JOINT AUDIO-VIDEO DIFFUSION TRANSFORMER WITH HIERARCHICAL SPATIO-TEMPORAL PRIOR SYNCHRONIZATION

Anonymous authorsPaper under double-blind review

000

001

002

004

006

008 009

010 011

012

013

014

015

016

017

018

019

021

025

026

027 028

029

031

033

034

037

040

041

042

043

044

046 047

048

051

052

ABSTRACT

This paper introduces JavisDiT, a novel Joint Audio-Video Diffusion Transformer designed for synchronized audio-video generation (JAVG). Based on the powerful Diffusion Transformer (DiT) architecture, JavisDiT simultaneously generates high-quality audio and video content from open-ended user prompts in a unified framework. To ensure audio-video synchronization, we introduce a fine-grained spatio-temporal alignment mechanism through a Hierarchical Spatial-Temporal Synchronized Prior (HiST-Sypo) Estimator. This module extracts both global and fine-grained spatio-temporal priors, guiding the synchronization between the visual and auditory components. Furthermore, we propose a new benchmark, **JavisBench**, which consists of 10,140 high-quality text-captioned sounding videos and focuses on synchronization evaluation in diverse and complex realworld scenarios. Further, we specifically devise a robust metric for measuring the synchrony between generated audio-video pairs in real-world content. Experimental results demonstrate that JavisDiT significantly outperforms existing methods by ensuring both high-quality generation and precise synchronization, setting a new standard for JAVG tasks. Our code and model will be made available at https://is.gd/iwWDzB.

1 Introduction

In the field of AI-generated content (AIGC), multimodal generation — covering images, videos, and audio — has gained increasing attention (Rombach et al., 2022; Liu et al., 2024b; Yang et al., 2024b; Polyak et al., 2024), where diffusion-based models (Ho et al., 2020; Liu et al., 2023b; Peebles & Xie, 2023) have demonstrated remarkable performance. While early studies mainly focused on single-modality generation, recent work has shifted toward the simultaneous generation of multiple modalities (Gadre et al., 2024; Wu et al.). Notably, synchronized audio and video generation (Ruan et al., 2023; Xing et al., 2024; Sun et al., 2024) has emerged as a crucial area of study. Audio and video are inherently interconnected in most real-world scenarios, making their joint generation highly valuable for applications, such as movie production and short video creation. Current approaches to synchronized audio-video generation can be broadly categorized into two types. The first involves asynchronous pipelines, where audio is generated first and then used to synthesize video (Yariv et al., 2024; Jeong et al., 2023; Zhang et al., 2025), or vice versa (Comunità et al., 2024; Zhang et al., 2024; Xie et al., 2024). The second type involves end-to-end Joint Audio-Video Generation (namely JAVG) (Ruan et al., 2023; Xing et al., 2024; Sun et al., 2024), which, by avoiding noise accumulation during cascading processes of the first type, has thus attracted more research attention. Overall, it is strongly and widely believed that promoting JAVG requires two equally critical criteria: (1) ensuring high-quality generation of audio and video, and (2) maintaining perfect synchronization between two modalities.

On one hand, for high-quality audio-video generation, the backbone modules of both branches need to be as strong as possible. Recent JAVG research (Ruan et al., 2023; Xing et al., 2024; Sun et al., 2024) attempts to take Diffusion Transformer (DiT) (Peebles & Xie, 2023) as backbone architectures, due to the remarkably enhanced performance for vision and even audio tasks. In particular, AV-DiT (Wang et al., 2024b) and MM-LDM (Sun et al., 2024) inherit an image DiT (Peebles & Xie, 2023; Esser et al., 2024) to generate video and audio, where the fine-grained spatio-temporal modeling capability can be however limited. Concurrent works like Uniform (Zhao et al., 2025) and SyncFlow (Liu et al., 2024a) take the enhanced STDiT3 (Zheng et al., 2024) blocks, but they either simply concatenate

video/audio tokens or utilize a mono-directional video-to-audio adapter, lacking sufficient mutual information exchange between the two channels.

On the other hand, in terms of synchronization between the audio and video channels in JAVG, extensive investigations have been made to explore more effective alignment strategies. However, most of these works still fall prey to insufficient modeling of synchronization. For example, some emphasize coarse-grained temporal alignment by implementing straightforward parameter sharing (Ruan et al., 2023; Wang et al., 2024b; Zhao et al., 2025) or temporal modulating (Wang et al., 2024b; Liu et al., 2024a) between the audio and video channels. Others focus on coarse-grained semantic alignment, such as Seeing-Hearing (Xing et al., 2024), which carries out simple audiovideo embedding alignment, and MM-LDM (Sun et al., 2024) performing straightforward representation align-

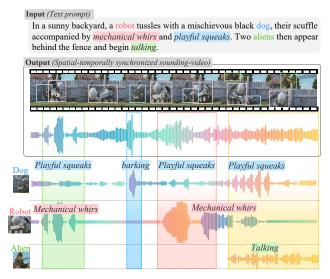


Figure 1: Given the text prompt, a JAVG system generates a spatial-temporally synchronized sounding video. The sounds align perfectly with the temporal progression of the actions.

ment in the latent space of the diffusion model. These approaches fail to consider a *fine-grained spatio-temporal alignment*, which is essential for realistic synchronized audio-video generation. Specifically, in a realistic-sounding video, all visual and auditory contents should be dictated by spatio-temporal characteristics of various objects, *i.e.*, **a**) spatially distinct visual content (*e.g.*, where events occur and how objects move), and **b**) temporally corresponding auditory attributes, such as timbre and synchronized duration. We illustrate this in Fig. 1: A dog and a robot are playing on the ground, while an alien appears and starts to talk. In this scene, the robot's mechanical whirs and the dog's squeaks and barks keep consistent, and the alien's talking emerges later and continues to the end. Overall, the visual content and audio stay synchronized both temporally and spatially.

To address the above dual challenges, this paper proposes a novel niche-targeting JAVG system, called **JavisDiT** (cf. Fig. 2). First, JavisDiT adopts the DiT architecture as the backbone, where the audio and video channels share AV-DiT blocks to enable high-quality audio-video generation. Within JavisDiT, we design three infrastructural blocks: Spatio-Temporal Self-Attention, Coarse-Grained Cross-Attention, and Fine-Grained Spatio-Temporal Self-Attention Cross-Attention. To implement the fine-grained spatio-temporal alignment idea, we design a Hierarchical Spatial-Temporal Synchronized Prior (HiST-Sypo) estimation module. The HiST-Sypo Estimator hierarchically extracts the following from the input conditional prompt: • global coarse-grained spatio-temporal priors, such as the semantic framework of the overall sounding video; • fine-grained spatio-temporal priors, such as the distinct visual content triggered by various sounds and their corresponding temporal priors.

These global and fine-grained priors serve as spatio-temporal features, injected into different AV-DiT blocks to guide the spatial semantic and temporal synchronization of both audio and video. Built on JavisDiT, we further design a contrastive learning-based (Chen et al., 2020b) HiST-Sypo estimation strategy to learn robust spatial-temporal synchronized prior knowledge from large-scale data of sounding videos.

Further, we observe that the existing JAVG evaluation benchmarks, *e.g.*, Landscape (Lee et al., 2022) and AIST++ (Li et al., 2021) suffer from certain limitations, including overly simplistic audio-video content and limited scene diversity (Mao et al., 2024). This creates an evident gap compared to the complex audio-video content encountered in real-world applications, *i.e.*, models (Ruan et al., 2023; Sun et al., 2024) trained there largely suffer from out-of-domain issue, thereby partially hindering the development and practical applicability on open environment. To bridge these gaps, we introduce a new JAVG benchmark, called **JavisBench**. After rigorous manual inspection, we collect a total of 10,140 high-quality text-captioned sounding videos, which span over 5 dimensions with 19 scene categories. Notably, more than 50% of these videos feature highly complex and challenging

scenarios, making the dataset more representative of diverse real-world applications. Meanwhile, we find that existing JAVG evaluation method can also be limited by failing to adequately assess the JAVG systems on complex sounding videos. Thus, we devise a novel metric (JavisScore) based on a temporal-aware semantic alignment mechanism.

Extensive experiments on both existing JAVG benchmarks and our JavisBench dataset demonstrate that JavisDiT significantly outperforms state-of-the-art methods in generating high-quality sounding videos. Our system is especially effective in handling complex-scene videos, thanks to the HiST-Sypo estimation mechanism. In summary, our main contributions are as follows:

- We introduce JavisDiT, a novel JAVG model with a hierarchical spatio-temporal prior estimation mechanism to achieve audio-video synchronization in both spatial and temporal dimensions.
- We contribute JavisBench, a new large-scale JAVG benchmark dataset with challenging scenarios, along with robust metrics to evaluate audio-video synchronization, offering a comprehensive baseline for future research.
- Empirically, our system achieves state-of-the-art performance across both existing closed-set and our open-world benchmarks, setting a new standard for JAVG.

2 Preliminaries

Task Definition and Formulation. Joint Audio-Video Generation (JAVG) requires diffusion models \mathcal{G}_{θ} to simultaneously generate videos \mathbf{v} and corresponding audios a given a text input s. The forward process at timestamp t, either with DDPM (Ho et al., 2020) or FlowMatching (Lipman et al., 2023), is formulated as:

$$(\mathbf{v}_{t-1}, \mathbf{a}_{t-1}) = \mathcal{G}_{\theta}(\mathbf{v}_t, \mathbf{a}_t, s, t). \tag{1}$$

 $\mathbf{v} \in \mathbb{R}^{T_v \times (H \times W) \times 3}$ denotes a video with T_v frames and 3 RGB channels, where H, W are the frame height and width. On the other hand, $\mathbf{a} \in \mathbb{R}^{T_a \times M \times 1}$ encodes audio in the image-like mel-spectrogram with T_a temporal frames and 1 channel, where M is the frequency bins.

Spatio-Temporal Alignment. For an ideal model \mathcal{G} , the generated video and audio should maintain coarse-grained semantic alignment with the input text s, while further achieving fine-grained spatio-temporal alignment between video \mathbf{v} and audio \mathbf{a} , which are formally defined as:

- Spatial Alignment means the occurrence of a specific event in a region of the video frame (at the $H \times W$ dimension in \mathbf{v}), which corresponds to the appearance of matching frequency components in the audio (at the M dimension in \mathbf{a}).
- Temporal Alignment means the onset or termination of an event at a specific frame or timestamp in the video (at the T_v dimension in \mathbf{v}), which must coincide with the corresponding start or stop of the response in the audio (at the T_a dimension in \mathbf{a}).

We focus on the joint generation of high-quality and spatio-temporally aligned video-audio pairs.

3 THE PROPOSED JAVISDIT 🎒 SYSTEM

3.1 JAVISDIT MODEL ARCHITECTURE

Fig. 2 illustrates the overall architecture. Spatio-temporal attention is employed for effective cross-modal alignment while ensuring high-quality generation for videos and audios. Each branch uses *ST-SelfAttn* for intra-modal aggregation, incorporates coarse text semantics via *CrossAttn*, integrates fine-grained spatio-temporal priors through *ST-CrossAttn*, and enhances video-audio fusion with *Bi-CrossAttn*. Here we introduce the core components displayed in Fig. 2 (b):

Spatio-Temporal Self-Attention. Given that both videos and audios possess spatial and temporal attributes (Ruan et al., 2023), we employ a cascaded spatio-temporal self-attention mechanism for intra-modal information aggregation. As illustrated Fig. 2(b), MHSA is applied sequentially along the spatial $((H \times W))$ for \mathbf{v} , M for \mathbf{a}) and temporal (T_v) for \mathbf{v} , T_a for \mathbf{a}) dimensions. This efficiently achieves fine-grained spatio-temporal modeling with reduced computational cost.

Spatio-Temporal Cross-Attention. For a text prompt s, after injecting coarse semantics from the T5 encoder (Raffel et al., 2020) via vanilla cross attention, we estimate spatial and temporal priors with N_s and N_t learnable c-dimensional tokens by our ST-Prior Estimator (see details in Sec. 3.2). As Fig. 2(b) shows, spatial and temporal priors guide cross-attention in both branches along

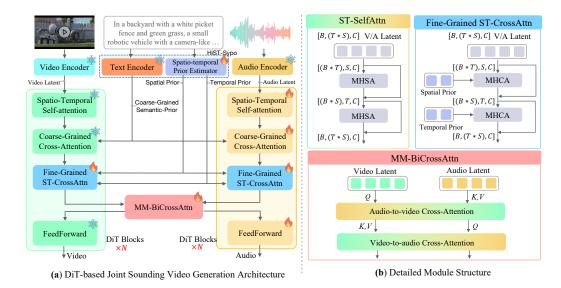


Figure 2: On the left, we present the overall DiT-based sounding video generation architecture of <code>JavisDiT</code>, consisting of a video generation branch, audio generation branch, <code>HiST-Sypo Estimator</code> module, and the <code>MM-BiCrossAttn</code> module. On the right, we illustrate the detailed structural design of Spatio-temporal Self-attention (<code>ST-SelfAttn</code>), Fine-grained Spatio-temporal Cross-attention (<code>Fine-Grained ST-CrossAttn</code>), and Multi-Modality Bidirectional Cross-attention (<code>MM-BiCrossAttn</code>).

the spatial and temporal dimensions, enabling unified, fine-grained conditioning for video-audio synchronization.

Cross-Modality Bidirectional-Attention. After aligning video and audio with ST-Prior, we incorporate a bidirectional attention module (Liu et al., 2025) to enable direct cross-modal interactions. As Fig. 2(b) illustrates, after computing the attention matrix A between q_v and k_a , we first multiply A with v_a to obtain the audio-to-video cross-attention. Similarly, multiplying A^T (the transpose of A) with v_v yields the video-to-audio cross-attention. This mechanism enhances cross-modal information aggregation to facilitate high-quality joint audio-video generation.

3.2 HIERARCHICAL SPATIAL-TEMPORAL SYNCHRONIZED PRIOR ESTIMATOR

Unlike previous works, our JavisDiT derives two hierarchical conditions (priors) from texts: a global semantic prior for coarse event suggestion (*what*) and a fine-grained spatio-temporal prior (ST-Prior) to specify event timing and location (*when* and *where*). These two complementary priors are unified into a HiST-Sypo Estimator, enabling precise synchronization between generated video and audio.

Coarse-Grained Spatio-Temporal Prior. Since the default semantic embeddings from T5 encoder (Raffel et al., 2020) are strong enough to coarsely describe the overall sounding event, we simply reuse T5 embeddings as our coarse-grained spatio-temporal prior (or, semantic prior).

Fine-Grained Spatio-Temporal Prior. Text inputs typically provide coarse-grained descrip-

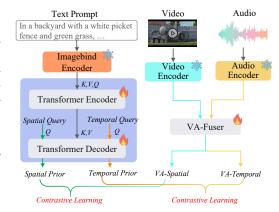


Figure 3: The framework of **Spatio-temporal Prior Estimator** with a 4-layer transformer encoder-decoder (referring to the purple region). Contrastive learning is utilized for optimization.

tions of events, *e.g.*, "a car starts its engine and leaves the screen". To enable fine-grained conditioning, we estimate spatio-temporal priors: *spatial prior* specifies where events occur (*e.g.*, "the car is in the top-left corner of screen"), and *temporal prior* defines when they start and end (*e.g.*, "sound starts at

2s, exits at 7s, fades at 9s"). Instead of generating explicit prompts from LLMs (Liu et al., 2024b), we efficiently estimate spatio-temporal priors as latent token conditions to guide the diffusion process.

Fine-Grained ST-Prior Estimation. As shown in Fig. 3, for a given text input s, we utilize the 77 hidden states from ImageBind's (Girdhar et al., 2023) text encoder, and use $N_s=32$ spatial tokens \mathbf{p}_s and $N_t=32$ temporal tokens \mathbf{p}_t to query a 4-layer transformer \mathcal{P} to extract spatio-temporal information. Since input text prompts usually may not specify an event's happening time and location, it may almost occur and cease in arbitrary location and timing, and the same text s should correspondingly yield different ST-Priors $(\mathbf{p}_s,\mathbf{p}_t)$. To capture this variability, our ST-Prior Estimator \mathcal{P} outputs the mean and variance of a Gaussian distribution, from which a plausible $(\mathbf{p}_s,\mathbf{p}_t)$ is sampled: $(\mathbf{p}_s,\mathbf{p}_t)\leftarrow \mathcal{P}_\phi(s;\epsilon)$, where ϵ is a Gaussian noise. Furthermore, we carefully devised a contrastive learning approach to learn a robust ST-Prior Estimator, which involves a series of negative sample (asynchronous video-audio pairs) construction strategies and specifically-designed loss functions. Implementation details are provided in the Sec. C.2.

3.3 MULTI-STAGE TRAINING STRATEGY

Our JavisDiT focuses on two core objectives: achieving high-quality single-modal generation and ensuring fine-grained spatio-temporal alignment between generated videos and audios. To do this, we leverage the pretrained weights of OpenSora (Zheng et al., 2024) for the video branch, and adopt a three-stage training strategy for robust joint video-audio generation:

- 1) **Audio Pretraining**. We initialize the audio branch with OpenSora's video branch weights and train it on 0.8M audio-text samples to ensure superior single-modal generation quality.
- 2) **ST-Prior Training**. We train the ST-Prior Estimator \mathcal{P}_{ϕ} using 0.6M synchronous text-video-audio triplets and synthesized asynchronous negative samples to strengthen representation.
- 3) **JAVG Training**. We freeze the video and audio branches' self-attention blocks and ST-Prior Estimator, training only the ST-CrossAttn and Bi-CrossAttn modules with 0.6M samples to enable synchronized video-audio generation.

In addition, dynamic temporal masking (Zheng et al., 2024) enables flexible adaptation of JavisDiT to x-conditional tasks (v2a/a2v generation, image animation, video-audio extension, *etc.*). Details refer to Sec. E.6.

4 A CHALLENGING JAVISBENCH BENCHMARK

A strong generative model must ensure diverse video content, audio types, and fine-grained spatio-temporal synchrony. However, current JAVG evaluation benchmarks fall short in encompassing the diversity of scenes, the complexity of real-world environments, and the presence of multiple sounding events within a single audio-video instance, thereby constraining comprehensive and realistic evaluation (detailed in Sec. D.1). To combat this, we propose a more challenging benchmark featuring complex multi-event video-audio pairs (Sec. 4.1), along with a robust synchronization-oriented metric tailored for evaluating fine-grained spatio-temporal alignment (Sec. 4.2).

4.1 DATA CONSTRUCTION

Taxonomy. We design five evaluation dimensions from coarse to fine:

- Event Scenario depicts scenarios where audio-visual events happen, such as nature or industry.
- Video Style describes the visual style of given videos, such as camera shooting or 2D animations.
- Sound Type: the sounding type of given audios, such as sound effects or music.
- Spatial Composition defines the sounding subjects appearing in videos and audios, divided by single or multiple sounding subjects that exist.
- Temporal Composition describes events' onsets and terminations in v-a pairs, distinguished by single or multiple sounding sources that sequentially or concurrently occur.

Using GPT-4 (Achiam et al., 2023), we develop a hierarchical categorization system with 5 dimensions and 19 categories (Fig. 4), and the detailed definitions are presented in Sec. D.2.

Data Curation. There are two data sources: (1) test sets of existing datasets (*e.g.*, Landscape/AIST++ and FAVDBench (Shen et al., 2023)), and (2) YouTube videos uploaded between June and December 2024 to prevent data leakage (Mao et al., 2024). To collect YouTube videos, we prompt GPT-4 to generate category-specific keywords using our defined taxonomy, enabling efficient and targeted

Benchmark	#Sample	Scenarios	Categories	Event Scenario Vis	ual Style Sound T
AIST++	20	1 (dancing)	1 (music)	28% 4% Living 20% 13%	3D Biological 32%
Landscape	100	1 (natural)	1 (ambient)	Natural 35% Camera Urban	2D Mechanical
TAVGBench	3,000	Unknown	Indiscriminate	Spatial Composition Single	Temporal Composition Single
Ours	10,140	Diversified	Hierarchical	75% Offscree	n Sequential 15% Simulta

Figure 4: **Left**: Comparison between JavisBench and other benchmarks. TAVGBench (Mao et al., 2024)'s evaluation set is not currently released (Unknown). **Right**: Detailed category distribution of JavisBench.

video collection while avoiding noisy data. After strict manual legal and ethical verification (see Sec. A.1), the above process yielded around 30K sounding video candidates. With several filtering tools to ensure quality and diversity, we use the advanced Qwen-family models (Yang et al., 2024a; Wang et al., 2024c; Chu et al., 2024) to generate captions and categorize video-audio pairs into desired taxonomy, and finally curate 10,140 diverse and high-quality video-audio pairs for JavisBench-10K. More details about data construction are presented in Sec. D.3. We also randomly select 1,000 samples to form the version of JavisBench-mini for efficient evaluation.

Benchmark Statistics. Fig. 4 highlights JavisBench's contributions: (1) offering more diverse data compared to AIST++ (Li et al., 2021) and Landscape (Lee et al., 2022), (2) providing a detailed taxonomy for comprehensive evaluations, surpassing TAVGBench (Mao et al., 2024), and (3) featuring the first evaluation benchmark focused on multi-event synchronization. As Fig. 4 shows, JavisBench covers diverse *event scenarios*, *visual styles*, and *audio types*, ensuring a balanced distribution for greater **diversity**. Most realistic but under-represented scenarios like 2D/3D animations (25%) and industrial events (13%) are included. Furthermore, JavisBench emphasizes spatial and temporal complexity. 75% of samples feature *multiple sounding events*, 28% involve *sequential events*, and 57% include *simultaneous events*. These diverse multi-event scenarios pose significant **challenges** for joint audio-video generative models, particularly in achieving accurate spatial and temporal alignment.

4.2 JAVISSCORE: A MORE ROBUST JAVG METRIC

Motivation. AV-Align (Yariv et al., 2024) is a widely-adopted JAVG metric, which uses video optical-flow estimation to match the onset detected in audios to measure temporal synchronization. However, AV-Align may struggle with complex scenarios (*i.e.*, with multiple sounding events or subtle visual movements) and produce misleading results (see Sec. D.4). Therefore, we propose a more robust evaluation metric, namely JavisScore, to measure spatiotemporal synchronization in diverse, real-world contexts.

Implementation and Verification. Technically, we chunk each video-audio pair into several clips with 2-second window size and 1.5-second overlap, use ImageBind (Girdhar et al., 2023) to compute the audio-visual synchronization for each segment, and average the scores as the final metric:

$$S_{Javis} = \frac{1}{W} \sum_{i=1}^{W} \sigma\left(E_v(V_i), E_a(A_i)\right), \qquad (2)$$

where E_v and E_a are vision and audio encoders, A_i and V_i are audio and video segments in i-th window, W is the number of windows, and σ is a specifically designed synchronization measure for video-audio segments (see Sec. D.4). Motivated by Mao et al. (2024), we calculate similarities between all frames and the audio within each segment, and select the 40% least synchronized frames to obtain the score of the current window. Sec. D.4 also presents a human-annotated evaluation dataset with 3,000 samples to verify JavisScore's efficacy against previous metrics.

5 EXPERIMENTS

5.1 SETUP

Evaluation Datasets and Metrics. The proposed JavisBench is used as the primary benchmark for comprehensive evaluation. Multiple metrics are reported: (1) audio/video quality, (2) semantic

Table 1: **Main results on proposed JavisBench.** Models are evaluated with 4-second videos at 240P/24fps and audio at 16kHz. JavisDiT comprehensively outperforms or gets on par with currently available SoTA approaches. **Best** and secondary results are marked **bold** and <u>underline</u> respectively.

Method	A	V-Quali	ty	Text-Consistency			AV-Consistency			AV-Synchrony	
17101104	FVD↓	KVD↓	FAD↓	TV-IB↑	TA-IB↑	CLIP↑	CLAP↑	AV-IB↑	CAVP↑	AVHScore ↑	JavisScore ↑
- T2A+A2V											
TempoTkn (Yariv et al., 2024)	539.8	7.2	-	0.085	-	0.198	-	0.137	0.787	0.122	0.103
TPoS (Jeong et al., 2023)	839.7	4.7	-	0.103	-	0.238	-	0.142	0.778	0.129	0.095
- T2V+V2A											
ReWaS (Jeong et al., 2024)	-	-	9.4	-	0.109	-	0.236	0.111	0.794	0.104	0.079
See&Hear (Xing et al., 2024)	-	-	7.6	-	0.072	-	0.248	0.164	0.798	0.143	0.112
FoleyCftr (Zhang et al., 2024)	-	-	9.1	-	0.173	-	0.299	0.204	0.800	0.186	0.151
- T2AV											
MM-Diff (Ruan et al., 2023)	2311.9	12.2	27.5	0.080	0.032	0.173	0.048	0.119	0.783	0.109	0.070
UniVerse-1 (Wang et al., 2025)	194.2	0.5	8.7	0.272	0.111	0.309	0.245	0.104	0.793	0.098	0.077
JavisDiT(Ours)	<u>203.2</u>	<u>1.4</u>	6.9	0.151	0.197	0.325	0.320	0.201	0.801	0.183	0.158

Table 2: **Experimental results on previous datasets.** Numbers are borrowed from their released papers, and our model still performs best.

Method	L	andscap	e	AIST++			
1,1011011	FVD↓	KVD↓	FAD↓	FVD↓	KVD↓	FAD↓	
MM-Diff	332.1	26.6	9.9	219.6	49.1	12.3	
See&Hear	326.2	9.2	12.7	-	-	-	
AV-DiT	172.7	15.4	11.2	68.8	21.0	10.2	
MM-LDM	105.0	8.3	9.1	105.0	27.9	10.2	
$\pmb{JavisDiT}(Ours)$	94.2	7.8	8.5	86.7	19.8	9.6	

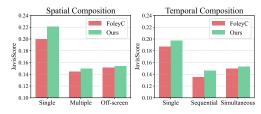


Figure 5: **Generative av-synchrony with Javis-Bench's taxonomy.** Current SOTA models still suffer from challenging scenarios.

consistency against conditional texts, (3) semantic consistency and (4) spatio-temporal synchrony between videos and audios. We also evaluate on AIST++ (Li et al., 2021) and landscape (Lee et al., 2022) for consistency, and follow prior works (Ruan et al., 2023; Sun et al., 2024) to report FVD, KVD, and FAD.

Compared Methods. For JavisBench, we reproduce and compare a series of baselines: (1) for cascaded T2A+A2V methods (Yariv et al., 2024; Jeong et al., 2023), we use AudioLDM2 (Liu et al., 2024b) for the prepositive T2A task; (2) for cascaded T2V+V2A methods (Jeong et al., 2024; Xing et al., 2024; Zhang et al., 2024), OpenSora (Zheng et al., 2024) is adopted to generate videos at first; and (3) for JAVG models, we currently compare with MM-Diffusion (Ruan et al., 2023) and UniVerse-1 Wang et al. (2025) since the others (Liu et al., 2024a; Zhao et al., 2025) are not currently open-sourced. For AIST++ and Landscape, we directly adopt the reported results.

Implementation Details. We collect 780K audio-text pairs for audio pertaining by 13 epochs, with 610K video-audio-text triplets to train ST-Prior Estimator for 1 epoch and JavisDiT for 2 epochs. The learning rate is 1e-5 for ST-Prior Estimator and 1e-4 for DiT. The video encoder-decoder and audio encoder-decoder are taken from OpenSora (Zheng et al., 2024) and AudioLDM2 (Liu et al., 2024b) and kept frozen. Detailed configurations refer to Sec. C.1.

5.2 Main Results and Observations

Our JavisDiT achieves superior single-Modal quality and video-audio synchrony. As shown in Tab. 1, our carefully designed spatiotemporal DiT architecture demonstrates exceptional unimodal generation quality, achieving significantly superior results compared to the UNet-based architecture (e.g., TempoToken (Yariv et al., 2024)) and the naive DiT architecture (e.g., MM-Diffusion (Ruan et al., 2023)), with remarkable FVD (203.2) and FAD (6.9) scores. Meanwhile, from the perspective of global semantic alignment, including text-consistency and video-audio consistency, our model also achieves state-of-the-art performance, as evidenced by its TA-IB score of 0.151 and CLIP similarity score of 0.325. We attribute the gap between our JavisDiT and UniVerse-1 (Wang et al., 2025) in video-text alignment to the gap of pretrained models (we use OpenSora while UniVerse-1 takes more powerful Wan2.1 (Wan et al., 2025)). Fig. 6 showcases some representative JAVG examples. Notably, in terms of audio-video synchrony, our end-to-end model outperforms various cascaded and joint

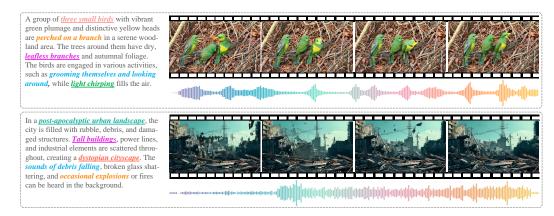


Figure 6: JavisDiT precisely captures the visual and auditory clues from text inputs to generate faithful sounding-videos with high-quality spatio-temporal alignments. Colored texts are spatio-temporal objects (underlined) and actions. More cases are shown in Sec. E.5.

Table 3: **Ablation on the model design.** The specifically developed DiT architecture and the HiST-Sypo estimator jointly contribute to single-modal quality and va-synchrony.

STDiT	HiST-Sypo	BiCA	Quality↓	Consist↑	Sync↑
×	×	×	9.371	1.140	0.118
\checkmark	×	×	7.293	1.155	0.130
V	\checkmark	×	6.127	1.191	0.150
V	×	\checkmark	6.581	1.157	0.133
	\checkmark	\checkmark	6.012	1.201	0.153

Table 4: **Ablation on token number and injection strategies of ST-Priors.** We provide more experimental results in Fig. A8 and Tab. A7.

N_s	N_t	Injection	Quality↓	Consist↑	Sync↑
0	0	-	6.581	1.157	0.133
1	1	CrossAttn	6.909	1.188	0.137
16	16	CrossAttn	6.322	1.200	0.151
32	32	CrossAttn	6.012	1.201	0.153
32	32	Addition	6.267	1.183	0.144
32	32	Modulate	6.190	1.191	0.145

audio-video generation approaches, achieving a JavisScore of **0.158**, surpassing the state-of-the-art cascaded method FoleyCrafter (Zhang et al., 2024). Tab. A8 additionally reports cascaded models' performance based on video/audio generated by our JavisDiT, where the joint JavisDiT shows consistent superiority against pipeline methods.

To ensure a rigorous comparison, we follow the standard settings from prior works (Ruan et al., 2023; Xing et al., 2024; Sun et al., 2024) and train our model for 300 epochs on two closed-set datasets, including Landscape (Lee et al., 2022) and AIST++ (Li et al., 2021). As demonstrated in Tab. 2, our method consistently achieves state-of-the-art performance, with FVD of **94.2** on Landscape and FAD of **9.6** on AIST++. These results further highlight the superiority of our meticulously designed DiT architecture and hierarchical spatial-temporal prior estimator.

Current models fail to simulate complex scenarios. Fig. 5 presents FoleyCrafter(Zhang et al., 2024) and our JavisDiT across JavisBench categories, showing that existing models — including ours — struggle with AV synchrony in complex scenarios. When the sounding video contains only a single sounding object (*e.g.*, a person playing the violin alone), JavisScore is generally higher than in multi-object cases (*e.g.*, a street performance with multiple instruments), as the latter requires identifying the correct visual-audio correspondence. Similarly, videos with multiple simultaneous events (*e.g.*, a dog barking while a car horn sounds) yield lower JavisScore than single-event cases (*e.g.*, a person clapping), due to the increased challenge of modeling event timing and interactions. Sec. D.5 and Fig. A7 thoroughly highlight the limitations in handling real-world complexity.

5.3 In-depth Analyses and Discussions

To efficiently evaluate our proposed method, we perform the 3rd stage (JAVG) training using a subset of 60K entries from our entire training data, and test the models on JavisBench-mini (with 1,000 randomly selected samples from JavisBench). For simplicity, we report three normalized scores for a clear comparison from three dimensions: (1) **Quality**: $S_{\text{AVQ}} = 0.01 \times S_{\text{FVD}} + S_{\text{KFD}} + 0.1 \times S_{\text{FAD}}$; (2) **Consistency**: $S_{\text{AVC}} = S_{\text{AV-IB}} + S_{\text{CAVP}} + S_{\text{AVHScore}}$; (3) **Synchrony**: $S_{\text{AVS}} = S_{\text{JavisScore}}$.

Well-designed DiT backbone is superior. In Tab. 3, we first build a vanilla baseline to replace the spatio-temporal transformer backbone by UNet (Liu et al., 2024b) for the audio branch, and the generation receives a pool AV-Quality at 9.371. After replacing the UNet backbone with STDiT modules (Zheng et al., 2024), $S_{\rm AVQ}$ immediately improves to 7.293, demonstrating the efficacy of the STDiT architecture. Then, we respectively incorporate the HiST-Sypo estimator and the bidirectional cross-attention (BiCA) modules for video-audio information sharing. Accordingly, HiST-Sypo brings significant enhancement on AV-Consistency (1.191 vs. 1.155) and AV-Synchrony (0.150 vs. 0.130), more than the simple BiCA module (1.157 vs. 1.155 for $S_{\rm AVC}$ and 0.133 vs. 0.130 for $S_{\rm AVS}$). This verifies our motivation that a simple channel-sharing mechanism in AV-DiT (Wang et al., 2024b) cannot effectively build audio-video synchronization, which can be achieved by our proposed fine-grained ST-Prior guidance instead. After bridging STDiT, HiST-Sypo, and BiCA modules, our JavisDiT reaches the best performance for both single-modal quality ($S_{\rm AVQ} = 6.012$) and video-audio synchronization ($S_{\rm AVC} = 1.201$ and $S_{\rm AVS} = 0.153$), demonstrating the superiority of the well-designed DiT backbone.

Spatio-temporal prior is effective and generic. In Tab. 4, we take a preliminary step to investigate the number of spatio-temporal priors and the way to inject ST-Priors for better video-audio synchronization. We first take the model without ST-CA modules (the 4th row in Tab. 3) as our baseline (the 1st row in Tab. 4), and add ST-CA modules by gradually increasing the prior number from 1 to 32. According to the results, both single-modal quality ($S_{\rm AVQ}$) and video-audio synchrony ($S_{\rm AVC}$ and $S_{\rm AVS}$) consistently increase as the prior number gains. Then, we try to utilize the 32 spatial-temporal priors for addition (adding to video/audio latent representations like a conditional embedding) and modulation (mapping priors to scales and biases to modulate video/audio representations). Although the final performance is inferior to the utilization of cross-attention, ST-priors still considerably enhance all metrics (*e.g.*, 0.144/0.145 *vs.* 0.133 for $S_{\rm AVS}$). Besides, we further experiment and evaluate the influence of prior numbers and dimensions in Fig. A8, and verify the optimization objectives during ST-priors' estimation in Tab. A7. All the empirical results verify the effectiveness and versatility of our estimated spatio-temporal priors.

How does ST-Prior ensure synchronized video&audio? Fig. 7 illustrates the synchronization mechanism for our ST-priors to guide the video-audio generation process. In particular, we visualize the spatial-temporal cross-attention map from the last block in JavisDiT at the last sampling step on both video and audio branches. The qualitative results in Fig. 7 show the spatial priors successfully help JavisDiT focus on the subject that would produce the sound (in this case, it is the bubbles rather than the diver that can make the sound), and the temporal priors bring nearly uniform attention scores along the timeline (as the

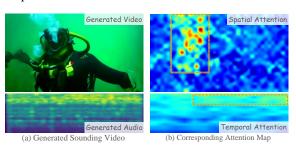


Figure 7: **Visualization of cross-attention maps by spatio-temporal priors.** Spatial priors successfully capture the sounding subjects (bubbles in this case), and temporal priors accurately cover the whole timeline for the continuous sounding event.

bubbling sound continues from beginning to end). The cross-attention mechanism for ST-priors is well-learned for synchronized video-audio generation.

6 CONCLUSION

This paper presents JavisDiT, a novel Joint Audio-Video Diffusion Transformer that simultaneously generates high-quality audio and video content with precise synchronization. We introduce the HiST-Sypo Estimator, a fine-grained spatio-temporal alignment module that extracts global and fine-grained priors to guide the synchronization between audio and video. We also propose the JavisBench dataset, comprising 10,140 high-quality text-captioned sounding videos with diverse scenes and real-world complexity, addressing limitations in current benchmarks. In addition, we introduce a temporal-aware semantic alignment mechanism to better evaluate JAVG systems on complex content. Experimental results show that JavisDiT outperforms existing approaches in both content generation and synchronization, establishing a new benchmark for JAVG tasks. Potential limitation and future work are discussed in Sec. A.5.

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.
- Honglie Chen, Weidi Xie, Andrea Vedaldi, and Andrew Zisserman. Vggsound: A large-scale audiovisual dataset. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 721–725. IEEE, 2020a.
- Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pp. 1597–1607. PMLR, 2020b.
- 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.
- 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.
- Marco Comunità, Riccardo F Gramaccioni, Emilian Postolache, Emanuele Rodolà, Danilo Comminiello, and Joshua D Reiss. Syncfusion: Multimodal onset-synchronized video-to-audio foley synthesis. In *ICASSP* 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), pp. 936–940. IEEE, 2024.
- Konstantinos Drossos, Samuel Lipping, and Tuomas Virtanen. Clotho: An audio captioning dataset. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 736–740. IEEE, 2020.
- Patrick Esser, Sumith Kulal, Andreas Blattmann, Rahim Entezari, Jonas Müller, Harry Saini, Yam Levi, Dominik Lorenz, Axel Sauer, Frederic Boesel, et al. Scaling rectified flow transformers for high-resolution image synthesis. In *Forty-first international conference on machine learning*, 2024.
- Samir Yitzhak Gadre, Gabriel Ilharco, Alex Fang, Jonathan Hayase, Georgios Smyrnis, Thao Nguyen, Ryan Marten, Mitchell Wortsman, Dhruba Ghosh, Jieyu Zhang, et al. Datacomp: In search of the next generation of multimodal datasets. *Advances in Neural Information Processing Systems*, 36, 2024.
- 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). IEEE, 2017a.
- 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), pp. 776–780. IEEE, 2017b.
- Rohit Girdhar, Alaaeldin El-Nouby, Zhuang Liu, Mannat Singh, Kalyan Vasudev Alwala, Armand Joulin, and Ishan Misra. Imagebind: One embedding space to bind them all. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 15180–15190, 2023.
- Akio Hayakawa, Masato Ishii, Takashi Shibuya, and Yuki Mitsufuji. Discriminator-guided cooperative diffusion for joint audio and video generation. *arXiv preprint arXiv:2405.17842*, 2024.
- Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020.
- Jonathan Ho, William Chan, Chitwan Saharia, Jay Whang, Ruiqi Gao, Alexey Gritsenko, Diederik P Kingma, Ben Poole, Mohammad Norouzi, David J Fleet, et al. Imagen video: High definition video generation with diffusion models. *arXiv preprint arXiv:2210.02303*, 2022.

- Yuchen Hu, Yu Gu, Chenxing Li, Rilin Chen, and Dong Yu. Video-to-audio generation with fine-grained temporal semantics. *arXiv preprint arXiv:2409.14709*, 2024.
 - Rongjie Huang, Jiawei Huang, Dongchao Yang, Yi Ren, Luping Liu, Mingze Li, Zhenhui Ye, Jinglin Liu, Xiang Yin, and Zhou Zhao. Make-an-audio: Text-to-audio generation with prompt-enhanced diffusion models. In *International Conference on Machine Learning*, pp. 13916–13932. PMLR, 2023.
 - Vladimir Iashin, Weidi Xie, Esa Rahtu, and Andrew Zisserman. Synchformer: Efficient synchronization from sparse cues. In *ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 5325–5329. IEEE, 2024.
 - Masato Ishii, Akio Hayakawa, Takashi Shibuya, and Yuki Mitsufuji. A simple but strong baseline for sounding video generation: Effective adaptation of audio and video diffusion models for joint generation. *arXiv* preprint arXiv:2409.17550, 2024.
 - Yujin Jeong, Wonjeong Ryoo, Seunghyun Lee, Dabin Seo, Wonmin Byeon, Sangpil Kim, and Jinkyu Kim. The power of sound (tpos): Audio reactive video generation with stable diffusion. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 7822–7832, 2023.
 - Yujin Jeong, Yunji Kim, Sanghyuk Chun, and Jiyoung Lee. Read, watch and scream! sound generation from text and video. *arXiv preprint arXiv:2407.05551*, 2024.
 - Chris Dongjoo Kim, Byeongchang Kim, Hyunmin Lee, and Gunhee Kim. Audiocaps: Generating captions for audios in the wild. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pp. 119–132, 2019.
 - Seung Hyun Lee, Gyeongrok Oh, Wonmin Byeon, Chanyoung Kim, Won Jeong Ryoo, Sang Ho Yoon, Hyunjun Cho, Jihyun Bae, Jinkyu Kim, and Sangpil Kim. Sound-guided semantic video generation. In *European Conference on Computer Vision*, pp. 34–50. Springer, 2022.
 - 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*, pp. 13401–13412, 2021.
 - Minghui Liao, Zhisheng Zou, Zhaoyi Wan, Cong Yao, and Xiang Bai. Real-time scene text detection with differentiable binarization and adaptive scale fusion. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2022.
 - Yaron Lipman, Ricky TQ Chen, Heli Ben-Hamu, Maximilian Nickel, and Matthew Le. Flow matching for generative modeling. In *The Eleventh International Conference on Learning Representations (ICLR)*, 2023.
 - Haohe Liu, Zehua Chen, Yi Yuan, Xinhao Mei, Xubo Liu, Danilo Mandic, Wenwu Wang, and Mark D Plumbley. Audioldm: Text-to-audio generation with latent diffusion models. *arXiv* preprint arXiv:2301.12503, 2023a.
 - Haohe Liu, Gael Le Lan, Xinhao Mei, Zhaoheng Ni, Anurag Kumar, Varun Nagaraja, Wenwu Wang, Mark D Plumbley, Yangyang Shi, and Vikas Chandra. Syncflow: Toward temporally aligned joint audio-video generation from text. *arXiv preprint arXiv:2412.15220*, 2024a.
 - Haohe Liu, Yi Yuan, Xubo Liu, Xinhao Mei, Qiuqiang Kong, Qiao Tian, Yuping Wang, Wenwu Wang, Yuxuan Wang, and Mark D Plumbley. Audioldm 2: Learning holistic audio generation with self-supervised pretraining. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 2024b.
 - Shilong Liu, Zhaoyang Zeng, Tianhe Ren, Feng Li, Hao Zhang, Jie Yang, Qing Jiang, Chunyuan Li, Jianwei Yang, Hang Su, et al. Grounding dino: Marrying dino with grounded pre-training for open-set object detection. In *European Conference on Computer Vision*, pp. 38–55. Springer, 2025.

- Xingchao Liu, Chengyue Gong, and Qiang Liu. Flow straight and fast: Learning to generate and transfer data with rectified flow. In *The Eleventh International Conference on Learning Representations (ICLR)*, 2023b.
- Xubo Liu, Qiuqiang Kong, Yan Zhao, Haohe Liu, Yi Yuan, Yuzhuo Liu, Rui Xia, Yuxuan Wang, Mark D Plumbley, and Wenwu Wang. Separate anything you describe. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 2024c.
- Yuxin Mao, Xuyang Shen, Jing Zhang, Zhen Qin, Jinxing Zhou, Mochu Xiang, Yiran Zhong, and Yuchao Dai. Tavgbench: Benchmarking text to audible-video generation. In *Proceedings of the 32nd ACM International Conference on Multimedia*, pp. 6607–6616, 2024.
- Irene Martín-Morató and Annamaria Mesaros. What is the ground truth? reliability of multi-annotator data for audio tagging. In 2021 29th European Signal Processing Conference (EUSIPCO), pp. 76–80. IEEE, 2021.
- Xinhao Mei, Chutong Meng, Haohe Liu, Qiuqiang Kong, Tom Ko, Chengqi Zhao, Mark D Plumbley, Yuexian Zou, and Wenwu Wang. Wavcaps: A chatgpt-assisted weakly-labelled audio captioning dataset for audio-language multimodal research. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 2024.
- William Peebles and Saining Xie. Scalable diffusion models with transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 4195–4205, 2023.
- Karol J Piczak. Esc: Dataset for environmental sound classification. In *Proceedings of the 23rd ACM international conference on Multimedia*, pp. 1015–1018, 2015.
- Adam Polyak, Amit Zohar, Andrew Brown, Andros Tjandra, Animesh Sinha, Ann Lee, Apoorv Vyas, Bowen Shi, Chih-Yao Ma, Ching-Yao Chuang, et al. Movie gen: A cast of media foundation models. *arXiv preprint arXiv:2410.13720*, 2024.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of machine learning research*, 21(140):1–67, 2020.
- Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*, 1(2):3, 2022.
- Nikhila Ravi, Valentin Gabeur, Yuan-Ting Hu, Ronghang Hu, Chaitanya Ryali, Tengyu Ma, Haitham Khedr, Roman Rädle, Chloe Rolland, Laura Gustafson, et al. Sam 2: Segment anything in images and videos. *arXiv preprint arXiv:2408.00714*, 2024.
- Yong Ren, Chenxing Li, Manjie Xu, Wei Liang, Yu Gu, Rilin Chen, and Dong Yu. Sta-v2a: Videoto-audio generation with semantic and temporal alignment. *arXiv preprint arXiv:2409.08601*, 2024.
- Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 10684–10695, 2022.
- Ludan Ruan, Yiyang Ma, Huan Yang, Huiguo He, Bei Liu, Jianlong Fu, Nicholas Jing Yuan, Qin Jin, and Baining Guo. Mm-diffusion: Learning multi-modal diffusion models for joint audio and video generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 10219–10228, 2023.
- Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily L Denton, Kamyar Ghasemipour, Raphael Gontijo Lopes, Burcu Karagol Ayan, Tim Salimans, et al. Photorealistic text-to-image diffusion models with deep language understanding. *Advances in neural information processing systems*, 35:36479–36494, 2022.
- Justin Salamon, Christopher Jacoby, and Juan Pablo Bello. A dataset and taxonomy for urban sound research. In *Proceedings of the 22nd ACM international conference on Multimedia*, pp. 1041–1044, 2014.

- Christoph Schuhmann, Richard Vencu, Romain Beaumont, Robert Kaczmarczyk, Clayton Mullis, Aarush Katta, Theo Coombes, Jenia Jitsev, and Aran Komatsuzaki. Laion-400m: Open dataset of clip-filtered 400 million image-text pairs. *arXiv preprint arXiv:2111.02114*, 2021.
 - Xuyang Shen, Dong Li, Jinxing Zhou, Zhen Qin, Bowen He, Xiaodong Han, Aixuan Li, Yuchao Dai, Lingpeng Kong, Meng Wang, et al. Fine-grained audible video description. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 10585–10596, 2023.
 - Bob L Sturm. The gtzan dataset: Its contents, its faults, their effects on evaluation, and its future use. arXiv preprint arXiv:1306.1461, 2013.
 - Mingzhen Sun, Weining Wang, Yanyuan Qiao, Jiahui Sun, Zihan Qin, Longteng Guo, Xinxin Zhu, and Jing Liu. Mm-ldm: Multi-modal latent diffusion model for sounding video generation. In *Proceedings of the 32nd ACM International Conference on Multimedia*, pp. 10853–10861, 2024.
 - Yi Tay, Mostafa Dehghani, Vinh Q Tran, Xavier Garcia, Jason Wei, Xuezhi Wang, Hyung Won Chung, Siamak Shakeri, Dara Bahri, Tal Schuster, et al. Ul2: Unifying language learning paradigms. *arXiv* preprint arXiv:2205.05131, 2022.
 - Team Wan, Ang Wang, Baole Ai, Bin Wen, Chaojie Mao, Chen-Wei Xie, Di Chen, Feiwu Yu, Haiming Zhao, Jianxiao Yang, et al. Wan: Open and advanced large-scale video generative models. *arXiv preprint arXiv:2503.20314*, 2025.
 - Ao Wang, Hui Chen, Zijia Lin, Jungong Han, and Guiguang Ding. Repvit: Revisiting mobile cnn from vit perspective. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 15909–15920, 2024a.
 - Duomin Wang, Wei Zuo, Aojie Li, Ling-Hao Chen, Xinyao Liao, Deyu Zhou, Zixin Yin, Xili Dai, Daxin Jiang, and Gang Yu. Universe-1: Unified audio-video generation via stitching of experts. *arXiv preprint arXiv:2509.06155*, 2025.
 - Kai Wang, Shijian Deng, Jing Shi, Dimitrios Hatzinakos, and Yapeng Tian. Av-dit: Efficient audiovisual diffusion transformer for joint audio and video generation. *arXiv preprint arXiv:2406.07686*, 2024b.
 - Peng Wang, Shuai Bai, Sinan Tan, Shijie Wang, Zhihao Fan, Jinze Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, et al. Qwen2-vl: Enhancing vision-language model's perception of the world at any resolution. *arXiv preprint arXiv:2409.12191*, 2024c.
 - Yaohui Wang, Xinyuan Chen, Xin Ma, Shangchen Zhou, Ziqi Huang, Yi Wang, Ceyuan Yang, Yinan He, Jiashuo Yu, Peiqing Yang, et al. Lavie: High-quality video generation with cascaded latent diffusion models. *International Journal of Computer Vision*, pp. 1–20, 2024d.
 - 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*.
 - Zhifeng Xie, Shengye Yu, Qile He, and Mengtian Li. Sonicvisionlm: Playing sound with vision language models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 26866–26875, 2024.
 - Yazhou Xing, Yingqing He, Zeyue Tian, Xintao Wang, and Qifeng Chen. Seeing and hearing: Opendomain visual-audio generation with diffusion latent aligners. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 7151–7161, 2024.
 - Haofei Xu, Jing Zhang, Jianfei Cai, Hamid Rezatofighi, Fisher Yu, Dacheng Tao, and Andreas Geiger. Unifying flow, stereo and depth estimation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2023.
 - 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*, 2024a.

- Zhuoyi Yang, Jiayan Teng, Wendi Zheng, Ming Ding, Shiyu Huang, Jiazheng Xu, Yuanming Yang, Wenyi Hong, Xiaohan Zhang, Guanyu Feng, et al. Cogvideox: Text-to-video diffusion models with an expert transformer. *arXiv preprint arXiv:2408.06072*, 2024b.
 - Guy Yariv, Itai Gat, Sagie Benaim, Lior Wolf, Idan Schwartz, and Yossi Adi. Diverse and aligned audio-to-video generation via text-to-video model adaptation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pp. 6639–6647, 2024.
 - Lijun Yu, José Lezama, Nitesh B Gundavarapu, Luca Versari, Kihyuk Sohn, David Minnen, Yong Cheng, Vighnesh Birodkar, Agrim Gupta, Xiuye Gu, et al. Language model beats diffusion—tokenizer is key to visual generation. *arXiv preprint arXiv:2310.05737*, 2023.
- Lin Zhang, Shentong Mo, Yijing Zhang, and Pedro Morgado. Audio-synchronized visual animation. In *European Conference on Computer Vision*, pp. 1–18. Springer, 2025.
- Yiming Zhang, Yicheng Gu, Yanhong Zeng, Zhening Xing, Yuancheng Wang, Zhizheng Wu, and Kai Chen. Foleycrafter: Bring silent videos to life with lifelike and synchronized sounds. *arXiv* preprint arXiv:2407.01494, 2024.
- Lei Zhao, Linfeng Feng, Dongxu Ge, Fangqiu Yi, Chi Zhang, Xiao-Lei Zhang, and Xuelong Li. Uniform: A unified diffusion transformer for audio-video generation. arXiv preprint arXiv:2502.03897, 2025.
- Zangwei Zheng, Xiangyu Peng, Tianji Yang, Chenhui Shen, Shenggui Li, Hongxin Liu, Yukun Zhou, Tianyi Li, and Yang You. Open-sora: Democratizing efficient video production for all, March 2024. URL https://github.com/hpcaitech/Open-Sora.
- Jinxing Zhou, Jianyuan Wang, Jiayi Zhang, Weixuan Sun, Jing Zhang, Stan Birchfield, Dan Guo, Lingpeng Kong, Meng Wang, and Yiran Zhong. Audio-visual segmentation. In *European Conference on Computer Vision*, 2022.

A DISCUSSIONS

A.1 ETHICAL STATEMENT

In this work, we construct our JavisBench from publicly available academic datasets as well as YouTube videos. To ensure ethical compliance, we strictly adhere to YouTube's Terms¹ of Service and licensing policies. Specifically:

- Privacy Protection: We have taken measures to remove any personally identifiable or privacysensitive information from the collected data. No private, confidential, or user-specific metadata has been retained.
- Copyright Compliance: All data collection respects the original content licenses. We only
 utilize publicly accessible videos that are either explicitly licensed for research use or fall under
 fair-use considerations. No copyrighted content is redistributed or modified in violation of
 licensing terms.
- 3. Responsible Data Release: Any potential dataset release will fully comply with YouTube's data policies and ethical guidelines. We will first release the self-curated caption data to support the evaluation of all metrics except the AV-Quality terms in Tab. 1, and then ensure that shared sounding-videos are either appropriately anonymized or restricted in accordance with relevant regulations and platform policies.

By implementing these measures, we strive to maintain high ethical standards in data collection, use, and dissemination.

A.2 SOCIETAL IMPACT STATEMENT

This work proposes JavisDiT, a unified model for synchronized audio-video generation from openended prompts. While primarily intended as a contribution to foundational multimodal generation research, it has both positive applications and potential societal risks.

Positive Impact. JavisDiT may support creative industries by enabling efficient generation of audio-video content. Addressing the challenge of audio-video alignment has practical implications for animation, video conferencing, television broadcasting, and video editing—domains where synchronization is traditionally achieved through offline computation or extensive manual post-processing.

Potential Risks. As with other generative models, JavisDiT may be misused to produce highly realistic synthetic media for malicious purposes, including misinformation, impersonation, and manipulation. Its ability to generate temporally coherent sounding videos may increase the realism and persuasiveness of deepfakes. Specifically, fabricated audio-video clips could be used to spread false information or simulate events that never occurred, posing risks to public trust, political discourse, and media integrity.

Mitigation Strategy. To mitigate these concerns, we recommend responsible release practices, including usage gating and dataset transparency. Future work may explore watermarking and detection tools to distinguish synthetic output.

A.3 REPRODUCIBILITY STATEMENT

We provide detailed descriptions of model design, training, and evaluation in both the main paper and the appendix. Furthermore, all code, pretrained checkpoints, and processed datasets will be publicly released to ensure full reproducibility of our results.

A.4 LLM USAGE STATEMENT

Large Language Models (LLMs) were used solely as writing assistants, including tasks such as language polishing and presentation refinement. They were not involved in the conception of core ideas or designs.

¹https://www.youtube.com/t/terms

Table A1: Latency analysis for 240P4s sounding-video generation on H100 with 30 sampling steps.

Video Branch	Audio Branch	ST-Prior Modulation	AV-Interaction	Overall
13s	2s	13s	2s	30s

A.5 POTENTIAL LIMITATION AND FUTURE WORK

Despite the strong performance in joint audio-video generation, our JavisDiT has several limitations that present opportunities for future research:

- 1. **Scalability of Training Data**: Our model was trained on 0.6M text-video-audio triplets, which, while substantial, is still limited compared to the scale of some large vision-language models (*e.g.*, OpenSora (Zheng et al., 2024) takes over 60M data to build a foundation mode for video generation). Expanding the dataset with more diverse and higher-quality real-world audio-video samples could further enhance the model's ability to generalize across different domains and fine-grained synchronization patterns.
- 2. **Synchronization Evaluation Metrics**: While we introduce JavisScore for evaluating audio-video synchronization, its current accuracy of 75% suggests room for improvement. More robust synchronization metrics—potentially incorporating perceptual alignment assessments or human-in-the-loop evaluation—could further refine synchronization quality measurement in JAVG research.
- 3. **Efficiency and Computation Overhead**: JavisDiT employs a Diffusion Transformer for high-quality generation, but diffusion-based models tend to be computationally intensive, as analyzed in Tab. A1. While our approach achieves state-of-the-art results, generation speed and efficiency remain challenges (*e.g.*, generating a 2-second sounding-video at 720P/24fps/16kHz takes 6 minutes on one H100 GPU). Exploring accelerated sampling strategies or hardware optimization could improve efficiency.
- 4. **Benchmarking Across Resolutions and Durations**: Our evaluations primarily focus on a fixed resolution (240P) and duration (4s) setting in JavisBench. However, real-world applications may require generation at higher resolutions (*e.g.*, 1080P) and longer durations. Conducting benchmark tests across multiple settings would provide a more comprehensive understanding of current models' strengths and limitations.

By addressing these limitations, future iterations of JavisDiT could further enhance scalability, synchrony, efficiency, and adaptability, pushing the boundaries of joint audio-video generation.

B RELATED WORK

In the field of AIGC, multimodal generation has become a key topic, encompassing text-to-image (Rombach et al., 2022; Ramesh et al., 2022; Saharia et al., 2022), text-to-video (Ho et al., 2022; Wang et al., 2024d; Yu et al., 2023), and text-to-audio (Liu et al., 2023a; 2024b; Huang et al., 2023) generation, *etc.*. Among these, sounding video generation (Ruan et al., 2023; Xing et al., 2024; Wang et al., 2024b; Sun et al., 2024; Liu et al., 2024a) is drawing increasing attention due to its strong alignment with real-world applications. To generate synchronized audio and video, early works decomposed the task into two cascaded subtasks. a) generating video from text, then adding audio (Comunità et al., 2024; Jeong et al., 2024; Zhang et al., 2024; Ren et al., 2024; Hu et al., 2024); or b) generating audio first, then synchronizing video (Yariv et al., 2024; Jeong et al., 2023; Zhang et al., 2025). As an instance, MovieGen (Polyak et al., 2024) achieves movie-grade generation quality using this cascaded approach. From a methodological perspective, the community focus has shifted from UNet-based models (Rombach et al., 2022; Liu et al., 2024b) to DiT-based methods (Zheng et al., 2024; Polyak et al., 2024) to achieve more state-of-the-art performance.

Another line of work considers the limitations of pipeline approaches, such as error propagation and reliability issues, by focusing on end-to-end Joint Audio-Video Generation (JAVG). This paradigm aims to improve generation quality by modeling audio-video synchronization, including semantic consistency and temporal alignment, and has received growing attention (Ruan et al., 2023; Hayakawa et al., 2024; Ishii et al., 2024). For example, Ruan et al. (2023) introduce a diffusion block that enables cross-modal interaction during denoising. Xing et al. (2024) propose using ImageBind to align the semantic representations of the two modalities. Sun et al. (2024) introduce a hierarchical VAE, mapping audio and video to a shared semantic space while reducing their dimensionality.

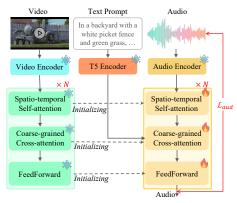


Figure A1: **Audio Pretraining**. Parameters are initialized from the video branch.

Table A2: Detailed settings for three-stage training.

Setting	Stage-1	Stage-2	Stage-3
trainable params	1.11B	29.3M	923.8M
learning rate	1e-4	1e-5	1e-4
warm-up steps	1000	1000	1000
weight decay	0.0	0.0	0.0
ema decay	0.99	-	0.99
dropout rate	0.0	0.0	0.0
training samples	788K	611K	611K
batch size	dynamic	dynamic	dynamic
epoch	13	1	2
training objective	rect. flow	contrastive	rect. flow
GPU days (H100)	16	8	256

AV-DiT (Wang et al., 2024b) employs a single DiT model to generate both video and audio modalities simultaneously, improving efficiency. Uniform (Zhao et al., 2025) also takes a single DiT for JAVG, by simply concatenating the video and audio latent tokens during the diffusion process, without any explicit synchronization guidance. SyncFlow (Liu et al., 2024a) utilizes STDiT (Zheng et al., 2024) blocks to enhance the video generation quality, with a temporal adapter to guide the audio generation process. Unfortunately, current JAVG methods either lack a strong backbone for audio and video generation or insufficiently model audio-video synchronization, resulting in suboptimal generation quality. To address these issues, we construct a novel DiT-based JAVG model, further enhanced by modeling hierarchical spatial-temporal synchronized prior features. Moreover, this paper aims to further advance JAVG by providing a more comprehensive, robust, and challenging benchmark.

C IMPLEMENTATION DETAILS

C.1 JAVISDIT MODEL CONFIGURATION

Model Architecture. As Fig. 2 illustrates, JavisDiT consists of two branches for video and audio generation, each comprising N=28 DiT blocks. Within each DiT block, the latent video and audio representations are processed through several modules, where all attention modules utilize 16 attention heads with a hidden size of 1152. The intermediate dimension of the FFN is $4\times$ the hidden size. Each attention and FFN module is preceded by a LayerNorm layer. In practical implementation, every attention module is followed by an FFN module to further enhance feature modeling. The video/audio latent is sequentially processed with: Spatial-SelfAttn-CrossAttn-Spatial-CrossAttn-Bi-CorssAttn-FFN-Temporal-SelfAttn-CrossAttn-Temporal-CrossAttn-Bi-CorssAttn-FFN for <math>N=28 times, resulting in our JavisDiT with 3.14B parameters in total.

Training Strategy. As Sec. 3.3 states, we carefully design a three-stage training strategy to achieve high-quality single-modal generation and ensure fine-grained spatio-temporal synchrony on generated videos and audios. In particular, the video branch of JavisDiT (including the ST-SelfAttn module, the Coarse-Graind CrossAttn module, and the FFN module) is initialized from OpenSora (Zheng et al., 2024) and frozen during the whole training stages. We use the weights of the video branch to initialize the audio branch for better convergence (see Fig. A1). The detailed configuration for the three-stage training strategy is displayed in Tab. A2.

Training Data Curation. For audio-pretraining (stage-1), the audio-caption pairs come from various public datasets, including AudioSet (Gemmeke et al., 2017b), AudioCaps (Kim et al., 2019), VGGSound (Chen et al., 2020a), WavCaps (Mei et al., 2024), Clotho (Drossos et al., 2020), ESC50 (Piczak, 2015), GTZAN (Sturm, 2013), MACS (Martín-Morató & Mesaros, 2021), and UrbanSound8K (Salamon et al., 2014). They contribute to a total of 788K training entries. For the ST-Prior estimator and the final JavisDiT training (stage-2&3), the data source comes from two newly-proposed sounding video datasets: MMTrail (Chi et al., 2024) and TAVGBench (Mao et al., 2024). MMTrail provides slightly more caption annotations (2M) than TAVGBench (1.7M) but has lower quality. Due to the resource limit, we collect a part of the original sounding-videos from

YouTube and filtered out around 80% human-talking videos with the FunASR ² tool. We finally collect 136K high-resolution video-audio-caption triplets from MMTrail and 475K from TAVGBench, resulting in 611K data for training in total. The data construction for ST-Prior estimator's contrastive learning is discussed in Sec. C.2.4.

Training Objective and Inference. The ST-Prior Estimator is trained with contrastive learning objectives, which will be detailed in Sec. C.2. For the DiT model, we use rectified flow (Liu et al., 2023b) as the denoising scheduler for better performance (Zheng et al., 2024; Polyak et al., 2024). The inference sample step is 30, with classifier-free guidance at 7.0. The video and audio latent are concurrently sampled at each inference step. The video encoder-decoder comes from OpenSora (Zheng et al., 2024), and the audio encoder-decoder comes from AudioLDM2 (Liu et al., 2024b). Both of them are frozen during training.

Detailed Evaluation Setup. In Sec. 5 in our manuscript, we conduct a comprehensive evaluation across multiple video and/or audio generation models on two settings: (1) our proposed Javis-Bench benchmark and (2) previously used AIST++ (Li et al., 2021) and Landscape (Lee et al., 2022) datasets. For our JavisBench, generative models are required to generate videos and 240p and 24fps, with audios at 16kHz sample rate. The duration is 4 seconds. For AIST++ (Li et al., 2021), we follow the literature (Ruan et al., 2023; Sun et al., 2024) to generate 2-second clips at the visual resolution of 256×256 , with 240p on landscape (Lee et al., 2022).

C.2 ST-PRIOR ESTIMATOR CONFIGURATION

C.2.1 OVERVIEW

 Since ImageBind (Girdhar et al., 2023) has built a consistent semantic space across multiple modalities, we use ImageBind's text encoder to extract the potential spatio-temporal prior from the input text caption. Formally, for a given input text s, we leverage the 77 hidden states from the last layer of ImageBind's text encoder, and learn N_s spatial tokens \mathbf{p}_s and N_t temporal tokens \mathbf{p}_t from a 4-layer transformer block \mathcal{P} . Notably, since the same event may occur at different locations and timestamps in different video-audio generations, the same text s can also produce different ST-Priors $(\mathbf{p}_s, \mathbf{p}_t)$. We adopt a sampling strategy akin to VAE, where our ST-Prior Estimator \mathcal{P} outputs the mean and variance of a Gaussian distribution, from which we sample a plausible $(\mathbf{p}_s, \mathbf{p}_t)$ to describe specific spatio-temporal events. This can be formalized as: $(\mathbf{p}_s, \mathbf{p}_t) = \mathcal{P}_{\phi}(s; \epsilon)$, where ϵ refers to the normalized Gaussian distribution.

As Fig. 3 states, the ST-Prior Estimator \mathcal{P}_{ϕ} is trained with contrastive learning. Specifically, we take the text prompts s and estimated prior $(\mathbf{p}_s, \mathbf{p}_t)$ as an *anchor*, and treat the corresponding synchronous video-audio (\mathbf{v}, \mathbf{a}) pairs from training datasets as *positive* samples. Then, we either synthesize an asynchronous video \mathbf{v}^- or audio \mathbf{a}^- , and take the original \mathbf{a} or \mathbf{v} to form *negative* pairs.

Formally, we also extract N_s/N_t positive or negative spatial/temporal embeddings from synchronous or asynchronous video-audio pairs: $(\mathbf{e}_s^+, \mathbf{e}_t^+) = \mathcal{E}_{\psi}(\mathbf{v}, \mathbf{a}), \ (\mathbf{e}_s^-, \mathbf{e}_t^-) = \mathcal{E}_{\psi}(\mathbf{v}, \mathbf{a}^-)$ or $\mathcal{E}_{\psi}(\mathbf{v}^-, \mathbf{a})$. Sec. C.2.2 displays more details.

The training objective is to make the priors $(\mathbf{p}_s, \mathbf{p}_t)$ closer to the synchronous embeddings $(\mathbf{e}_s^+, \mathbf{e}_t^+)$ while pushing them further from the asynchronous $(\mathbf{e}_s^-, \mathbf{e}_t^-)$, thereby equipping the priors $(\mathbf{p}_s, \mathbf{p}_t)$ with fine-grained conditioning capabilities. Several contrastive losses with the Kullback-Leibler (KL) divergence loss are utilized to optimize the spatial and temporal priors (see Sec. C.2.3):

$$\mathcal{L}_{prior} = \mathcal{L}_{contrast}(\mathbf{p}_s, \mathbf{e}_s^+, \mathbf{e}_s^-) + \mathcal{L}_{kl}(\mathbf{p}_s) + \mathcal{L}_{contrast}(\mathbf{p}_t, \mathbf{e}_t^+, \mathbf{e}_t^-) + \mathcal{L}_{kl}(\mathbf{p}_t). \tag{A1}$$

The details for negative sample construction are provided in Sec. C.2.4.

C.2.2 SPATIAL-TEMPORAL VA-ENCODING

For a given video clip with the corresponding audio, we first use OpenSora's (Zheng et al., 2024) and AudioLDM2's (Liu et al., 2024b) VAE-encoders to respectively map the raw data to the feature space, and derive a lightweight VA-Fuser module to extract the spatio-temporal embeddings from the video-audio pair. As shown in Fig. A2, the raw video $V \in \mathbb{R}^{T \times (H \times W) \times C}$ will be encoded

²https://github.com/modelscope/FunASR

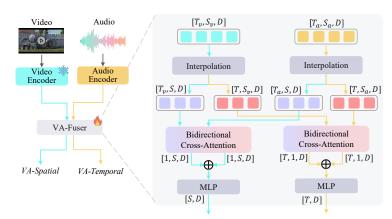


Figure A2: Architecture of VA-Fuser to encode spatio-temporal embeddings.

and reshaped to a video embedding $\mathbf{v} = E_v(V) \in \mathbb{R}^{T_v \times S_v \times D_v}$, where the downsampling rate for (T,H,W) is (4,8,8), and the output channel D=4. Similarly, the raw audio will be first transformed to the MelSpectrogram $A \in \mathbb{R}^{R \times M \times 1}$, and then encoded to an audio embedding $\mathbf{a} = E_a(A) \in \mathbb{R}^{T_a \times S_a \times D_a}$. We use a linear layer to unify video and audio channels into the hidden size D=512.

Then, we align the video and audio embeddings with the pre-defined spatial-temporal token numbers via linear interpolation. Specifically, we will obtain a video-spatial embedding $\mathbf{v}_s \in \mathbb{R}^{T_v \times S \times D}$, a video-temporal embedding $\mathbf{v}_t \in \mathbb{R}^{T \times S_v \times D}$, a audio-spatial embedding $\mathbf{a}_s \in \mathbb{R}^{T_a \times S \times D}$, and a audio-temporal embedding $\mathbf{a}_t \in \mathbb{R}^{T \times S_a \times D}$. Subsequently, we apply bidirectional attention (see Fig. 2(d)) for the spatial va-embedding-pair $(\mathbf{v}_s, \mathbf{a}_s)$ and the temporal va-embedding-pair $(\mathbf{v}_t, \mathbf{a}_t)$. After averaging the temporal dimension (T_v, T_a) and merging the video-audio embeddings, we use a 2-layer MLP module to project to the desired video-audio spatial embedding $\mathbf{e}_s \in \mathbb{R}^{S \times D}$. With similar operations, we can obtain the video-audio temporal embedding $\mathbf{e}_t \in \mathbb{R}^{T \times D}$.

As defined in Sec. 2, the extracted *spatial embedding* e_s should determine the occurrence of a specific event in a region of the video frame, with the corresponding appearance of matching frequency components in the audio spectrogram. On the other hand, the extracted *temporal embedding* e_t should identify the onset or termination of an event at a specific frame or timestamp in the video, with the corresponding start or stop of the response in the audio. We achieve these goals by a dual contrastive learning objective.

C.2.3 CONTRASTIVE TRAINING LOSSES

Our core objective is to align the spatial/temporal priors $(\mathbf{p}_s, \mathbf{p}_t)$ (extracted from text) more closely with the positive spatial/temporal embeddings $(\mathbf{e}_s^+, \mathbf{e}_t^+)$ derived from synchronized video-audio pairs, while pushing them away from the embeddings $(\mathbf{e}_s^-, \mathbf{e}_t^-)$ of negative samples (asynchronous video-audio pairs). This ensures that the priors $(\mathbf{p}_s, \mathbf{p}_t)$ provide the fine-grained spatio-temporal condition required for synchronized video-audio generation.

To achieve this, we designed four contrastive loss functions. For simplicity, we omit the subscripts of s, t in this part, as we apply the same loss functions on both spatial and temporal priors/embeddings.

- 1. Token-level hinge loss: $\mathcal{L}_{token}(\mathbf{p}, \mathbf{e}^+, \mathbf{e}^-) = |1.0 sim(\mathbf{p}, \mathbf{e}^+)| + |1.0 + sim(\mathbf{p}, \mathbf{e}^-)|$, where sim refers to the cosine similarity.
- 2. Auxiliary discriminative loss: $\mathcal{L}_{disc}(\mathbf{p}, \mathbf{e}^+, \mathbf{e}^-) = \mathcal{L}_{BCE}(\mathcal{D}_{\theta}(\mathbf{p}, \mathbf{e}^+), 1) + \mathcal{L}_{BCE}(\mathcal{D}_{\theta}(\mathbf{p}, \mathbf{e}^+), 0)$, where \mathcal{D}_{θ} is a 1-layer attention module parameterized with θ that use the [CLS] token to gather the information in \mathbf{p} and \mathbf{e} .
- 3. VA-embedding discrepancy loss: $\mathcal{L}_{vad}(\mathbf{e}^+, \mathbf{e}^-) = |1.0 + sim(\mathbf{e}^+, \mathbf{e}^-)|$, which aims to enlarge the discrepancy of positive va-embedding \mathbf{e}^+ and negative \mathbf{e}^- .
- 4. L2-regularization loss: $\mathcal{L}_{reg}(\mathbf{p}, \mathbf{e}^+) = \|\mathbf{p} \mathbf{e}^+\|_2$, which directly pushes the prior \mathbf{p} towards the positive (synchronous) va-embedding \mathbf{e}^+ .

Overall, the contrastive learning loss is combined with:

$$\mathcal{L}_{contrast}(\mathbf{p}, \mathbf{e}^+, \mathbf{e}^-) = \mathcal{L}_{token}(\mathbf{p}, \mathbf{e}^+, \mathbf{e}^-) + \mathcal{L}_{disc}(\mathbf{p}, \mathbf{e}^+, \mathbf{e}^-) + \mathcal{L}_{vad}(\mathbf{e}^+, \mathbf{e}^-) + \mathcal{L}_{reg}(\mathbf{p}, \mathbf{e}^+). \tag{A2}$$

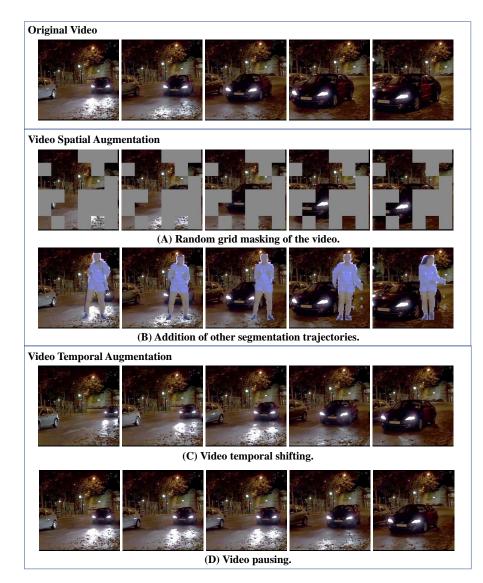


Figure A3: Video augmentation for spatial and temporal negative samples.

Specifically, we take the text prompts s and estimated prior $(\mathbf{p}_s, \mathbf{p}_t)$ as an *anchor*, and treat the corresponding synchronous video-audio (\mathbf{v}, \mathbf{a}) pairs from training datasets as *positive* samples. Then, we either synthesize an asynchronous video \mathbf{v}^- or audio \mathbf{a}^- , and take the original \mathbf{a} or \mathbf{v} to form *negative* pairs. It can be formulated as $(\mathbf{e}_s^+, \mathbf{e}_t^+) = \mathcal{E}_{\psi}(\mathbf{v}, \mathbf{a}), (\mathbf{e}_s^-, \mathbf{e}_t^-) = \mathcal{E}_{\psi}(\mathbf{v}, \mathbf{a}^-)$ or $\mathcal{E}_{\psi}(\mathbf{v}^-, \mathbf{a})$. The VA-Encoder \mathcal{E}_{ψ} is detailed in Sec. C.2.2, and the negative sample construction is introduced in Sec. C.2.4. We also discuss the efficacy of each loss function in Sec. E.3

C.2.4 NEGATIVE SAMPLE CONSTRUCTION

Given a synchronized video-audio pair from the training set, we design two approaches to synthesize *easy* and *hard* asynchronous videos/audios for optimization. (1) For *easy* negatives, we take AudioLDM2 (Liu et al., 2024b) to generate arbitrary audios from the text without referencing the corresponding video, which naturally results in asynchronous video-audio pairs. (2) For *hard* negatives, various augmentation strategies are employed to add, remove, or modify elements in either the original video or audio from GT-pairs, creating more fine-grained spatio-temporally asynchronous pairs. As detailed below, these methods further enhance the discriminative power of the ST-Prior, and Fig. A3 and Fig. A4 showcase some representative augmentation examples.

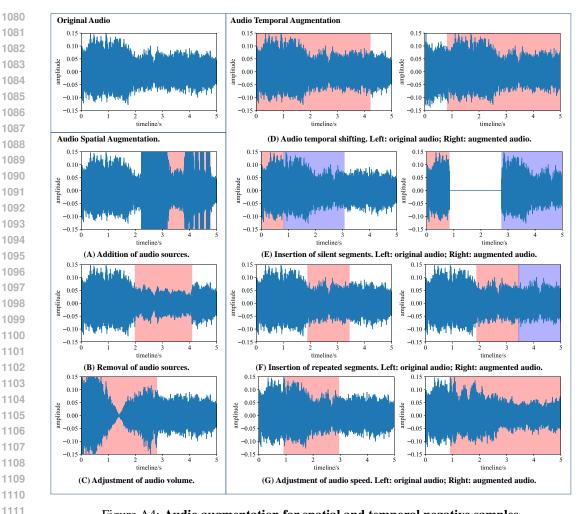


Figure A4: Audio augmentation for spatial and temporal negative samples.

Video Spatial Augmentation

- Random Masking (Fig. A3(A)). Videos are divided into 6×6 grids, with a mask ratio p uniformly sampled from (0.2, 0.8)). A proportion p of the grids is randomly masked.
- Adding Subject Trajectories (Fig. A3(B)). The trajectories from the SA-V dataset (Ravi et al., 2024) are overlaid to simulate new sound-producing objects. Trajectories are preprocessed using RepViT-SAM (Wang et al., 2024a) for interpolation from 6fps to 24fps, retaining objects with an average size > 32 × 32 pixels. During augmentation, a random trajectory is selected from the pool and added to the video.

Video Temporal Augmentation

- *Video Temporal Shifting* (Fig. A3(C)). The video is cyclically rearranged from a random starting point to create temporal misalignment.
- *Video Pausing* (Fig. A3(D)). A random frame is duplicated for at least 0.5 seconds at a chosen point, pausing the video and disrupting temporal synchronization.

Audio Spatial Augmentation

All audio augmentations (except temporal shifting) focus on sound-producing intervals, avoiding silent or noisy segments. We use QwenPlusAPI³ to extract potential sound-producing objects from audio descriptions, and utilize AudioSep (Liu et al., 2024c) to separate sound sources and labels intervals using RMS. Separation results are precomputed and stored as metadata, guiding target augmentations.

³https://help.aliyun.com/zh/model-studio/developer-reference/use-qwen-by-calling-api

- Audio Source Addition (Fig. A4(A)). Additional audio from other dataset sources is mixed with the original one, causing spatial content misalignment between audio and video.
- Audio Source Removal (Fig. A4(B)). A separated sound source from the original audio is randomly removed to induce more spatial misalignments.
- Audio Volume Adjustment (Fig. A4(C)). The target audio's volume is modified by a randomly selected transformation (e.g., cosine, sine, linear increase/decrease) to adjust its amplitude.

Audio Temporal Augmentation

- Audio Temporal Shifting (Fig. A4(D)). Audio is shifted by extracting a new interval from the original, creating temporal misalignment.
- Silent Segment Insertion (Fig. A4(E)). Silent segments are inserted into the target interval, disrupting synchronization with the video timeline in subsequent frames.
- Repeated Segment Insertion (Fig. A4(F)). Target segments are repeated after a random transformation interval, making another type of timeline misalignment.
- Audio Speed Adjustment (Fig. A4(G)). Playback speed within the target interval is altered to $0.5 \times$ or $2 \times$, causing another type of timeline desynchronization.

D MORE DETAILS ON JAVISBENCH

D.1 MOTIVATION

A strong generative model must produce diverse video content and audio types while ensuring fine-grained spatio-temporal synchronization. However, we have noticed that current benchmarks and metrics lack sufficient diversity and robustness for comprehensive evaluation. For instance, the commonly used AIST++ (Li et al., 2021) and Landscape (Lee et al., 2022) benchmarks focus only on limited scenarios (human dancing for AIST++ and natural scenarios for landscape) with merely 20-100 test samples. On the other hand, although TAVGBench (Mao et al., 2024) provides 3,000 test samples (but not released yet), it lacks detailed analysis of different scenarios in practical audio-video generations. While larger-scale datasets such as VGGSound (Chen et al., 2020a) and AudioSet (Gemmeke et al., 2017a) provide millions of video-audio pairs for training and testing, they are typically limited to a single-event taxonomy (e.g., human voice, engine) and do not reflect multi-event dynamics within a single video clip. Consequently, they are not suitable for evaluating spatial-temporal alignment in complex real-world environments. Moreover, all videos in VGGSound are captured in the wild, resulting in a lack of in-house scenes. Both of the two datasets fall short in scene diversity, underrepresenting diverse categories such as 2D/3D animations and industrial settings, which are crucial for evaluating generative models under varied and challenging conditions. In addition, the video-audio synchronization metric AV-Align (Yariv et al., 2024) cannot produce accurate evaluation on complicated scenarios, e.g., with multiple sounding events (Ishii et al., 2024; Mao et al., 2024).

We are therefore motivated to propose a more challenging benchmark to evaluate generative models from various dimensions on a larger data scale in Sec. 4.1, with a particular focus on spatio-temporal synchronization. We also introduce a robust metric to quantify the fine-grained spatio-temporal alignment between generated video and audio pairs in Sec. 4.2. Here we provide more details of the benchmark construction and metric verification.

D.2 TAXONOMY

To comprehensively evaluate the capabilities of joint video-audio generation models, we designed five evaluation dimensions from different aspects: (1) *Event Scenario*, (2) *Video Style*, (3) *Sound Type*, (4) *Spatial Subjects*, and (5) *Temporal Composition*. Leveraging GPT-4 (Achiam et al., 2023), we enumerated 3-5 primary categories under each aspect, contributing to a hierarchical categorization system with 5 dimensions and 19 categories. Tab. A3 presents the detailed definitions and clarifications.

D.3.2 QUALITY-BASED PRE-FILTERING

Following OpenSora (Zheng et al., 2024), we utilize a series of filtering tools to improve data quality, including:

Table A3: Clarification of the category taxonomy of our JavisBench.

Aspect	Category	Description and Examples
	Natural Scenario	Scenes dominated by natural environments with minimal human interference, such as forests, oceans, and mountains.
	Urban Scenario	Outdoor spaces shaped by human activity, including cities, villages, streets, and parks.
Event Scenario	Living Scenario	Indoor environments where daily human activities occur, like houses, schools, and shopping malls.
	Industrial Scenario	Work-oriented spaces related to industrial or energy activities, such as factories, construction sites, and mines.
	Virtual Scenario	Imaginative or abstract settings, including virtual worlds, sci-fi cities, and artistic installations.
	Camera Shooting	Filmed with handheld, fixed, or drone cameras, including slow-motion footage.
Y' 10.1	2D-Animate	Styles like hand-drawn animation, flat animation, cartoon styles, or watercolor illustrations.
Visual Style	3D-Animate	Photorealistic styles, sci-fi/magical effects, CG (Computer Graphics), or steam- punk aesthetics.
	Ambient Sounds	Sounds that occur naturally in the environment, including both natural and human-influenced surroundings. This category includes sounds like wind, rain, water flow, animal sounds, human activity (e.g., traffic, construction), and urban noise.
	Biological Sounds	Sounds produced by living creatures (e.g. animals, birds). This includes vocalizations such as barking, chirping, growling, as well as non-vocal human sounds like heartbeat, and other physical noises.
Sound Type	Mechanical Sounds	Sounds generated by man-made machines, devices, or mechanical processes. This includes the noise of engines, motors, appliances, and any mechanical or electronic noise. This category also includes malfunction sounds (e.g., malfunctioning machinery or alarms).
Sound Type	Musical Sounds	Sounds related to music or musical performance, including both human- generated and instrument-generated sounds and melodies. This category covers singing, instrumental performances, as well as background music used in vari- ous media formats.
	Speech Sounds	Sounds generated from human speech, whether in conversation, dialogue, public speeches, debates, interviews, or monologues. This category specifically covers linguistic communication in various contexts, whether formal, informal, or contentious.
	Single Subject	There is only one primary object or source producing sound in the scene.
Spatial Composition	Multiple Subject	There are multiple primary objects that (or potentially can) make sounds in the scene.
Spauai Composition	Off-screen Sound	The source of the sound is not visible in the scene but logically exists (<i>e.g.</i> , a car engine outside the camera view).
	Single Event	The audio contains only one event, with no overlapping sounds. For example, "a single dog barking without background noise."
Temporal Composition	Sequential Events	There are multiple events occurring sequentially, with no overlap. For example, "the applause begins after the music performance ends."
	Simultaneous Events	Multiple audio sources are present simultaneously, such as "a person speaking

D.3 DATA CURATION

D.3.1 COLLECTION

We collected data from two sources to construct the benchmark. First, we incorporate existing datasets' test sets, including those in the JAVG domain like Landscape (Lee et al., 2022) and AIST++ (Li et al., 2021) (despite their limited coverage) and others from sounding-video understanding tasks such as FAVDBench (Shen et al., 2023) and AVSBench (Zhou et al., 2022). Second, we expanded the benchmark by crawling videos uploaded to YouTube between June 2024 and December 2024 to prevent data leakage in previous methods' training scope (Mao et al., 2024). Using the previously defined categorization system, we prompt GPT4 to generate potential keywords for specific categories, enabling targeted video collection and avoiding indiscriminate gathering of homogeneous data, which significantly improves curation efficiency. This stage yields around 30K sounding video candidates.

while music plays in the background."

- 1. Scene Cutting. Given the extended duration of some YouTube videos (up to several hours), we employed PySceneDetect⁴ for scene detection and segmented the videos at identified scene transitions. Each clip was constrained to a length of 2–60 seconds, resulting in approximately 230K sounding video clips.
 - 2. *Aesthetic Filtering*. We filter out videos with aesthetic scores (Schuhmann et al., 2021) lower than 4.5, remaining in 70K clips.
 - 3. *Optical-flow Filtering*. We use UniMatch (Xu et al., 2023) to estimate the motion quality of the given videos, and remove the (static) videos whose score is lower than 0.1. This produces 46K candidates.
 - 4. *OCR Filtering*. We use DBNet (Liao et al., 2022) to detect and filter out the videos containing more than 5 text regions, resulting in a smaller 30K scope.
 - 5. Speech Filtering. As the Internet-available videos contain too much human speech (including talking and voiceover), we use FunASR⁵ to detect and remove the speech videos. In this step, there are still around 20% speech videos remaining, due to the non-perfect speech detection by FunASR. We keep this part of videos in our final benchmark to ensure the diversity of video-audio sources. This step leads to 22.3K sounding video clips.

D.3.3 ANNOTATION

Since most of the collected data lacks captions, we designed a generic pipeline to generate captions for the video-audio pairs, and then categorize them into corresponding classes from different aspects, as illustrated in Sec. D.2.

- First, we use Qwen2-VL-72B (Wang et al., 2024c) to generate detailed captions for videos. Metadata from the data source (e.g., object labels or YouTube keywords) is included in the context to enhance caption quality.
- Next, we use Qwen2-Audio-7B (Chu et al., 2024) to caption the audio in detail. We do not incorporate the previously generated video captions as contextual input, as it produces even more hallucinations with the input bias (*i.e.*, wrongly identifying a sound that will happen in corresponding visual scenarios but actually does not exist in the current audio).
- Then, we employ Qwen2.5-72B-Instruct (Yang et al., 2024a) to merge video and audio captions into a unified text prompt, serving as the textual condition for the JAVG task. During this process, Qwen2.5-72B is prompted to reason on the video-audio captions and identify apparent logical mistakes (*e.g.*, there is a dog in the video but Qwen2-Audio gives sounds from a car engine) or missing captions (*e.g.*, unknown sound source).
- Finally, we query Qwen2.5-72B-Instruct again to classify the data points based on the video, audio, and generated captions, assigning each entry to the appropriate category in the classification system. We do not prompt multimodal LLMs to do this categorization because the generative models to evaluate will only receive the text caption as inputs.

After removing the logical conflict captions and classification results that fail to parse, we obtain 19.4K sounding video clips with detailed captions and hierarchical categorization results.

D.3.4 CONTENT-BASED POST-FILTERING

To build a diversified and balanced benchmark to evaluate joint audio-video generation, we conduct another post-filtering based on the categorization results obtained above. In particular, we further remove videos that only contain background music and speech voice. By doing so, nearly half of videos are filtered out, since a large part of YouTube videos rely on music and voice acting to attract viewers. After human checking, we obtained 10,140 samples in our JavisBench with diverse data sources and fine-grained category annotations, setting a new standard to facilitate a comprehensive evaluation for future joint audio-video generation (JAVG) methods. The category statistics can be found in Fig. 4, and Tab. A4 displays the data sources before and after our filtering strategies.

⁴https://github.com/Breakthrough/PySceneDetect

⁵https://github.com/modelscope/FunASR

Table A4: Data source composition before and after filtering strategies of our JavisBench

	YouTube	FAVDBench	AVSBench	Landscape	AIST++	Total
Before Filtering	30,107	1,000	804	100	20	32,031
After Filtering	8,507	833	680	100	20	10,140

D.4 EXTENDED SPECIFICATION ON JAVISSCORE EVALUATION SUITE

D.4.1 IMPLEMENTATION DETAILS

 Given a generated video-audio pair (V, A), we use ImageBind (Girdhar et al., 2023) to estimate audio-visual synchrony with following steps:

- 1. Temporal Segmentation via Sliding Windows. (V, A) is chunked into several segments with 2-seconds window size and 1.5-seconds overlap $(\mathcal{C} = \{(V_1, A_1), (V_2, A_2), \cdots (V_W, A_W)\})$. The 2-second window ensures compatibility with the audio encoder's minimum processing length (Girdhar et al., 2023), preventing suboptimal feature extraction. The 1.5-second overlap enhances continuity and robustness, allowing each frame to be evaluated within multiple temporal contexts. This mitigates artifacts caused by abrupt segmentation boundaries.
- 2. Frame-wise Audio-Visual Similarity Computation. Inspired by Mao et al. (2024), we calculate the cosine similarity between each frames ($\mathcal{F} = \{V_{i,1}, V_{i,2}, \cdots V_{i,w}\}$) and the whole audio clip A_i by using ImageBind's vision and audio encoders (E_v, E_a). Frame-wise similarity captures fine-grained temporal dynamics, ensuring transient events (e.g., lip movements, object interactions) are accurately modeled.
- 3. **Segment-wise Synchronization Estimation**. Instead of averaging all similarity scores, we select the 40% least synchronized frames (with lower similarities) and compute their mean to obtain the synchronization score for each window. By focusing on the least synchronized frames, the metric becomes more sensitive to local resynchronization, as a simple mean can be biased by a considerable proportion (*e.g.*, 70%) of synchronized frames.
- 4. **Global Synchronization Estimation**. The final JavisScore is computed by averaging the window-level scores across all segments, balancing local variations while maintaining sensitivity to desynchronization patterns. Due to window overlap, each video frame is evaluated multiple times, reducing the influence of outliers and providing a more stable final score.

These steps can be formulated as:

$$S_{Javis} = \frac{1}{W} \sum_{i=1}^{W} \sigma(V_i, A_i), \quad \sigma(V_i, A_i) = \frac{1}{k} \sum_{j=1}^{k} \underset{\min}{\text{top-}k} \{ \cos\left(E_v(V_{i,j}), E_a(A_i)\right) \}. \tag{A3}$$

This approach effectively differentiates synchronized and desynchronized video-audio pairs.

D.4.2 VERIFICATION ON METRIC QUALITY

This paper provides a quantitative comparison between the accuracy of video-audio synchrony metrics. We constructed a validation dataset consisting of 3,000 audio-video pairs to evaluate the effectiveness of our proposed metric, JavisScore. The validation set is initialized by 1,000 positive (synchronous) video-audio pairs from our JavisBench dataset, and we construct around 2,000 negative pairs from three different sources:

- 1. The first 1,000 negative pairs are generated through online augmentation/transformation on the positive pairs. Given a synchronized video-audio pair (V^+, A^+) , we reuse the augmentation strategies in Sec. C.2.4 to generate asynchronous video-audio pairs, *i.e.*, (V^+, A^-) or (V^-, A^+) . The four augmentation types (video-spatial, video-temporal, audio-spatial, audio-temporal) produce 250 negative samples, resulting in the first 1,000 negative pairs.
- 2. The second 500 negative pairs come from separate generations. In particular, we randomly split the 1,000 positive samples into two parts, and select the first 500 samples to prompt AudioLDM2 (Liu et al., 2024b) to generate audios conditioned solely on text captions, without accessing the corresponding videos. This naturally leads to 500 asynchronous video-audio pairs.
- 3. The third 500 negative video-audio pairs are generated by a preliminary JAVG model conditioned on the other 500 text prompts from positive samples. These generated pairs were then



Figure A5: Complex Scenario with multi-source sounds at a time.

manually annotated to determine whether they were synchronized or not. We finally obtained 411 hard negative samples with 89 positive video-audio pairs.

In the evaluation dataset, each positive sample is related to 2 negative video-audio pairs. One is constructed by online augmentation on the positive entry (the first negative source), and the other is generated from text captions of the positive data point (the second or third negative source).

Then, we compare our proposed JavisScore with the previously-used AV-Align (Yariv et al., 2024) metric on the 3,000 evaluation samples, and the result is displayed in Tab. A5. We calculate the AV-Align scores and our JavisScore on all the 3000 samples, and compute the AUROC value for the binary classification (ideally, a VA-synchrony metric should produce higher scores on positive video-audio pairs), which shows JavisScore achieves approximately 0.13 higher AUROC than AV-Align. Then, as each positive sample is related to 2 negative video-audio pairs,

Table A5: Comparison of VA-synchrony metrics. † means we relax the criteria to $Acc = \mathbb{I}\{s^+ \geq s^-\}$.

Metric	AUROC	Accuracy
Random-Guess	0.5000	0.5000
AV-Align	0.5296	0.5254
DeSync	0.5742	0.3961
DeSync †	0.5742	0.7797
JavisScore	0.6533	0.7514

we can calculate the prediction accuracy of $1000 \times 2 = 2000$ paired positive-negative samples. It is viewed as correct when a metric assigns a higher score on the positive sample than the corresponding negative video-audio pair. According to Tab. A5, our metric significantly surpasses AV-Align by 23% accuracy, demonstrating the efficacy of video-audio synchrony measurement.

Notably, AV-Align nearly performs as a random-guess (with a near 0.5 AUROC and Accuracy). This is because AV-Align simply uses optical flow to capture the video dynamics, and match with audio dynamics estimated by onset detection results. If there is a concurrent "pulse" at the video and audio timelines, AV-Align will produce a high synchrony score. However, in real-world scenarios, the video optical flow and audio onset detection cannot precisely capture the start and end of the visual-audio event. As exemplified in Fig. A5, when a person in the video speaks, the movements of their mouth are subtle and often too minimal to be effectively captured by optical flow. Moreover, in environments with strong background noise (*e.g.*, TV programs), the onset of human talking also becomes difficult to detect. The scores generated by AV-Align are unreliable in such complex scenarios. In contrast, our proposed JavisScore leverages the high-dimensional semantic space of ImageBind (Girdhar et al., 2023) to compute synchronization at the second-level, robustly distinguishing between synchronized and unsynchronized cases. It is worth noting, however, that our metric does not achieve 100% accuracy. Developing a more precise evaluation metric remains a significant challenge in the JAVG domain, and we hope this work will inspire further advancements in this critical area.

In addition, SynchFormer (Iashin et al., 2024) employs a 21-category classifier to estimate the temporal shift degree between video and audio, which is then used to compute a synchrony score (lower is better). However, as shown in Tab. A5, its discrete predictions result in poor separation between positive and negative samples, yielding low accuracy and AUROC. Only when we relax the evaluation metric to a $Acc = \mathbb{I}\{s^+ \geq s^-\}$ criterion does its accuracy approach that of JavisScore. Besides, SynchFormer is inherently limited to predicting temporal shifts and is not applicable to the diverse spatiotemporal misalignment scenarios considered in the real world (see Sec. C.2.4). Nevertheless, its approach suggests a promising direction: learning a neural network–based synchrony metric via supervision, which we leave as future work.

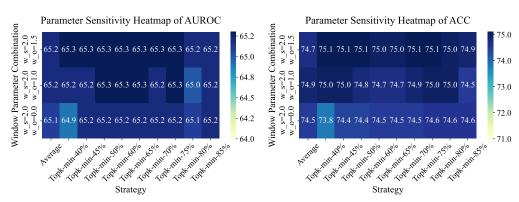


Figure A6: **Parameter sensitivity evaluation of our JavisScore.** Our method presents stable and robust video-audio synchrony estimate at various settings. We finally choose (2-second window size, 1.5-second overlap, topmin-40%) due to the relatively better performance.

Besides, we also conducted an ablation study to evaluate the parameter sensitivity of our JavisScore, including the sliding window size, overlap length, and score computation strategy. As suggested by Fig. A6, the optimal parameters were identified as a sliding window size of 2, an overlap length of 1.5, and selecting the top 40% minimum strategy. However, other settings do not significantly reduce the performance of our metric. The worst accuracy, for example, still achieves around 74% and outperforms AV-Align of 52.5% by a large margin. The results further demonstrate the robustness of our JavisScore metric.

D.5 IN-DEPTH ANALYSIS OF EVALUATION RESULTS

Fig. A7 provides a more fine-grained, multi-dimensional, and comprehensive evaluation of our model on JavisBench, offering an in-depth analysis of the current SOTA models' limitations in real-world joint audio-video generation tasks. Based on these results, we summarize the following key insights:

- Insufficient unimodal modeling capability in rare scenarios: (1) The FVD in Event Scenario (1st row, 1st column) indicates poor video quality in *industrial* and *virtual* scenes, suggesting weaker modeling capability in these domains. (2) The FVD in Visual Style (1st row, 2nd column) shows a significant disparity between generated and real videos in *3D-animate* scenes, potentially due to insufficient training data. (3) The FAD in Sound Type (r3c3) suggests weaker audio modeling in *ambient*, *biological*, and *mechanical* categories, likely because these categories are too broad. This observation aligns with Event Scenario (r3c1), where *natural* and *industrial* scenes also exhibit high FAD values.
- Unimodal quality does not directly correlate with text consistency: (1) The TV-IB in Visual Style (r4c2) reveals that 2D/3D-animate scenes exhibit poor text-following capability, despite their unimodal quality. (2) The CLAPScore in Sound Type (r7c3) suggests weak audio-text alignment in ambient and mechanical scenes, potentially due to a lack of corresponding audio-text pairs in the first-stage (Audio Branch) training data.
- **Poor AV-synchronization in challenging scenarios**: (1) The JavisScore in Event Scenario, Visual Style, and Sound Type (r10c1-3) show that scenarios with poor unimodal quality also suffer from weak audio-video synchronization, particularly in *virtual environments*, 2D/3D-animate styles, and musical sound types. (2) The JavisScore in Spatial and Temporal Composition (r10c4-5) indicate significantly lower AV synchronization performance for *complex events* (e.g., sequential/simultaneous events, multiple sound sources, or off-screen sounds) compared to *simpler scenes* (e.g., single events with a single sound source).

In conclusion, current SOTA models still struggle with *rare and complex scenarios*, both in terms of *audio-video generation quality* and *synchronization performance*. The JAVG community still faces significant challenges in bridging the gap between research models and real-world applications.

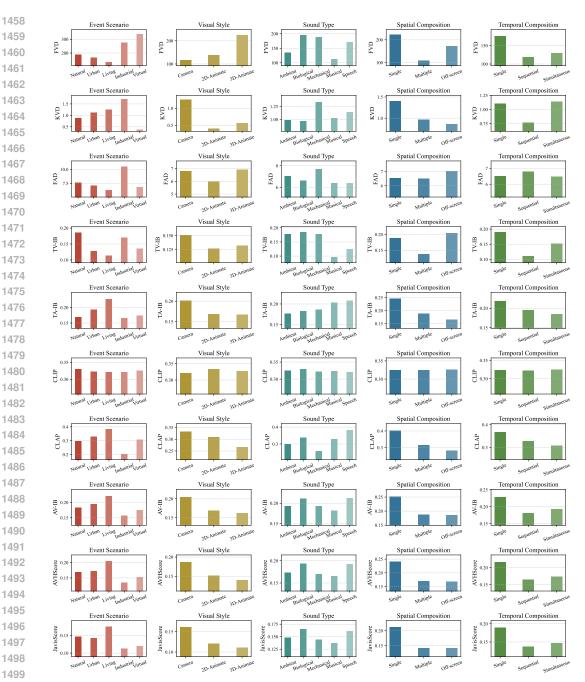


Figure A7: Metric distribution on all JavisBench's taxonomy.

E ADDITIONAL EXPERIMENTS

E.1 EVALUATION ON AUDIO GENERATION QUALITY

Tab. A6 presents the evaluation of our model's performance in the first training stage, concerning both unimodal text-to-audio generation quality and text-following capability. Specifically, we employ two datasets for evaluation:

• AudioCaps test set (Kim et al., 2019): Filtered by AudioLDM2 (Liu et al., 2024b), this dataset consists of 964 samples, each 8–10 seconds long. Models are required to generate 10-second audio clips.

Table A6: **Evaluation on audio generation.** After sufficient training iterations at stage 1, our JavisDiT presents moderate audio generation performance on in-domain test set AudioCaps (Kim et al., 2019) and comparable quality on out-of-domain test set JavisBench-mini.

Method	Aı	udioCaps (l	(nD)	JavisBench-mini (OoD)		
	FAD↓	TA-IB↑	CLAP↑	FAD↓	TA-IB↑	CLAP↑
AudioLDM2 (Liu et al., 2024b)	2.01	0.205	0.487	4.62	0.198	0.428
JavisDiT-audio-ep13	5.88	0.145	0.319	5.11	0.165	0.401
JavisDiT-audio-ep55	<u>5.19</u>	0.164	0.356	4.68	<u>0.194</u>	0.432

• **JavisBench-mini**: A subset containing 1,000 samples from JavisBench ranging from 2 to 10 seconds in length, which is also utilized in ablation studies in Sec. 5.3). Models are required to generate 4-second audio clips.

On these datasets, we first analyze the performance progression of our model across different training stages and then compare it against the current state-of-the-art (SOTA) model, AudioLDM2 (Liu et al., 2024b).

As shown in Tab. A6, with increased training iterations (from epoch 13 to epoch 55), our model demonstrates improvements in both *audio generation quality* (*e.g.*, FAD decreases from 5.88 to 5.19 on AudioCaps) and *text consistency* (*e.g.*, TA-IB increases from 0.165 to 0.194, while CLAPScore improves from 0.401 to 0.432 on JavisBench-mini). Furthermore, our model achieves comparable audio generation performance to AudioLDM2 on JavisBench-mini, further validating the effectiveness of our DiT-based architecture.

Notably, our model performs worse than AudioLDM2 on AudioCaps, which can be attributed to two primary factors: (1) *Training strategy*: AudioLDM2 is explicitly trained for 10-second audio generation, whereas our model supports variable-length audio generation. Consequently, our model currently lacks sufficient training data and iterative steps specifically for 10-second audio generation. (2) *Potential data leakage*: The training data of AudioLDM2 includes the training set of AudioCaps, leading to a smaller in-domain test gap, which increases the risk of overfitting. In contrast, JavisBench is collected from more diverse real-world YouTube audio sources, and its domain-specific soundscapes are not exposed to either AudioLDM2 or our model during training. Thus, performance on the *out-of-domain* JavisBench-mini provides a more realistic reflection of a model's usability in real-world scenarios.

In future iterations, we will continue enhancing the unimodal audio generation quality, as this is a crucial prerequisite for achieving precise *video-audio synchronization*.

E.2 INVESTIGATION OF ST-PRIOR'S EXTRACTION AND REPRESENTATION

In Fig. A8, we investigate the representation capability of ST-priors from three aspects.

Firstly, since the latent representations of video (v) and audio (a) contain different numbers of spatial and temporal tokens $(e.g., \text{ for a } 240\text{P}, 24\text{fps}, 16\text{kHz}, 4\text{-second sounding video}, <math>v \in \mathbb{R}^{400 \times 30 \times C_v}$ while $a \in \mathbb{R}^{16 \times 64 \times C_a}$), we also experimented with various token allocation strategies for the spatial and temporal priors. These included ratios such as 1:2 (n16x32), 1:1 (n32x32), and 2:1 (n32x16). As shown in Fig. A8, the 1:1 ratio achieved the best performance. This could be attributed to the inherent disparity in the latent shapes: video latent contains significantly more spatial tokens than temporal tokens, while audio latent stays in the opposite situation. Thus, a balanced 1:1 prior token ratio helps minimize bias.

Next, using the 1:1 ratio, we explored different configurations of prior *numbers* and *dimensions*. The results in Fig. A8 indicate that increasing the prior number and dimension consistently improve performance, demonstrating the scalability of our approach. In this work, we chose the "n32x32+d128" configuration as it offers a good trade-off between performance and training cost, considering the diminishing returns of further scaling.

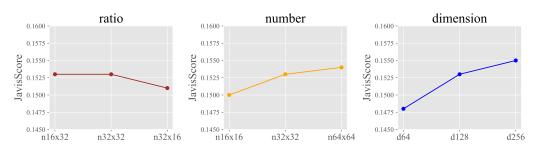


Figure A8: **Further ablation on ST-Prior hyper-parameters,** including (1) spatial-temporal token ratio, (2) token number, and (3) embedding dimension. Our default setting of "n32x32-d128" is a good trade-off between performance and training cost.

Table A7: **Ablation of ST-Prior loss functions.** All training objectives jointly contribute to video-audio synchronization.

\mathcal{L}_{token}	\mathcal{L}_{disc}	\mathcal{L}_{vad}	\mathcal{L}_{reg}	1	AV-Consist	tency	AV-Syı	nchrony
≈token	≈aisc	≈vaa	~reg	IB-AV↑	CAVP↑	AVHScore↑	AV-Align↑	JavisScore↑
√	×	×	×	0.190	0.799	0.167	0.097	0.133
V	√	×	×	0.193	0.799	0.170	0.102	0.136
v/	v /	1/	×	0.196	0.800	0.174	0.096	0.140
V	$\sqrt{}$	×	√	0.202	0.800	0.179	0.107	0.159
$\sqrt{}$	$\sqrt{}$	\checkmark	$\sqrt{}$	0.211	0.801	0.190	0.122	0.153

E.3 TRAINING LOSSES FOR ST-PRIOR'S ESTIMATION

In Sec. C.2.3, we introduce four loss functions to train our ST-Prior Estimator: (1) Token-level hinge loss: \mathcal{L}_{token} , (2) Auxiliary discriminative loss: \mathcal{L}_{disc} , (3) VA-embedding discrepancy loss: \mathcal{L}_{vad} , and (4) L2-regularization loss: \mathcal{L}_{reg} . This section presents a detailed ablation study on the efficacy of utilized loss functions.

According to Tab. A7, \mathcal{L}_{disc} brings slight improvement in addition to the original \mathcal{L}_{token} (e.g., 0.193 vs. 0.190 for IB-AV), as they share the same optimization purpose —— pushing the text prior anchor to the positive video-audio samples while pushing away from negative samples —— and differ only in the specific gradient back-propagation mechanism.

On the other hand, \mathcal{L}_{vad} considerably enhances the synchrony of generated video-audio pairs (e.g., 0.140 vs. 0.136 on JavisScore), thanks to its ability to maximize the divergence between positive and negative video-audio samples themselves. However, \mathcal{L}_{reg} offers even greater benefits due to its smoother regularization, facilitating the convergence of the text prior to the positive embeddings.

Moreover, the combination of all loss functions achieves the best performance (*e.g.*, 0.153 of JavisScore), as they collaboratively address the same goal from different perspectives: embedding video-audio synchrony into the text prior and ensuring it captures the semantic meaning of synchronization. This comprehensively validates both our motivation and the effectiveness of our methodology.

E.4 AUXILIARY COMPARISON WITH BASELINE MODELS

In Sec. 5 in our manuscript, we demonstrate a comprehensive comparison with cascaded audio-video generators and our joint JavisDiT model. Since our model takes AudioLDM2 (Liu et al., 2024b) and OpenSora (Zheng et al., 2024) as backbone audio and video encoder-generator, for T2A+A2V pipeline, we take AudioLDM2 to generate pre-positive audios to support final audio-video generation; for T2V+V2A, OpenSora is utilized to generate pre-positive videos.

This demonstrates another setting: we take the audios and videos generated from our JavisDiT as pre-positive content to support cascaded audio-video generation. We report the performance on

Table A8: **Auxiliary comparison on JavisBench-mini.** Prepositive audios (for A2V methods) and videos (for V2A methods) are provided by our JavisDiT. The joint JavisDiT shows consistent superiority against pipeline methods.

Method	AV-Quality			Text-Consistency				AV-Consistency			AV-Synchrony
	FVD↓	KVD↓	FAD↓	TV-IB↑	TA-IB↑	CLIP↑	CLAP↑	AV-IB↑	CAVP↑	AVHScore ↑	JavisScore ↑
- T2A+A2V											
TempoTkn (Yariv et al., 2024)	913.5	7.2	-	0.091	-	0.200	-	0.133	0.790	0.122	0.105
TPoS (Jeong et al., 2023)	1294.5	<u>5.1</u>	-	0.100	-	0.238	-	0.128	0.749	0.114	0.080
- T2V+V2A											
ReWaS (Jeong et al., 2024)	-	-	13.3	-	0.123	-	0.237	0.109	0.793	0.103	0.078
See&Hear (Xing et al., 2024)	-	-	7.6	-	0.072	-	0.249	0.168	0.795	0.090	0.071
FoleyCftr (Zhang et al., 2024)	-	-	10.2	-	0.149	-	0.301	0.204	0.798	0.179	0.149
- T2AV											
JavisDiT(Ours)	327.8	1.9	7.6	0.141	0.184	0.322	0.314	0.203	0.799	0.181	0.153

JavisBench-mini in Tab. A8 for efficient evaluation, where our joint JavisDiT shows consistent superiority against pipeline methods, *e.g.*, achieving the highest JavisScore of 0.153.

E.5 MORE GENERATION EXAMPLES

Fig. A9 presents realistic audio-visual pairs generated by our JavisDiT across diverse scenarios, including industrial, outdoor, and natural environments. These results encompass a wide range of visual styles, such as camera-captured footage, 2D/3D animations, as well as various audio types, including mechanical, musical, ambient, and biological sounds. Our model effectively maintains video-audio synchronization across cases involving single or multiple sounding subjects, as well as single, sequential, or simultaneous sounding events. All multimedia resources are available in the supplementary materials.

E.6 DISCUSSION: EXTENSION TO X-CONDITIONAL GENERATION

Built on diffusion models with transformers (DiT), our JavisDiT can be easily extended to support various conditional video-audio generation tasks, as shown in Fig. A10. Inspired by UL2 (Tay et al., 2022) and OpenSora (Zheng et al., 2024), we propose a dynamic masking strategy to support video and audio conditioning, from the basic (I) *text-to-audio-video* (t2av) generation to (II) *audio-to-video* (a2v) generation, (III) *video-to-audio* (v2a) generation, (IV) *audio-image-to-video* (ai2v) generation, (V) *image-to-audio-video* (i2av) generation, and (VI) *audio-video-extension* (av_ext) generation. Note that the text condition still works in all kinds of conditional generations, which provides both coarse-grained global semantic embeddings and fine-grained spatio-temporal priors on targeted visual-sounding events. In particular, we unmask the specific video and frames to be conditioned on for X-conditional generation. During both model training and inference, unmasked frames will have timestep 0, while others remain the same (t):

- Unmasking all video/audio frames for v2a/a2v generation.
- Unmasking the first frame of the video (and all audio frames) for i2av (and ia2v) generation.
- Unmasking a preceding part of video and audio frames for va-extension.

We view the integration of all audio-video interactive generation tasks in one unified model as our future work.



Figure A9: Extensive JAVG cases on diverse event scenarios, visual styles, audio types, sounding subjects, and temporal compositions. Our JavisDiT achieves high-quality and text-consistency for single-modality generation and keeps a good video-audio synchronization.

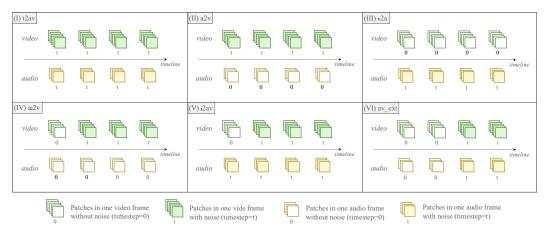


Figure A10: **Masking strategies for X-conditional generation.** The DiT architecture allows feasible conditional video-audio generation by replacing the noisy latent representation with reference videos and/or audios with specific strategies during training and inference.