

Tri-Ergon: Fine-Grained Video-to-Audio Generation with Multi-Modal Conditions and LUFS Control

Bingliang Li^{1,2*}, Fengyu Yang^{2*}, Yuxin Mao³, Qingwen Ye¹, Hongkai Chen^{1†}, Yiran Zhong^{4†}

¹vivo Mobile Communication Co., Ltd

²The Chinese University of Hong Kong, Shenzhen

³Northwestern Polytechnical University

⁴OpenNLPLab

{bingliangli,fengyuyang1}@link.cuhk.edu.cn, maoyuxin@mail.nwpu.edu.cn, qingwen.ye@vivo.com, allenhkchen@gmail.com, zhongyiran@gmail.com

Abstract

Video-to-audio (V2A) generation utilizes visual-only video features to produce realistic sounds that correspond to the scene. However, current V2A models often lack fine-grained control over the generated audio, especially in terms of loudness variation and the incorporation of multi-modal conditions. To overcome these limitations, we introduce Tri-Ergon, a diffusion-based V2A model that incorporates textual, auditory, and pixel-level visual prompts to enable detailed and semantically rich audio synthesis. Additionally, we introduce Loudness Units relative to Full Scale (LUFS) embedding, which allows for precise manual control of the loudness changes over time for individual audio channels, enabling our model to effectively address the intricate correlation of video and audio in real-world Foley workflows. Tri-Ergon is capable of creating 44.1 kHz high-fidelity stereo audio clips of varying lengths up to 60 seconds, which significantly outperforms existing state-of-the-art V2A methods that typically generate mono audio for a fixed duration.

Project website — <https://tri-ergon.github.io/Tri-Ergon/>

Introduction

When we experience visual events, we naturally expect to hear the corresponding sounds, enhancing our understanding and perception of the world. However, current video generation models mainly focus on creating visual content from text, often neglecting the integration of audio. The video-to-audio (V2A) task aims to fill this gap by generating semantically relevant and temporally synchronized audio from video frames. Due to its potential applications in movie automatic dubbing, video content production, game production, and other fields, the V2A task has attracted growing interest. As video generation models undergo continual refinement, the generation of fine-grained, high-fidelity, synchronized audio that aligns with prescribed loudness standards has emerged as a pivotal focal point for research.

*These authors contributed equally.

†Corresponding author.

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

Recent innovations in V2A generation have demonstrated considerable progress; however, current state-of-the-art methods still face limitations that hinder their practical application in industrial settings. Three primary issues have been identified in existing methods: i) The generated audio content may not be closely aligned with the corresponding visual or textual inputs due to single/double modality inputs. ii) Lack of precise control over audio attributes to effectively synchronize the audio and visual contents. iii) The absence of high fidelity, stereo, long-duration, and open-domain audio generation capability. Overcoming these challenges is essential for improving the effectiveness of V2A models in accurately generating audio alongside visual content.

In this paper, we posit that these challenges can be effectively mitigated through the utilization of *fine-grained audio control and generation*. Our concept of fine-grained encompasses two key aspects: i) the application of multi-modal conditions for audio generation, and ii) the implementation of a meticulous audio generation process that allows for exact control over loudness, audio quality, and duration. The former enables the model to leverage a broader range of information for audio generation, resulting in more coherent audio output. For instance, text information can provide semantic context for the audio, which is particularly beneficial for off-screen audio generation. Audio input serves as a style reference for the audio generation process. The latter aspect of fine-grained control ensures that the generated audio is closely aligned with the visual input and possesses enhanced quality.

In light of the aforementioned observation, we propose a novel framework called Tri-Ergon¹ to address the above limitations, as depicted in Figure 1. Tri-Ergon introduces three multi-modal conditioning and precise loudness control mechanisms. Specifically, the multi-modal conditioning utilizes textual, auditory, visual prompts to steer the audio generation process, enabling the creation of audio tracks that are semantically rich and contextually appropriate. An unbalanced multi-modal prompting module is introduced to en-

¹In memory of the sound-on-film system developed around 1919 (Lipton 2021) by three German inventors. Tri-Ergon means “the work of three”, derived from the Greek word $\tau\rho\acute{\iota}\alpha$, which corresponds to the three modalities in our work.

able Tri-Ergon to perform audio generation from any modality with aligned, semantic-rich features (Zhu et al. 2023) from video, text, and audio.

Tri-Ergon incorporates Loudness Units relative to Full Scale (LUFS) for precise loudness control mechanisms. LUFS provides a standardized metric closely aligned with human auditory perception, allowing fine-grained control over audio loudness throughout its entire duration. The momentary LUFS can be derived from a reference audio, crafted by an audio engineer, or predicted using our LUFS prediction module, which utilizes DINO-V2 to extract frame features from video and a transformer encoder to obtain the predicted LUFS. By integrating LUFS into the diffusion process, we enable precise adjustments to loudness variations over time, crucial for maintaining audio consistency and quality, especially in complex Foley workflows where synchronization between audio and visual elements is essential.

We have created the first comprehensive V2A dataset for training Tri-Ergon, featuring high-fidelity, long-duration, open-vocabulary multi-modal labeling, named MM-V2A. Through training on this dataset, Tri-Ergon can generate top-quality stereo audio at a sampling rate of 44.1 kHz. It is capable of producing audio clips of any length, with a maximum duration of 60 seconds. This represents a significant advancement over current V2A methods, which are generally limited to generating mono audio for a fixed duration.

The main contributions can be summarized as follows:

- We propose Tri-Ergon, a novel framework for video-to-audio (V2A) generation that introduces fine-grained multi-modal conditioning to generate semantically rich and contextually aligned audio tracks.
- We incorporate Loudness Units relative to Full Scale (LUFS) into the diffusion process, enabling precise control over loudness variations throughout the audio generation process.
- We introduce MM-V2A, a comprehensive V2A dataset with high-fidelity, long-duration, open-vocabulary multi-modal labeling for training our proposed Tri-Ergon.
- Our framework outperforms existing state-of-the-art methods, as demonstrated by both qualitative and quantitative results.

Related Work

Video-to-Audio Generation. Video-to-audio (V2A) generation focuses on synthesizing audio that aligns with the visual content of a video. SpecVQGAN (Iashin and Rahtu 2021) uses a Transformer-based autoregressive model with ResNet50 or RGB+Flow as the visual backbone to predict spectrograms from visual features, thereby generating audio. Im2Wav (Sheffer and Adi 2023) employs a dual-transformer model conditioned on CLIP features to produce sound directly from visual data. Difffoley (Luo et al. 2024) enhances the coherence of generated audio by using contrastive audio-visual pretraining to align audio and visual features. Some approaches also incorporate multi-modal and joint training encoders. V2A-Mapper (Wang et al. 2024) uses a lightweight mapper to convert CLIP embeddings of

videos into CLAP (Wu et al. 2023) embeddings, conditioning the audio generation on visual input. Additionally, Xing et al. (Xing et al. 2024) leverage an ImageBind-based latent aligner for conditional guidance in audio generation.

Multi-Modal Representation. Multi-modal representation typically starts with vision and language pretraining. CLIP (Radford et al. 2021) pioneered this approach by aligning images and texts using a large-scale dataset of 400 million samples, effectively bridging the gap between these two domains. Building on CLIP’s success, various models have extended this alignment to other modalities. CLIP4Clip (Luo et al. 2022) aligns videos with text, and CLAP aligns audio with text. Furthermore, more than two modalities can be aligned simultaneously with some modifications. ImageBind (Girdhar et al. 2023) broadens multi-modal alignment pretraining to include six different modalities. LanguageBind (Zhu et al. 2023), a language-based multi-modal pretraining approach, aligns all modalities with the language modality through contrastive learning, unifying them within a shared embedding space during the pretraining process.

High-Quality and Variable-Length Audio Generation.

In terms of quality, AudioLDM2 (Liu et al. 2024) demonstrates the ability to generate 48kHz mono audio, while Levy et al. (Levy et al. 2023) achieve the generation of 44.1kHz stereo music. Regarding the duration of generated audio, previous non-autoregressive models were constrained to producing music of up to 20 seconds (Parker et al. 2024). End-to-end and latent diffusion models extended this limit to 30 seconds (Levy et al. 2023; Huang et al. 2023). Stable Audio (Evans et al. 2024a) further advances the state-of-the-art by generating stereo signals up to 95 seconds in length. Our proposed method, Tri-Ergon, utilizes Diffusion Transformers to deliver high-fidelity stereo audio at 44.1kHz, with variable-length outputs of up to 60 seconds.

Methodology

Preliminary

Our framework is based on the latent diffusion model, which follows the standard formulation outlined in DDPM, and comprises a forward diffusion process and a backward reverse denoising process. Initially, a data sample $\mathbf{x} \sim p(\mathbf{x})$ undergoes processing by an autoencoder, consisting of an encoder \mathcal{E} and a decoder \mathcal{D} . The autoencoder projects \mathbf{x} into a latent variable \mathbf{z} via $\mathbf{z} = \mathcal{E}(\mathbf{x})$. Subsequently, the diffusion and denoising process takes place within the latent space.

The denoised latent variable is recovered to the input space by $\hat{\mathbf{x}} = \mathcal{D}(\hat{\mathbf{z}}_0)$.

Inspired by non-equilibrium thermodynamics, diffusion models (Ho, Jain, and Abbeel 2020) are a class of latent variable (z_1, \dots, z_T) models of the form $p_\theta(z_0) = \int p_\theta(z_{0:T}) dz_{1:T}$, where the latent variables are of the same dimensionality as the input data z_0 . The joint distribution $p_\theta(z_{0:T})$ is also called the *reverse process*:

$$p_\theta(z_{0:T}) = p_\theta(z_T) \prod_{t=1}^T p_\theta(z_{t-1}|z_t), \quad (1)$$

$$p_\theta(z_{t-1}|z_t) = \mathcal{N}(z_{t-1}; \mu_\theta(z_t, t), \Sigma_\theta(z_t, t)). \quad (2)$$

Here, μ_θ and Σ_θ are determined through a denoiser network $\epsilon_\theta(z_t, t)$, typically structured as a UNet or a transformer (Ronneberger, Fischer, and Brox 2015; Peebles and Xie 2023; Mao et al. 2023).

The approximate posterior $q(z_{1:T}|z_0)$ is called the *forward process*, which is fixed to a Markov chain that gradually adds noise according to a predefined noise scheduler $\beta_{1:T}$:

$$q(z_{1:T}|z_0) = \prod_{t=1}^T q(z_t|z_{t-1}), \quad (3)$$

$$q(z_t|z_{t-1}) = \mathcal{N}(z_t; \sqrt{1 - \beta_t}z_{t-1}, \beta_t \mathbf{I}). \quad (4)$$

The training is performed by minimizing a variational bound on the negative log-likelihood, and the final training objective for θ is a noise estimation loss with a conditional variable \mathbf{E} . It can be formulated as:

$$\mathcal{L}_{\text{diffusion}}(\theta) := \mathbb{E}_{\mathbf{z}, \mathbf{c}, \epsilon \sim \mathcal{N}(0, \mathbf{I}), t} [\|\epsilon - \epsilon_\theta(\mathbf{z}_t, \mathbf{E}, t)\|^2]. \quad (5)$$

Conditioning

The conditional variable \mathbf{E} , as highlighted in the diffusion loss function, plays a crucial role in guiding the generation.

Unbalanced Multi-Modal Prompting Module. We introduce an unbalanced multi-modal prompting module as a conditioning mechanism, leveraging LanguageBind as the multi-modal encoder. We denote the LanguageBind encoder as τ_θ , which encodes the input condition y (e.g., language, audio, and video) into an embedding $\tau_\theta(y) \in \mathbb{R}^{M \times d_\tau}$, represented as \mathbf{e}_L , \mathbf{e}_A , and \mathbf{e}_V , respectively.

To enhance the ability to generate audio based on various combinations of conditions, we propose a random masking mechanism. During training, the embeddings \mathbf{e}_L , \mathbf{e}_A , and \mathbf{e}_V are randomly masked, leaving between 0 and 3 conditions active. Additionally, a task ID x_{task} is assigned based on the combination of available conditions (e.g., $\{\mathbf{e}_L, \text{mask}, \mathbf{e}_V\}$ is assigned the integer 5). This task ID is then encoded into a task embedding \mathbf{e}_{task} .

The final multi-modal feature representation $\mathbf{E}_M \in \mathbb{R}^{4M \times d_\tau}$ is constructed by stacking \mathbf{e}_{task} , \mathbf{e}_L , \mathbf{e}_A , and \mathbf{e}_V . This representation contains aligned, semantically rich information, enabling our model to generate audio based on prompts from up to three modalities.

Stereo Loudness Control with LUFS. Loudness Units relative to Full Scale (LUFS) is a standardized metric for audio loudness that accounts for human perception and signal intensity, widely used for audio normalization in media. LUFS can be measured in three ways: momentary (400 ms window), short-term (3-second moving average), and integrated (mean of short-term LUFS). We utilize a customized momentary LUFS measurement with a window size of $\frac{1}{6}$ second.

For a given audio signal x_a with duration L_a , the LUFS information is extracted separately for the left and right channels, resulting in $\mathbf{e}'_{\text{left}}$ and $\mathbf{e}'_{\text{right}}$, both with a length of $6L_a$. These values are clipped to the range of -70 (extremely

quiet) to 0 (extremely loud) and normalized to a range of 0 to 1. The normalized values are then padded and interpolated to match the dimensions of \mathbf{E}_M as $\{\mathbf{e}_{\text{left}}, \mathbf{e}_{\text{right}}\}$, forming the LUFS embedding $\mathbf{E}_L \in \mathbb{R}^{2M \times d_\tau}$.

We propose three methods for generating \mathbf{E}_L : (1) directly measuring a reference audio using a LUFS meter, (2) manually crafting $\mathbf{e}'_{\text{left}}$ and $\mathbf{e}'_{\text{right}}$ based on the video, and (3) utilizing our proposed LUFS prediction module, which can also be modified by the user for enhanced control.

The LUFS prediction module employs the DINO-V2 model (Oquab et al. 2023) D to extract image features from a video at 6 FPS, which are then processed by a transformer encoder T_{LUFS} to predict the LUFS embedding \mathbf{E}_L . This module is trained using MSE loss and supports videos up to 60 seconds in length.

Timing Embeddings for Variable Audio Length. The current V2A model often utilizes video encoders that can only process a fixed number of frames, and their audio VAE usually operates in the Mel spectrogram space. Due to these limitations, they can only generate audio with a limited and fixed duration, usually 10 seconds. To enable generating audio with variable length, we adopt the timing embeddings in (Evans et al. 2024a).

When extracting a segment of audio from our training data, we determine two properties: the starting second of the segment, x_{start} , and the total duration of the original audio file, x_{total} . For instance, for a given audio, if we extract a 60-second segment from a 180-second audio file, starting at the 21.5-second mark, then x_{start} would be 21.5, and x_{total} would be 180. These values are subsequently converted into learned timing embeddings $\mathbf{E}_T \in \mathbb{R}^{2M \times d_\tau}$ and concatenated along the sequence dimension with \mathbf{E}_M and \mathbf{E}_L as \mathbf{E} before being input into the DiT’s cross-attention layers.

During training, we randomly cut the audio of a video to perform augmentation. During inference, x_{start} is always set to 0, and x_{total} is set to the duration of the video. This enables the model to generate audio with variable length to exactly match the duration of the video.

Multi-Conditioned Latent Diffusion

Latent Autoencoder. We adopt the Variational Autoencoder (VAE) \mathcal{E}_θ in (Evans et al. 2024b) to achieve latent space encoding and decoding. Specifically, the encoding process aims to compress the raw high-fidelity stereo 44.1kHz audio x_a into the lower-dimensional latent code $z_0 = \mathcal{E}_\theta(x_a) \in \mathbb{R}^{C \times \frac{L}{r}}$, where C is the dimension of the encoded latent, L is the length of the audio signal, r is the downsampling factor, $\frac{r}{2 \times C}$ is the compression ratio. We use a combination of multiple datasets to train the VAE, with specific details provided in the experiment section.

Denoising Network Architecture. We utilize the Diffusion Transformer (DiT) (Peebles and Xie 2023) as our denoising network ϵ_θ . To effectively integrate conditioning into the diffusion process, cross-attention mechanisms are incorporated within the transformer’s layers. These mechanisms allow the model to incorporate additional contextual information, such as text, visual, or auditory cues.

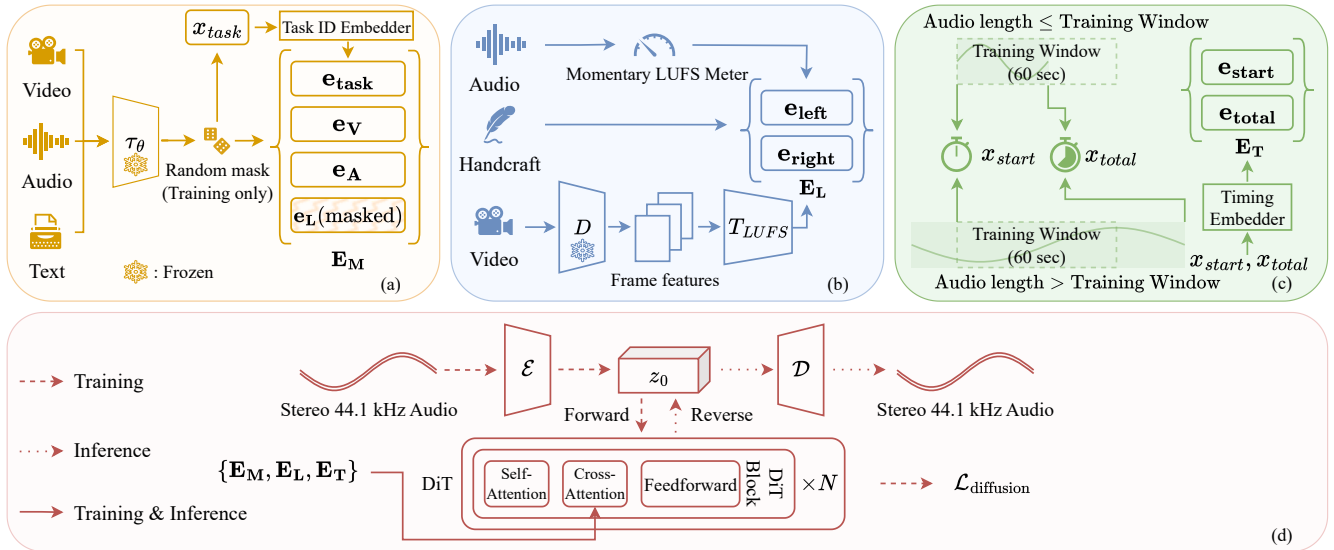


Figure 1: **Tri-Ergon Overview.** (a) Unbalanced multi-modal prompting module, which takes any combination of prompts. (b) Stereo loudness control with LUFS. (c) Timing embeddings for variable audio length. (d) The DiT backbone.

Denoising with LUFS Information. For video to audio generation, the objective is to generate synchronized, high-fidelity stereo audio x_a given a multi-modal feature representation \mathbf{E}_M .

Latent Diffusion Model Objective. In the context of multi-conditioned latent diffusion, the Latent Diffusion Model (LDM) aims to reverse the process of noise addition, which transforms the original data distribution into a standard Gaussian distribution over a series of timesteps. The goal of the LDM is to recover the original latent representation by minimizing the discrepancy between the actual noise added during the forward process and the model’s predicted noise.

Given the latent variable z_t at timestep t , the model predicts the noise $\epsilon_\theta(z_t, t, \{\mathbf{E}_M, \mathbf{E}_L, \mathbf{E}_T\})$, where $\mathbf{E}_M, \mathbf{E}_L, \mathbf{E}_T$ represent the multi-modal feature embeddings that provide additional contextual information. The model is trained to minimize the denoising objective:

$$\mathcal{L}_{LDM} = \mathbb{E}_{z_0, t, \epsilon} \|\epsilon - \epsilon_\theta(z_t, t, \{\mathbf{E}_M, \mathbf{E}_L, \mathbf{E}_T\})\|_2^2 \quad (6)$$

This objective function, based on L_2 loss (mean squared error), ensures that the model learns to effectively reverse the diffusion process, gradually removing noise from the latent code and reconstructing a clean representation.

Reverse Process and Generation. During generation, the LDM starts with a latent code sampled from a standard Gaussian distribution and iteratively applies the learned denoising steps in reverse. Each step refines the latent code, guided by the multi-modal inputs, until a denoised latent representation is achieved. This refined latent code is then decoded by the VAE into the final high-fidelity output, such as stereo audio synchronized with video content. The integration of LUFS information and other contextual features ensures that the generated audio is both high-quality and consistent with the input modalities.

Experiments

Model Configuration and Architecture Details

VAE. We first train the audio VAE using automatic mixed precision for 1M steps with an effective batch size of 128 on 8 A100 GPUs. Similar to (Evans et al. 2024a), we use the multi-resolution sum and difference STFT loss (Steinmetz et al. 2021) with A-weighting (Fletcher and Munson 1933), which is designed for stereo signals. We also utilize adversarial and feature matching losses with a modified multi-scale STFT EnCodec based discriminator (Défossez et al. 2022) that is adapted to process stereo audio. The losses are weighted as follows: 1.0 for spectral losses, 0.1 for adversarial losses, 5.0 for feature matching loss, and 1×10^{-1} for KL loss. The total number of parameters for our VAE is 157M.

Multi-Conditioned Latent Diffusion. We define two model variants, Tri-Ergon-S and Tri-Ergon-L, based on their scale and number of parameters. For the Tri-Ergon-S model, we employ a DiT architecture comprising 20 DiT blocks as the diffusion backbone, with 24 attention heads, an embedding dimension of 1536, and a total of 881M parameters. This model is trained for 7.2×10^5 steps. Additionally, for Tri-Ergon-L, we utilize 24 DiT blocks while maintaining the same block configuration. It contains 1.1B parameters and is trained for 2.2×10^5 steps. Both models are trained using the v-objective (Salimans and Ho 2022), with a cosine noise schedule and continuous denoising timesteps. The training process was conducted on 32 A100 GPUs, with an effective batch size of 256. For the conditioning set E_m , we apply a 30% probability of independently dropping the embeddings $\mathbf{e}_L, \mathbf{e}_A$, and \mathbf{e}_V . Additionally, a 10% dropout rate is applied to the overall conditioning set E to enable classifier-free guidance.

LUFS Prediction Module. We train two separate LUFS prediction modules for Tri-Ergon-S and Tri-Ergon-L, re-

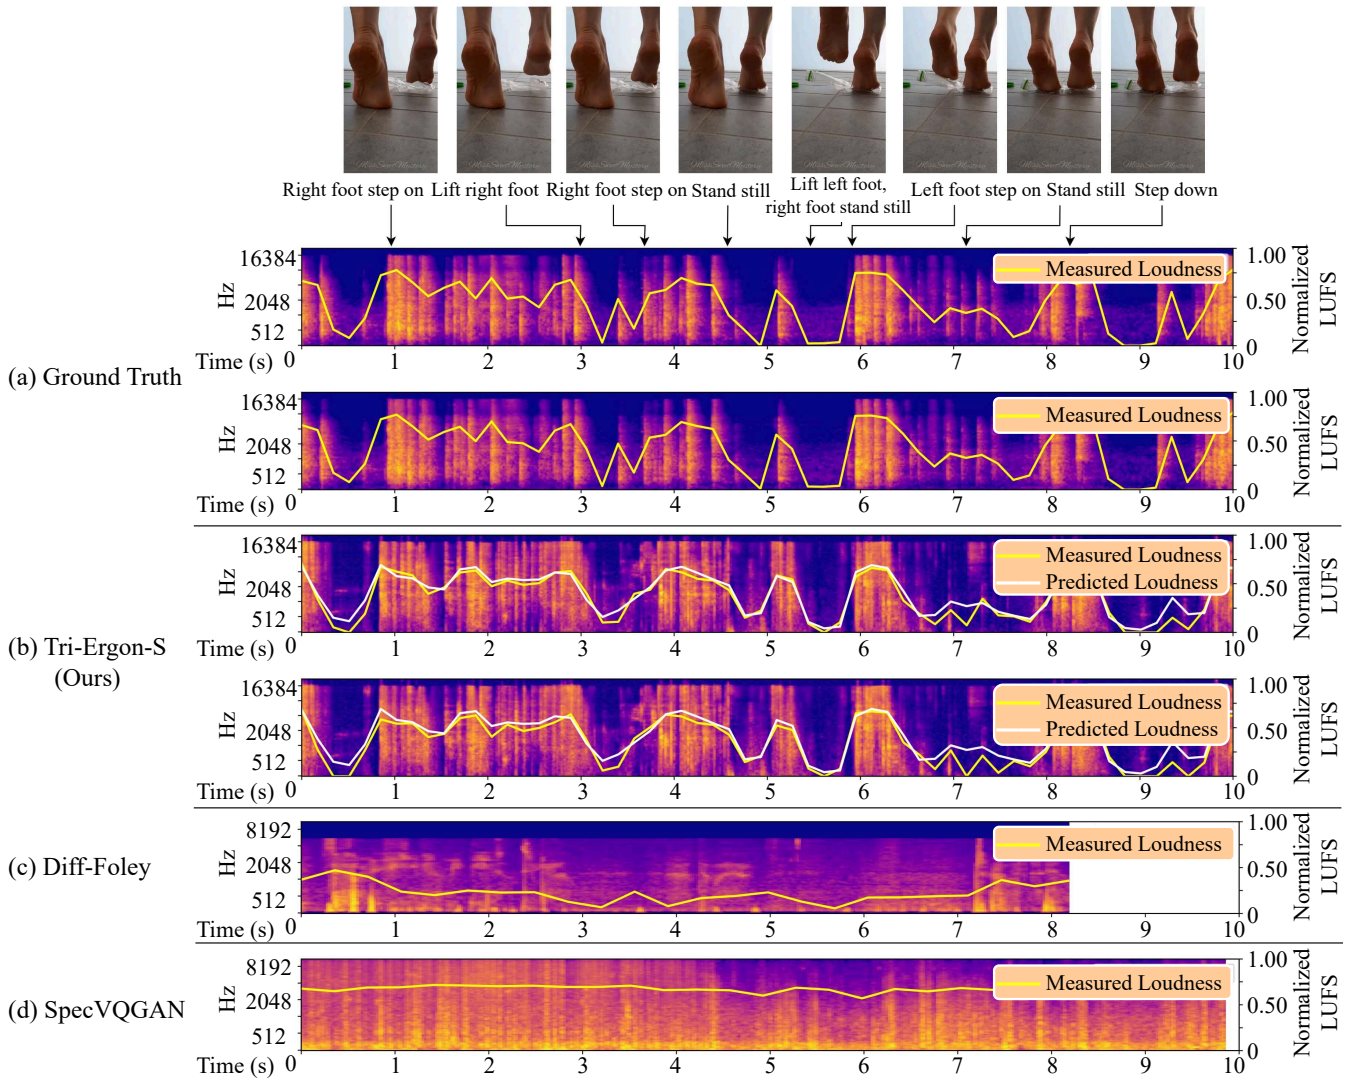


Figure 2: Video-to-Audio generation results on VGGSound. Given a video of a person alternately stepping on a plastic bottle with both feet, our model can successfully detect the loudness variation and generate highly synchronized audio. Note that the predicted and the actual LUFS for the generated audio are closely match, demonstrate the controllability Tri-Ergon provides.

ferred to as T_{LUFS-S} and T_{LUFS-L} , respectively. The T_{LUFS-S} module is based on a transformer encoder architecture with 12 layers, 8 attention heads, and an embedding dimension of 768. In contrast, T_{LUFS-L} features a more complex architecture with 14 layers and 12 attention heads, while maintaining the same embedding dimension of 768. T_{LUFS-S} is trained on the VGGSound dataset for 10 epochs, whereas T_{LUFS-L} is trained on our proposed dataset for 30 epochs. Both modules are trained on 8 A100 GPUs with a batch size of 128.

Evaluation Metrics

SpecVQGAN (Iashin and Rahtu 2021) relies on metrics such as Mean MKL and Melception-based FID for evaluating relevance and fidelity, using audio that is resampled to 22.05 kHz, mono, and fixed at 10 seconds. However, our

work focuses on generating long-form full-band stereo signals, for which we utilize FD_{openl3} , KL_{passt} (Evans et al. 2024a), and AV-Align (Yariv et al. 2024) as quantitative metrics. The FD_{openl3} metric evaluates the similarity between generated and reference audio by projecting them into the Openl3 feature space and calculating the Fréchet Distance, making it suitable for assessing stereo signals up to 44.1 kHz. The KL_{passt} metric leverages the PaSST model to compute the Kullback-Leibler (KL) divergence between the label probabilities of generated and reference audio, thereby assessing their semantic similarity by analyzing long-form audio segments. Additionally, AV-Align is utilized to measure the alignment between audio and visual modalities by detecting and comparing energy peaks in both.

Dataset	Video	Text	Audio	Video Duration	Unified Data
VGGSound	199K	199K	199K	10 sec	199K
VALOR-32K	25.3K	25.3K	25.3K	10 sec	25.3K
VIDAL-10M	69.7K	69.7K	69.7K	< 20 sec	69.7K
BBC Sound Effect Library	-	31.1K	31.1K	-	31.1K
FAVDBench	7.5K	7.5K	7.5K	5 sec ~ 10 sec	7.5K
HD-VILA-100M	44.6K	44.6K	44.6K	< 60 sec	44.6K
AudioSet	128.1K	128.1K	128.1K	10 sec	128.1K
MM-V2A	474.2K	505.3K	505.3K	5 sec ~ 60 sec	505.3K

Table 1: Statistics for MM-V2A, which consolidates seven existing datasets, are presented in the table. The table numerically displays the count of data integrated from each source dataset.

	Channels/sr	Output length	FD_{openl3} ↓	KL_{passt} ↓	AV-Align ↑
Training data (upper bound)	2/44.1kHz	10 sec	16.87	-	-
Autoencoded training data	2/44.1kHz	10 sec	28.34	-	-
SpecVQGAN	1/22.05kHz	9.8 sec	183.07	3.12	0.204
Diff-Foley (Classifier-Free Guidance)	1/16kHz	8 sec	249.15	3.07	0.177
Diff-Foley (Double Guidance)	1/16kHz	8 sec	242.90	3.03	0.182
Tri-Ergon-S (Ours)	2/44.1kHz	10 sec	144.57	2.63	0.229
Tri-Ergon-L (Ours)	2/44.1kHz	Up to 60 sec	113.21	1.82	0.231

Table 2: Video-to-Audio generation results on VGGSound dataset. Our model achieve impressive results on both audio quality and temporal synchronization.

Datasets

The dataset used for training the diffusion model is systematically partitioned into two distinct segments: one dedicated to the training of the Tri-Ergon-S model and the other to the training of the Tri-Ergon-L model. The Tri-Ergon-S model is trained and evaluated exclusively on the VGG-Sound dataset (Chen et al. 2020), which consists of 200,000 video clips across 309 audio classes. Each 10-second clip features visual frames of the sound-producing object, paired with an audio track that corresponds to the visual content.

In contrast, the Tri-Ergon-L model is trained on a newly proposed and more comprehensive dataset, termed MM-V2A. This dataset is designed to extend beyond VGG-Sound by incorporating additional datasets with multi-modal annotations, including VALOR-32K (Chen et al. 2024), VIDAL-10M (Zhu et al. 2023), the BBC Sound Effect Library², FAVDBench (Shen et al. 2023), HD-VILA-100M (Xue et al. 2022), and AudioSet (Gemmeke et al. 2017). As detailed in Table 1, MM-V2A aggregates 69,749 videos from VIDAL-10M, 44,637 videos from HD-VILA-100M, 128,079 videos from AudioSet, and the complete sets of VGG-Sound, VALOR-32K, BBC Sound Effects, and FAVDBench.

For the VAE model, we collect several dataset in various domain to achieve optimal reconstruction results, including general audio: AudioSet, FSD50K³, FAVDBench, Urban-Sound8K (Salamon, Jacoby, and Bello 2014), VGGSound, VIDAL-10M, HD-VILA-100M, sound effect: BBC Sound Effect Library, music: MTG-Jamendo (Bogdanov et al. 2019), and human voice: Common Voice Corpus 1 (En-

glish)⁴, results in total number of 996,033 audio files, we refer the above collection of datasets as the VAE ensemble dataset.

Baseline

We adopt two advanced V2A models as baselines: 1) SpecVQGAN (Iashin and Rahtu 2021), a transformer-based autoregressive model that generates spectrogram VQVAE indices from visual features; and 2) Diff-Foley (Luo et al. 2024), a powerful V2A model based on latent diffusion. For SpecVQGAN, we evaluate its best-performing variant using Resnet+Flow as visual input conditions. For Diff-Foley, we assess its performance both with and without classifier guidance to analyze the influence of its sophisticated external alignment mechanism.

Video-to-Audio Generation Results

We present the quantitative results on VGGSound test set in Table 2. For both Tri-Ergon-S and Tri-Ergon-L, we only use the video embedding e_V for fair compression. As we can observe that, our proposed Tri-Ergon-S outperforms previous methods significantly on both audio quality and temporal synchronization. Benefited from our newly proposed dataset, Tri-Ergon-L further pushes the boundaries of video-to-audio generation performance.

Figure 2 presents the qualitative results for Tri-Ergon-S using e_V , e_L , and E_L as the composed conditions, as predicted by T_{LIFS-S} . The T_{LIFS-S} successfully predicted the momentary LIFS change over time, guiding the model to generate synchronous audio. Additionally, the measured

²<https://sound-effects.bbcrewind.co.uk>

³<https://freesound.org/>

⁴<https://commonvoice.mozilla.org/en/datasets>

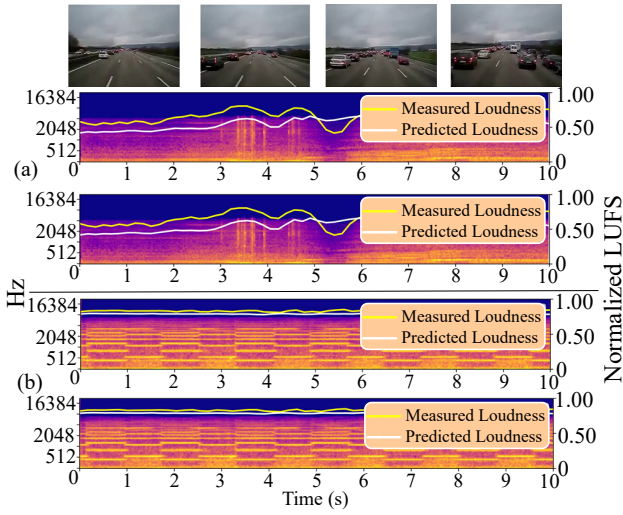


Figure 3: Fine-grained Condition Analysis. We present a scenario featuring an ambulance with no visible cue on a roadway. (a) Generated audio with visual information only. (b) Generated audio after adjusting the LUFS embedding along with additional text prompt.

LUFS of the generated audio closely aligns with the input \mathbf{E}_L condition, demonstrating the fine-grained controllability of the proposed Tri-Ergon.

Fine-Grained Condition Analysis

Our proposed model captures fine-grained multimodal information to enable more precise control. As illustrated in Figure 3, the absence of an ambulance in the scene results in the model’s LUFS prediction lacking any visual information about an ambulance, leading to the generation of traffic sounds on the street, as shown in Figure 3(a). However, by manually adjusting the LUFS embedding and incorporating a text prompt specifying the ambulance sound “Sound of loud ambulance siren”, the model can generate audio that includes the sound of an ambulance, as shown in Figure 3(b).

Ablation Study

Compose Different Conditions. As discussed above, different condition types provide diverse guidance for the final generation results. In this section, we study the impact of different combinations of e_V , e_L and e_A on the VGGSound dataset, as shown in Table 3. We also compare the results between using \mathbf{E}_L predicted by T_{LUFS} and the ground truth of the testing audio, the latter can be seen as handcrafted momentary LUFS by audio engineers. We observe the following: 1) Additional text embedding will improve the audio semantic relevance with the video, but have marginal contribution to the audio-video synchronization. 2) Using the ground truth e_A significantly improves the semantic similarity, demonstrating the one-shot flexibility of Tri-Ergon in audio generation. 3) Utilizing ground truth \mathbf{E}_L further enhances both audio quality and alignment, demonstrating the versatility and flexibility of our model, which allows for

e_V	e_L	\bar{e}_A	$\hat{\mathbf{E}}_L$	$\bar{\mathbf{E}}_L$	$FD_{openl3} \downarrow$	$KL_{passt} \downarrow$	$AV\text{-Align} \uparrow$
✓	✗	✗	✓	✗	144.57	2.63	0.229
✓	✓	✗	✓	✗	141.47	2.61	0.229
✓	✗	✓	✓	✗	103.82	1.83	0.236
✓	✓	✓	✓	✗	96.23	1.47	0.240
✓	✗	✗	✗	✓	89.92	1.79	0.232
✓	✓	✗	✗	✓	89.65	1.608	0.234
✓	✗	✓	✗	✓	82.29	1.06	0.238
✓	✓	✓	✗	✓	78.19	0.97	0.242

Table 3: The effect of using different combinations of conditions with Tri-Ergon-S on VGGSound. We use \bar{e}_A denote ground truth e_A , \mathbf{E}_L denote $\mathbf{E}_L(T_{LUFS})$, and $\bar{\mathbf{E}}_L$ denote $\mathbf{E}_L(GT)$.

manual intervention when the results do not meet the user’s requirements.

CFG Scale & Number of Inference Steps analysis. As shown in Figure 4, our model achieves optimal performance with a CFG scale ω set to 7 and an inference step count of 100. Further increasing the number of inference steps beyond 100 results in only marginal improvements in audio quality. Therefore, for balancing computational efficiency and audio fidelity, 100 inference steps with a CFG scale of 7 is the recommended configuration.




Figure 4: CFG Scale & Number of Inference Steps. Our model reach the best performance with $\omega = 7$, inference steps of 100.

Conclusion

The Tri-Ergon model is a significant advancement in V2A synthesis, addressing key limitations by integrating multimodal inputs and fine-grained loudness control. It incorporates textual, auditory, and visual prompts, achieving a richer, more nuanced audio generation aligned with the video context. LUFS embedding enhances its utility in professional settings, particularly in Foley workflows. It can produce high-fidelity stereo audio up to 60 seconds, surpassing the current state-of-the-art and opening new possibilities for detailed sound synthesis in various applications. These advancements have transformative potentials for film, gaming, virtual reality, and automated video processing industries. Moving forward, refining these techniques and exploring their applications in different contexts will be crucial, such as audible-video generation (Mao et al. 2024) or techniques for efficient generation (Qin et al. 2024).

References

- Bogdanov, D.; Won, M.; Tovstogan, P.; Porter, A.; and Serra, X. 2019. The MTG-Jamendo Dataset for Automatic Music Tagging. In *Machine Learning for Music Discovery Workshop, International Conference on Machine Learning (ICML 2019)*. Long Beach, CA, United States.
- Chen, H.; Xie, W.; Vedaldi, A.; and Zisserman, A. 2020. Vggsound: A large-scale audio-visual dataset. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 721–725. IEEE.
- Chen, S.; Li, H.; Wang, Q.; Zhao, Z.; Sun, M.; Zhu, X.; and Liu, J. 2024. Vast: A vision-audio-subtitle-text omnimodality foundation model and dataset. *Advances in Neural Information Processing Systems*, 36.
- Défossez, A.; Copet, J.; Synnaeve, G.; and Adi, Y. 2022. High fidelity neural audio compression. *arXiv preprint arXiv:2210.13438*.
- Evans, Z.; Carr, C.; Taylor, J.; Hawley, S. H.; and Pons, J. 2024a. Fast timing-conditioned latent audio diffusion. *arXiv preprint arXiv:2402.04825*.
- Evans, Z.; Parker, J. D.; Carr, C.; Zukowski, Z.; Taylor, J.; and Pons, J. 2024b. Stable Audio Open. *arXiv preprint arXiv:2407.14358*.
- Fletcher, H.; and Munson, W. A. 1933. Loudness, its definition, measurement and calculation. *Bell System Technical Journal*, 12(4): 377–430.
- Gemmeke, J. F.; Ellis, D. P.; Freedman, D.; Jansen, A.; Lawrence, W.; Moore, R. C.; Plakal, M.; and Ritter, M. 2017. Audio set: An ontology and human-labeled dataset for audio events. In *2017 IEEE international conference on acoustics, speech and signal processing (ICASSP)*, 776–780. IEEE.
- Girdhar, R.; El-Nouby, A.; Liu, Z.; Singh, M.; Alwala, K. V.; Joulin, A.; and Misra, I. 2023. Imagebind: One embedding space to bind them all. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 15180–15190.
- Ho, J.; Jain, A.; and Abbeel, P. 2020. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33: 6840–6851.
- Huang, Q.; Park, D. S.; Wang, T.; Denk, T. I.; Ly, A.; Chen, N.; Zhang, Z.; Zhang, Z.; Yu, J.; Frank, C.; et al. 2023. Noise2music: Text-conditioned music generation with diffusion models. *arXiv preprint arXiv:2302.03917*.
- Iashin, V.; and Rahtu, E. 2021. Taming visually guided sound generation. *arXiv preprint arXiv:2110.08791*.
- Levy, M.; Di Giorgi, B.; Weers, F.; Katharopoulos, A.; and Nickson, T. 2023. Controllable music production with diffusion models and guidance gradients. *arXiv preprint arXiv:2311.00613*.
- Lipton, L. 2021. *The cinema in flux: the evolution of motion picture technology from the magic lantern to the digital era*. Springer.
- Liu, H.; Yuan, Y.; Liu, X.; Mei, X.; Kong, Q.; Tian, Q.; Wang, Y.; Wang, W.; Wang, Y.; and Plumbley, M. D. 2024. Audioldm 2: Learning holistic audio generation with self-supervised pretraining. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*.
- Luo, H.; Ji, L.; Zhong, M.; Chen, Y.; Lei, W.; Duan, N.; and Li, T. 2022. Clip4clip: An empirical study of clip for end to end video clip retrieval and captioning. *Neurocomputing*, 508: 293–304.
- Luo, S.; Yan, C.; Hu, C.; and Zhao, H. 2024. Diff-foley: Synchronized video-to-audio synthesis with latent diffusion models. *Advances in Neural Information Processing Systems*, 36.
- Mao, Y.; Shen, X.; Zhang, J.; Qin, Z.; Zhou, J.; Xiang, M.; Zhong, Y.; and Dai, Y. 2024. TAVGBench: Benchmarking text to audible-video generation. In *Proceedings of the 32nd ACM International Conference on Multimedia*, 6607–6616.
- Mao, Y.; Zhang, J.; Xiang, M.; Lv, Y.; Zhong, Y.; and Dai, Y. 2023. Contrastive conditional latent diffusion for audio-visual segmentation. *arXiv preprint arXiv:2307.16579*.
- Oquab, M.; Darcet, T.; Moutakanni, T.; Vo, H.; Szafraniec, M.; Khalidov, V.; Fernandez, P.; Haziza, D.; Massa, F.; El-Nouby, A.; et al. 2023. Dinov2: Learning robust visual features without supervision. *arXiv preprint arXiv:2304.07193*.
- Parker, J. D.; Spijkervet, J.; Kosta, K.; Yesiler, F.; Kuznetsov, B.; Wang, J.-C.; Avent, M.; Chen, J.; and Le, D. 2024. Stem-Gen: A music generation model that listens. In *ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 1116–1120. IEEE.
- Peebles, W.; and Xie, S. 2023. Scalable diffusion models with transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 4195–4205.
- Qin, Z.; Mao, Y.; Shen, X.; Li, D.; Zhang, J.; Dai, Y.; and Zhong, Y. 2024. You Only Scan Once: Efficient Multi-dimension Sequential Modeling with LightNet. *arXiv preprint arXiv:2405.21022*.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, 8748–8763. PMLR.
- Ronneberger, O.; Fischer, P.; and Brox, T. 2015. U-net: Convolutional networks for biomedical image segmentation. 234–241. Springer.
- Salamon, J.; Jacoby, C.; and Bello, J. P. 2014. A dataset and taxonomy for urban sound research. In *Proceedings of the 22nd ACM international conference on Multimedia*, 1041–1044.
- Salimans, T.; and Ho, J. 2022. Progressive distillation for fast sampling of diffusion models. *arXiv preprint arXiv:2202.00512*.
- Sheffer, R.; and Adi, Y. 2023. I hear your true colors: Image guided audio generation. In *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 1–5. IEEE.
- Shen, X.; Li, D.; Zhou, J.; Qin, Z.; He, B.; Han, X.; Li, A.; Dai, Y.; Kong, L.; Wang, M.; et al. 2023. Fine-grained audible video description. In *Proceedings of the IEEE/CVF*

Conference on Computer Vision and Pattern Recognition, 10585–10596.

Steinmetz, C. J.; Pons, J.; Pascual, S.; and Serra, J. 2021. Automatic multitrack mixing with a differentiable mixing console of neural audio effects. In *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 71–75. IEEE.

Wang, H.; Ma, J.; Pascual, S.; Cartwright, R.; and Cai, W. 2024. V2a-mapper: A lightweight solution for vision-to-audio generation by connecting foundation models. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 15492–15501.

Wu, Y.; Chen, K.; Zhang, T.; Hui, Y.; Berg-Kirkpatrick, T.; and Dubnov, S. 2023. Large-scale contrastive language-audio pretraining with feature fusion and keyword-to-caption augmentation. In *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 1–5. IEEE.

Xing, Y.; He, Y.; Tian, Z.; Wang, X.; and Chen, Q. 2024. Seeing and hearing: Open-domain visual-audio generation with diffusion latent aligners. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 7151–7161.

Xue, H.; Hang, T.; Zeng, Y.; Sun, Y.; Liu, B.; Yang, H.; Fu, J.; and Guo, B. 2022. Advancing high-resolution video-language representation with large-scale video transcriptions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 5036–5045.

Yariv, G.; Gat, I.; Benaim, S.; Wolf, L.; Schwartz, I.; and Adi, Y. 2024. Diverse and aligned audio-to-video generation via text-to-video model adaptation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 6639–6647.

Zhu, B.; Lin, B.; Ning, M.; Yan, Y.; Cui, J.; Wang, H.; Pang, Y.; Jiang, W.; Zhang, J.; Li, Z.; et al. 2023. Language-bind: Extending video-language pretraining to n-modality by language-based semantic alignment. *arXiv preprint arXiv:2310.01852*.