9.

Automatic Annotation After human annotation, in the automatic annotation phase, the MLLM will receive the video and audio content, human annotation results, and required metadata as input. An example of the metadata is shown in Figure 8, where the video title and audio tags will be used in the automatic annotation process. The MLLM will generate detailed video-music alignment annotations, including semantic alignment and temporal synchronization understanding. We use Gemini 1.5 Pro as the MLLM for the automatic annotation phase, with specific instructions shown in Figure 10. In addition, to ensure the diversity of instructions and to avoid overestimation of performance, we employed multiple prompt templates for the instruction tuning data, as illustrated in Figure 11. These instructions convey the same underlying meaning while avoiding rigid patterns in sentence structure and word choice. The instructions in the dataset will be randomly assigned to one of these ten templates, promoting variability and enhancing the robustness of the training process. This not only helps in capturing a wider range of expressions but also mitigates the risk of the model becoming overly reliant on specific phrasing, thereby improving its generalization capabilities.

A.2. More statistics

The raw data crawled from the platform consists of 59,771 video-music pairs. After the first round of filtering, the number of video-music pairs was reduced to 49,610. Following a meticulous manual screening, the total number of videos was further reduced to 48,328. In addition to statistics including video categories, video duration, and the number of words

Main categories	Subclasses	Keywords
Life & Emotion	Family	childcare, familytime, familyfun, toddler, Babysitting, Preschool, Babycare, FamilyActivities, Parenting
	Relationships	lovestory, relationship, romance, dating, Lovejourney, Partnership, Courtship, couple, engagement
	Friendship	friendsforever, Bestfriends, ForeverFriends
	Social	Routine, Moments, Highlights, Challenge, Funfact, Social
	Memories	throwback, Nostalgia, Flashback, Flashbacks, Memories
	Emotion	emotions, feelings, sentiments, passion, emotional, Mood
	Cooking	baking, recipes, foodie, cuisine, Cookingrecipes, HomeBaking, FoodLover, CookingIdeas, FoodCulture, Gourmet, Gastronomy, cooking
	Pets	petlovers, AnimalLovers, PetCare, FurryFriends, pets
	Vlogs	dailyvlog, DailyLife, vloglife, TravelVlog
Art & Shows	Art	3DArt, illustration, art, graffiti
	Photography	streetphotography, portrait, photoshoot, photography, snapshots
	Sculpture	sculpture
	Movie	movie
	Programs	Programs, shows, performances
	Dance	LatinDance, HipHopDance, MusicalTheatre, ballet, hiphop, streetdance, choreography, latin, Dance
	Magic	illusion, trick, MagicTrick, MagicShow, magic
	Theatre	theatre, MusicalTheatre
	Music	instrumentals, Popmusic, BluesMusic, JazzMusic, electronicmusic, livemusic, guitar, piano, songwriter, rock, vocals, musical, jazz, pop, blues, melody, melodies, Concerts, Keyboard, Rockmusic, LiveShow, Showcase, melody, melodies, Music, Acoustic
Travel & Events	Scenery	Scenery, vista
	Nature	wildlife, nationalparks, NatureParks, naturesounds, NaturePhotography
	Travel	Sightseeing, Travelbug, Roadtrips, roadtrip, wanderlust, tourist, exploration, CityTour, destination, TravelVlog, backpacking
	CulturalHeritage	cultural relevance, HistoricalHeritage, heritage, CulturalHeritage
	Themepark	amusementpark, Funfair, themepark
	Celebrations	Xmas, CNY, FestiveSeason, Carnival, holidayseason, valentinesday, chinesenewyear, christmas, mothersday, fathersday, Celebrations, honeymoon, anniversary
	Festivals	Festival, festivals
Sports & Outdoors	Fitness	FitnessTips, WorkoutPlan, gymtime, training, cardio, bodybuilding, strength, meditation, Fitness
	Sports	WinterSports, SnowSports, Golfing, Diving, swimming, martialarts, athletics, tennis, golf, gym, workout, yoga, hiking, skiing, basketball, soccer, football, Crossfit, Sports
	Competitions	Competition, Competitions
	Wildlife	Wildlife
	Outdoors	Outdoors, Countryside
Tech & Fashion	Technology	TechUpdates, TechNews, TechLife, Technology, innovation, TechTrends
	Gadgets	Gadgets, TechGadgets, TechReviews
	Vehicles	automobile, carlovers, CarReviews, CarEnthusiasts, CarCulture, racing, car, Vehicles
	Fashion	attire, ensemble, Lookbook, Trends, FashionShow, fashiontrends, trendy, Fashion
	Style	style, outfit, accessories
	Makeup	MakeupLooks, cosmetics, beauty, beautyhacks, BeautyTips, Skincare, Makeup, BeautySecrets
Knowledge	Howto	howto, howtomake, Howto
	Tutoial	Tutoial, courses, teaching
	Productivity	Productivity, selfimproving
	Lifehacks	lifehacks, tipsandtricks, Hacks, LifeTips
	DIY	DIY, homedecor, Homemade
	History	History, legacy
	Economy	finance, Investment, FinancialAdvice, FinancialGoals, business, Economy

Figure 6. We employed a hierarchical keyword taxonomy for video collection, comprising six primary categories and 43 subcategories, yielding a total of 293 unique keywords. This taxonomy was meticulously designed to target videos with high-quality music. Keywords unrelated to music, such as *news broadcast* and *read*, were excluded to enhance the precision of the search and prioritize music-centric content. The resulting keyword set was then utilized to crawl and curate a collection of videos.

in the annotation text shown in Figure 2, we also compiled statistics on the frequency of music tags and keywords in video titles (with meaningless stop words like *I* and *the* removed). Figure 12 presents word clouds for music tags and video titles, illustrating the diversity in music genres, styles, and instruments, as well as the variety of content and themes showcased in the videos.

A.3. Benchmark

Metrics of HarmonySet-OE

Traditional language metrics are mainly sensitive to lexical variations and cannot identify changes in sentence semantics.

BLEU-4 BLEU [59], or Bilingual Evaluation Understudy, is a metric for evaluating machine translation by comparing N-grams of the translation to human references. BLEU-4 specifically evaluates the match of four-word sequences and

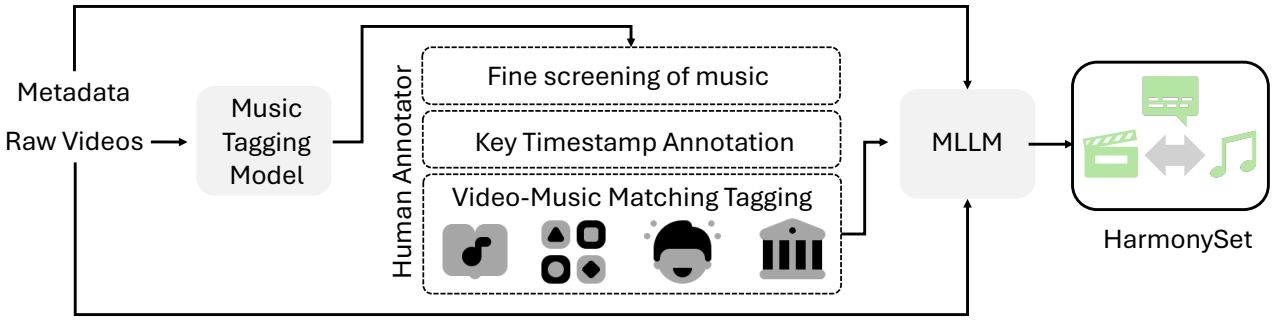

Figure 7. Multi-phase annotation process


```
{
    "title": "Do you know what kind of pets they are? #pets #animals #shorts",
    "author": "Sweetheart Zoo",
    "duration": 54,
    "video_height": 640,
    "video_width": 360,
    "audio_labels": [ "Music", "Speech", "Hip hop music", "Rapping", "Boing",
    "Singing", "Rhythm and blues", "Music for children", "Plop", "Zing" ]
}
```


Figure 8. An example of metadata of video-audio pairs. The top 10 most probable music labels generated by the PANNs are retained. Metadata used in the automated annotation generation pipeline includes video titles and music tags.

includes a Brevity Penalty to account for shorter translations.

ROUGE-L ROUGE [60], or Recall-Oriented Understudy for Gisting Evaluation, measures the overlap between generated and reference summaries. ROUGE-L focuses on the Longest Common Subsequence (LCS), which does not require consecutive word matches.

BLEURT BLEURT [69], or Bilingual Evaluation Understudy with Representations from Transformers, evaluates machine translation and natural language generation using pre-trained language models like BERT. It captures deep semantic information, addressing issues like synonym substitution and sentence rearrangement better than traditional metrics.

CIDEr CIDEr [70], or Consensus-based Image Description Evaluation, is designed for image captioning tasks. It uses TF-IDF weights to emphasize n-grams common in human annotations but rare in general descriptions, capturing more detailed information.

WUPS The Wu-Palmer similarity (WUPS) [71] measures the similarity of word senses based on their positions in the WordNet taxonomy. However, it struggles with words that are similar in form but different in meaning and cannot handle phrases or sentences effectively.

These traditional Natural Language Generation metrics lack the ability to understand and evaluate text with complex logic. In contrast, large language models are proven to be capable of comprehending text deeply. We employed GPT-4o to evaluate similarity between model responses and ground truth annotations in terms of their semantic meaning or the true intent they express. Specific prompt can be seen in Figure 13.

HarmonySet-MC Curation We followed the QA curation pipeline of widely used multiple-choice QA benchmarks like EgoSchema [10] and VideoMME [13] to design the extended benchmark HarmonySet-MC. We used a large language model to generate three incorrect answers for each annotation that is used as the correct answer. After iterating through multiple prompt versions, we ensured that the final multiple-choice questions were challenging yet reasonable, avoiding any form-based biases beyond semantics. Detailed prompt can be seen in Figure 14. For the LLM selection, we tested GPT4o, Claude, and GPT4, ultimately choosing GPT4o as the model for generating the multiple-choice questions.

B. More details of Experiment

B.1. Baseline Introduction

Our open-source model baselines include VideoLLaMA2 [21] and Video-SALMONN [22], both state-of-the-art video-audio multimodal large language models. Earlier MLLMs capable of processing both video and audio, such as Macaw-LLM [24], are no longer available due to a lack of maintenance and therefore were not included in our experiments.

Video-LLaMA2 VideoLLaMA2 is a Video Large Language Model designed for spatial-temporal modeling and audio understanding. VideoLLaMA2 comprises a vision-language branch, an audio-language branch, a Spatial-Temporal Convolution Connector (STC Connector), and a Large Language Model (LLM). The vision-language branch uses a CLIP image encoder (ViT-L/14) to process individual frames, then aggregates these features using a novel Spatial-Temporal Convolution Connector (STC Connector) designed to preserve spatial-temporal information efficiently. The audio-language branch transforms audio into spectrograms, encodes them using BEATs, and then uses an MLP to align the audio features with the LLM. The chosen LLMs for this

Rhythm & Synchronization	Music pace	fast, slow, medium, diverse
	Rhythm alignment	precise, appropriate, irrelevant, mismatch
Theme & Content	Content alignment	strongly related, indirectly related, unrelated, conflicting
	Narrative supplement	suggestion, enhancement, reversing, independent, no supplement
Emotion	Video Emotions	happy, sad, angry, fear, anxious, creative, hope, surprise, humor, relax, ···
	Emotion alignment	consistent, partially consistent, contrast, imbalanced, unrelated
	Emotional richness	rich, single, abrupt
Culture	Culture elements	clothing, customs, cuisines, landmarks, no culture element
	Cultural relevance	related, misuse, unrelated

Figure 9. A hierarchical labeling system is employed for manual annotation, encompassing four primary aspects: Rhythm and Synchronization, Theme and Content, Emotion, and Culture. These aspects are further refined into nine sub-aspects, with labels representing factual information or degrees of match.

architecture are Mistral-Instruct and Qwen2-Instruct. The STC Connector prioritizes maintaining spatial-temporal order, minimizing token count, and mitigating information loss during downsampling through the use of 3D downsampling and RegStage convolution blocks.

Video-SALMONN Video-SALMONN is designed for obtaining fine-grained temporal information required by speech understanding. Video-SALMONN uses pre-trained encoders for visual (InstructBLIP), speech (Whisper), and non-speech audio (BEATs) inputs. These features are temporally synchronized, aligning audio and visual features at the video frame rate (2Hz). A Multi-Resolution Causal (MRC) Q-Former then processes these synchronized features at different time scales (e.g., 1, 5, and 10-second windows) to capture fine-grained audio-visual joint representations.

B.2. Implementation Details

For VideoLLaMA2 training, we utilized 4 NVIDIA H800 GPUs (a total of 32 GPUs). The training configuration followed the default settings of VideoLLaMA2-AV, except for the learning rate (lr), which was adjusted to 1e-5. The model was trained for 2 epochs on our dataset. All testing was conducted on a single NVIDIA H800 GPU (a total of 8 GPUs). Experiments in Table 2 used 16 frames for testing. For experiments in Table 10, which explore the impact of varying frame numbers, the number of frames used for testing matched the number of frames used during training (16/32/64 frames).

B.3. More analysis of main results

In Table 2, we present scores for each model across six main video categories and four evaluation aspects. A horizontal comparison across the six video categories reveals that models generally score lowest on knowledge-based videos and highest on arts and performance videos. This may be because knowledge-based videos often require deeper semantic understanding and factual recall. The criteria for evaluating knowledge-based videos might be more stringent, reflecting the need for accurate information retrieval. On the other hand, the evaluation of arts and performance videos could be more subjective and open to interpretation, potentially leading to higher scores. A vertical comparison across the four evaluation aspects shows that all models consistently score lowest on the cultural aspect, suggesting that understanding and evaluating cultural nuances remains a significant challenge for current models. This could be due to the inherent complexity and subjectivity of cultural interpretations, and the current models lack sufficient training data that adequately represents the diversity and depth of cultural contexts. This deficiency hinders their ability to accurately assess culturally relevant aspects of the videos.

While Gemini-1.5 Pro generally performs well, the HarmonySet-enhanced VideoLLaMA2 demonstrates that open-source models can achieve comparable or even superior performance. This highlights the potential of open-source development in the MLLM domain. However, the base VideoLLaMA2 and video-SALMONN lag signifi-

Gemini Prompt for Automated Annotation Generation

```
Analyze the provided video and music, developing a preliminary interpretation by generating the most fitting
description for each of the following four aspects: Rhythm and Synchronization, Theme and Content, Emotion,
and Cultural Relevance.
• Rhythm and Synchronization: Analyze how well the music's tempo (fast or slow) matches the video's pace,
  the level of synchronization between the music's rhythm and the video's actions or cuts, and the overall
  rhythmic coherence between the two.
• Theme and Content: Analyze how well the music's theme complements the video's content, whether the music
  enhances the video's narrative or message, and if there is any thematic dissonance between the two.
• Emotion: Analyze how well the music's style and the video's emotional tone match, identify the specific
  emotions conveyed by both the music (e.g., joyful, melancholic, suspenseful) and the video, assess the
  level of emotional coherence between them, and consider whether the music elevates the emotional impact of
  the video.
• Cultural Relevance: Analyze whether the music and video share any common cultural elements, assess the
  relevance and accuracy of the cultural elements used, and identify any specific geographical references in
  both the music and video and how well they align.
Focus your analysis on the detailed reason of match or mismatch between the video and the music for each
aspect. Describe elements of both the video and the music and explain why you believe they match or mismatch.

Alongside this video and music, You will be provided with manually labeled tags representing the ground
truth. Use these tags to refine and enhance your initial interpretations. Specifically, integrate the
timestamps indicated in the sync_answer.value tag into your analysis of Rhythm and Synchronization,
explaining how these synchronized (or unsynchronized) moments shape your overall understanding of the video
and music's rhythmic relationship. The labels offer detailed insights, including the sync_answer.value tag
which specifies key video timestamps and their corresponding musical synchronization. Within this tag, 'A'
denotes precise synchronization between video transitions and musical changes; 'B' indicates a musical
change near the transition point, but without strong synchronization; and 'C' signifies no musical change
during the video transition. 0 means no key timestamps.

You will also be provided with the video's metadata including title and music tags. Use this metadata to
supplement and support your interpretation. Music tags are only used for reference, and latter tags in the
ten tags are generally not associated with music. It cannot be assumed that tag necessarily represents
musically related content.

Finally, articulate your comprehensive and nuanced interpretation in your own words, without explicitly
mentioning the labels themselves. Ensure your interpretation accurately and thoroughly reflects the
information conveyed by the labels.
For each aspect, provide a rich and detailed explanation, exploring the nuances and subtleties you observe,
always focusing on the interplay and degree of alignment between the visual and auditory elements. You can
include your own opinion to fulfill your answer.

The output format should be:

Rhythm and Synchronization: <your detailed explanation>
Theme and Content: <your detailed explanation>
Emotion: <your detailed explanation>
Cultural Relevance: <your detailed explanation>
```

Figure 10. Prompt for generating automated annotation using Gemini 1.5 Pro. Inputs include video and audio content, manual annotated labels, and metadata. The model is tasked with providing a detailed understanding of the match across four aspects.

cantly, indicating that further research and development are needed to close the gap with closed-source models without relying on additional datasets like HarmonySet.

B.4. Experiments on General Audio-Visual Tasks

Table 7 presents the performance of VideoLLaMA2 on AVSD dataset before and after training with HarmonySet. AVSD is a widely used dataset for audio-visual question answering. It includes general audio-visual QA tasks, such as "Do you hear any audio at all?" and "Is there a violin sound in the background of the video?" From the table, it is evident that the VideoLLaMA2 trained with HarmonySet has achieved improved performance across all metrics. This indicates that HarmonySet not only enhances the model's understanding of the relationships between

Table 7. Performance on AVSD for General Audio-Visual QA.

Metrics	BLEU	BLEU-4	ROUGE	BERT
VideoLLaMA2	0.28	0.19	0.33	0.87
VideoLLaMA2 (HarmonySet)	**0.30**	**0.21**	**0.35**	**0.89**

Diversified Instructions

```
"<video>\nExamine the given video and soundtrack, assessing their alignment in four critical areas: rhythm and
timing, thematic elements, emotional resonance, and cultural significance. Offer a detailed analysis for each
dimension, using specific timestamps to highlight how moments of synchronization (or lack thereof) enhance the
overall effectiveness of the audio-visual experience.",

"<video>\nEvaluate the supplied video and music, focusing on their compatibility in four essential aspects:
rhythm and coordination, thematic relevance, emotional effect, and cultural importance. Provide an in-depth
explanation for each category, referencing specific timestamps to demonstrate how instances of synchronization
(or their absence) influence the overall impact of the combined audio-visual presentation.",

"<video>\nAssess the provided video alongside the music, analyzing their compatibility across four main
dimensions: rhythm and alignment, thematic substance, emotional influence, and cultural context. Deliver a
thorough explanation for each aspect, incorporating specific timestamps to illustrate how moments of
synchronization (or their lack) contribute to the overall effectiveness of the audio-visual experience.",

"<video>\nAnalyze the video and accompanying music, evaluating their compatibility in four key areas: rhythm
and harmony, thematic expression, emotional depth, and cultural relevance. Provide a comprehensive breakdown
for each dimension, using specific timestamps to show how moments of synchronization (or their absence) affect
the overall effectiveness of the audio-visual experience.",

"<video>\nInvestigate the provided video and its music, assessing their compatibility across four fundamental
dimensions: rhythm and synchronization, thematic content, emotional impact, and cultural significance. Offer a
detailed analysis for each aspect, referencing specific timestamps to illustrate how moments of synchronization
(or their lack) enhance the overall effectiveness of the combined audio-visual experience.",

"<video>\nReview the given video and soundtrack, focusing on their compatibility in four critical areas: rhythm
and synchronization, thematic elements, emotional resonance, and cultural relevance. Provide a thorough
explanation for each dimension, incorporating specific timestamps to highlight how moments of synchronization
(or their absence) contribute to the overall effectiveness of the audio-visual experience.",

"<video>\nCritique the provided video and music, evaluating their compatibility across four key dimensions:
rhythm and timing, thematic content, emotional impact, and cultural significance. Deliver a comprehensive
analysis for each aspect, using specific timestamps to illustrate how moments of synchronization (or their
absence) influence the overall effectiveness of the audio-visual experience.",

"<video>\nExplore the relationship between the provided video and music, assessing their compatibility in four
main areas: rhythm and synchronization, thematic relevance, emotional effect, and cultural importance. Provide
an in-depth explanation for each dimension, referencing specific timestamps to demonstrate how moments of
synchronization (or their lack) affect the overall impact of the audio-visual presentation.",

"<video>\nAnalyze the video and its accompanying music, evaluating their compatibility across four essential
dimensions: rhythm and coordination, thematic substance, emotional influence, and cultural context. Offer a
detailed breakdown for each aspect, incorporating specific timestamps to illustrate how moments of
synchronization (or their absence) contribute to the overall effectiveness of the audio-visual experience.",

"<video>\nExamine the provided video and music, focusing on their compatibility in four critical areas: rhythm
and alignment, thematic expression, emotional depth, and cultural relevance. Provide a comprehensive
explanation for each dimension, using specific timestamps to show how moments of synchronization (or their lack)
enhance the overall effectiveness of the combined audio-visual experience.",

"<video>\nAnalyze the provided video and music, evaluating their compatibility across four key dimensions:
rhythm and synchronization, thematic content, emotional impact, and cultural relevance. Provide a comprehensive
explanation for each aspect, incorporating specific timestamps to illustrate how moments of synchronization (or
their absence) contribute to the overall effectiveness of the combined audio-visual experience."
```

Figure 11. Diversified instructions for HarmonySet data

video and music but also proves beneficial for conventional tasks, such as perception. We will further investigate the specific insights and capabilities that HarmonySet can provide, aiming to deepen our understanding of its impact on model performance and its potential applications in various audio-visual tasks.

B.5. More details of ablations

Full ablation on different training data Table 9 shows the results of VideoLLaMA2 when not trained, trained with 10k MLLM auto-generated data, trained with 10k HarmonySet data, and trained with the entire HarmonySet data. It provides a more intuitive comparison on different types of training data, demonstrating the effectiveness and importance of incorporating human knowledge through HarmonySet.

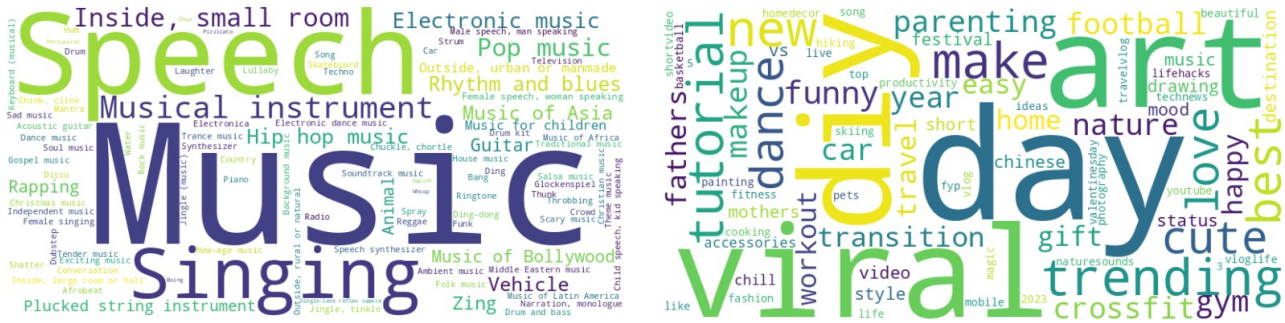

Figure 12. (Left) Word cloud visualizations of high-frequency music tags extracted by PANNs (excluding stop words). The larger the word, the higher its frequency. This showcases the diversity of musical genres, instruments, and cultures. (Right) Word cloud visualizations of video titles (excluding stop words). The word cloud demonstrates the wide range of video types and diverse scenes included in the dataset.

GPT Prompt for HarmonySet-OE Evaluation

```
Please evaluate the similarity between the following response and the correct answer in terms of their
semantic meaning or the true intent they express. Provide scores based on the following four aspects:
1. Rhythm and Synchronization
2. Thematic Content
3. Emotional Impact
4. Cultural Relevance

Provide a score for each aspect on a scale of 1 to 10, where a higher score indicates better semantic
similarity. Please output a single line containing only four values indicating the scores for four aspects,
respectively. Do not output  any words other than scores. The 4 scores are separated by a comma. Remember to
rate these 4 aspects respectively, and score for one aspect should not be affect be others. If an aspect is
not addressed or is irrelevant, give a lower score.

User Response: {response}\n\nCorrect Answer: {correct_answer}\n\n
```

Figure 13. Detailed prompt for using GPT4o as the evaluation metric for HarmonySet-OE. GPT4o will receive the correct answer and the model's output response and output the similarity of the response to the true answer in four aspects, assessing semantic and factual similarity rather than mere word matching.

Table 8. Performance of humans and models on 100 questions from HarmonySet-OE. Results show a noticeable gap between even the best model's performance and human performance, highlighting the limitations of current models in generating open-ended responses.

	R & S	T	E	C
VideoLLaMA2 (HarmonySet)	5.49	5.10	5.25	4.77
Human	**7.38**	**7.02**	**7.57**	**6.32**

Ablation on Number of Frames In Table 10 we provide the results of ablation on number of frames across all six categories. Increasing frames from 16 to 32 demonstrably improves performance, highlighting the importance of sufficient temporal context. However, the performance degradation with 64 frames reveals that more frames do not necessarily translate to better results, especially for short-form videos. This suggests potential overfitting, information redundancy, or an unfavorable cost-benefit ratio regarding computational resources. Crucially, this indicates that effectively tackling our dataset's challenges doesn't require excessive computation. A moderate frame count (32 in this instance) appears to strike an optimal balance, maximizing performance while minimizing computational burden. This underscores the possibility of creating efficient and effective solutions for short-form video analysis without resorting to computationally intensive strategies, and emphasizes the importance of optimizing frame selection based on video characteristics.

B.6. Human performance on HarmonySet-OE

We also evaluated both human and model performance on HarmonySet-OE, shown in Table 8. Due to the complexity of generating open-ended answers manually, we randomly selected 100 questions from HarmonySet-OE and

GPT Prompt for HarmonySet-MC Curation

```
Please design multiple-choice questions based on the following subtitles in four aspects: Rhythm and
Synchronization, Theme and Content, Emotion, and Cultural Relevance. For each aspect, use the given
subtitles as the correct option, and design three other options that are similar but have some errors. The
correct options should contain all information that the subtitles provide.

Note that the correct option can slightly modify the language of the subtitles but should not change the
meaning. Each option should have similar length. Especially in the rhythm section, where the correct option
may contain a timestamp, the confusingly wrong option may also contain some incorrect timestamp or incorrect
description of the correct timestamp to avoid the correct answer being too obvious.

Ensure that the incorrect options are distinguishable from the correct option but not too simple or
completely unrelated. Also, provide the correct answer's letter (A, B, C, or D) randomly positioned among
the options. The correct answer should only include the letter without any other words.

Provided subtitles:\n\n{input_subtitle}\n\n

The output format should be:
Rhythm and Synchronization:
A.<one option>
B.<one option>
C.<one option>
D.<one option>
Correct Answer: <The correct choice such as 'B'>

Theme and Content:
A.<one option>
B.<one option>
C.<one option>
D.<one option>
Correct Answer: <The correct choice such as 'A'>

Emotion:
A.<one option>
B.<one option>
C.<one option>
D.<one option>
Correct Answer: <The correct choice such as 'C'>

Cultural Relevance:
A.<one option>
B.<one option>
C.<one option>
D.<one option>
Correct Answer: <The correct choice such as 'D'>
```

Figure 14. Detailed prompt for generating HarmonySet-MC using GPT4o based on HarmonySet annotations. Each question includes one correct option derived from the dataset and three distractor options designed to be similar in structure, length, and theme, but containing identifiable factual errors.

collected answers from three different annotators per question. Human-generated answers were evaluated using the same methodology applied to the models. We compared human performance against VideoLLaMA2 (HarmonySet), the best-performing model in our main experiment. Results show a noticeable gap between even the best model's performance and human performance, highlighting the limitations of current models in generating open-ended responses. However, the performance gap between humans and models on the OE task is smaller (e.g., in the cultural aspect, the human score is 6.32, while the model score is 4.77) compared to the multiple-choice task (e.g., human accuracy on the cultural aspect is 93.81%, while the best model accuracy is only 50.40%). This smaller gap in open-ended responses might be attributed to the higher cost for humans to produce long-form text, whereas models can achieve higher scores by increasing text richness and length. This suggests that the OE task presents certain challenges even for humans.

Table 9. Full ablation on impact of different training data. Results reveal that training with 10,000 automatically generated annotations provides minimal performance improvement and even hinders performance on Theme and Emotion aspects, suggesting potential inaccuracies or misleading information in the auto-generated data. In contrast, training with HarmonySet data consistently enhances performance, with greater improvements observed with larger training sets. This demonstrates the effectiveness and importance of incorporating human knowledge through HarmonySet.

Models	Metrics	Life & Emotion	Art & Performance	Travel & Events	Sports & Outdoors	Knowledge	Tech & Fashion	Overall
VideoLLaMA2 (Vanilla)	R & S	3.89	4.80	4.56	4.01	3.39	3.54	4.15
	T	4.09	4.83	4.93	3.89	3.44	3.71	4.29
	E	4.36	5.01	5.02	4.08	3.44	3.49	4.38
	C	2.95	3.46	3.69	2.56	2.32	2.52	3.05
VideoLLaMA2 (10k, F.A.)	R & S	4.56	5.05	4.97	4.28	4.20	4.17	4.59
	T	4.20	5.01	4.85	3.49	3.36	3.41	4.16
	E	4.29	4.76	4.67	4.03	3.76	3.79	4.28
	C	3.53	3.98	3.67	2.93	3.13	3.02	3.44
VideoLLaMA2 (10k, HamrmonySet)	R & S	4.69	5.58	5.30	4.49	4.36	4.25	4.86
	T	4.66	5.02	4.98	4.40	4.39	4.29	4.70
	E	4.64	5.43	5.26	4.06	3.85	3.78	4.66
	C	3.99	4.30	4.25	2.97	3.79	3.34	3.89
VideoLLaMA2 (Full HamrmonySet)	R & S	5.43	6.35	6.03	4.94	5.33	4.83	**5.55**
	T	5.12	5.21	5.03	4.84	5.21	4.85	**5.06**
	E	5.25	6.41	5.84	4.00	4.88	4.47	**5.26**
	C	4.87	4.98	4.72	3.31	5.23	4.09	**4.62**

Table 10. Table 6 records the average results across six types of videos. This table presents the detailed results of the frame number ablation experiments. Experiments using VideoLLaMA2 trained with varying numbers of video frames (16, 32, and 64) show optimal performance with 32 frames. Performance degrades with 64 frames, indicating that using too many frames may lead to information redundancy and performance degradation.

Models	Metrics	Life & Emotion	Art & Performance	Travel & Events	Sports & Outdoors	Knowledge	Tech & Fashion	Overall
16 Frames	R & S	5.43	6.35	6.03	4.94	5.33	4.83	5.55
	T	5.12	5.21	5.03	4.84	5.21	4.85	5.06
	E	5.25	6.41	5.84	4.00	4.88	4.47	5.26
	C	4.87	4.98	4.72	3.31	5.23	4.09	4.62
32 Frames	R & S	5.47	6.37	6.07	4.98	5.39	4.87	**5.59**
	T	5.14	5.23	5.08	4.85	5.23	4.86	**5.08**
	E	5.26	6.44	5.88	4.05	4.90	4.50	**5.29**
	C	4.89	5.01	4.76	3.32	5.26	4.14	**4.65**
64 Frames	R & S	5.38	6.31	5.97	4.86	5.27	4.76	5.49
	T	5.03	5.11	4.88	4.72	5.10	4.71	4.94
	E	5.22	6.36	5.80	3.95	4.82	4.40	5.21
	C	4.76	4.89	4.68	3.22	5.10	4.02	4.53