

Video-Guided Foley Sound Generation with Multimodal Controls

Ziyang Chen^{1,2*} Prem Seetharaman² Bryan Russell² Oriol Nieto²
 David Bourgin² Andrew Owens¹ Justin Salamon²

¹University of Michigan ²Adobe Research

<https://ificl.github.io/MultiFoley/>

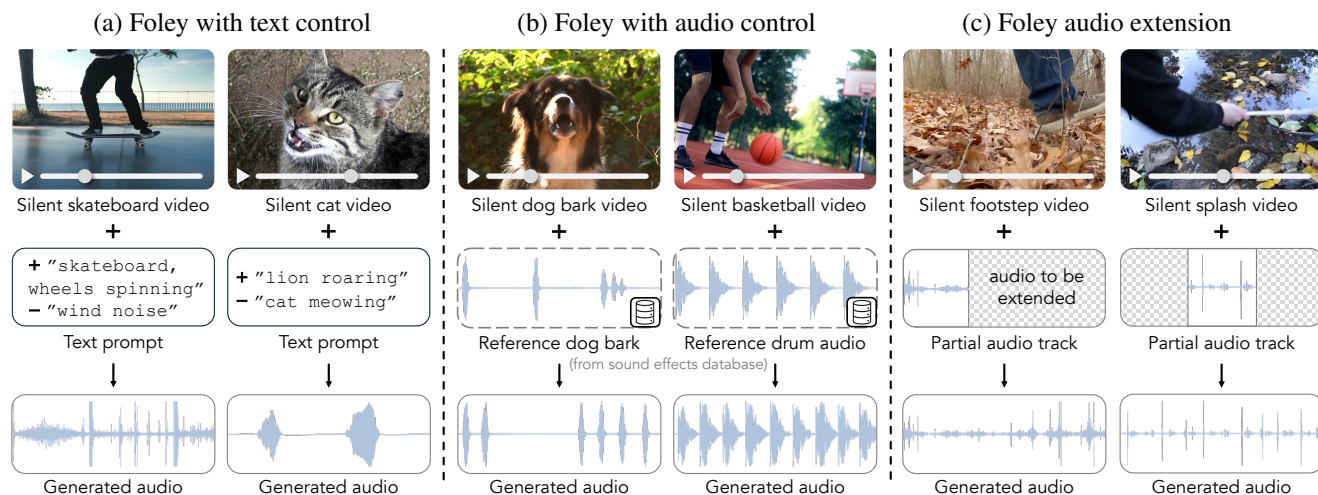


Figure 1. **MultiFoley for video-guided sound generation with multimodal controls.** We generate Foley sounds for silent videos with various control signals to shape their audio. (a) Text prompts, both positive and negative, guide synchronized Foley generation. (b) Reference audio from sound libraries defines the customized audio style. (c) A partial audio track is extended to produce a complete Foley sound. **We encourage the reader to watch and listen to the results in our [website](https://ificl.github.io/MultiFoley/).**

Abstract

Generating sound effects for videos often requires creating artistic sound effects that diverge significantly from real-life sources and flexible control in the sound design. To address this problem, we introduce MultiFoley, a model designed for video-guided sound generation that supports multimodal conditioning through text, audio, and video. Given a silent video and a text prompt, MultiFoley allows users to create clean sounds (e.g., skateboard wheels spinning without wind noise) or more whimsical sounds (e.g., making a lion’s roar sound like a cat’s meow). MultiFoley also allows users to choose reference audio from sound effects (SFX) libraries or partial videos for conditioning. A key novelty of our model lies in its joint training on both internet video datasets with low-quality audio and professional SFX recordings, enabling high-quality, full-bandwidth (48kHz) audio generation. Through automated evaluations and human studies, we demonstrate that MultiFoley successfully generates synchronized high-quality sounds

* Work done during an internship at Adobe.

across varied conditional inputs and outperforms existing methods. Please see our project page for video results: <https://ificl.github.io/MultiFoley/>

1. Introduction

Sound designers, often create sound effects from different sources that may not resemble the original video, such as everyday objects or pre-recorded sound libraries, to create a soundtrack for a video that synchronizes with on-screen actions. For example, they use crinkling paper for a warm fire, a coconut shell for horse hooves, or the cracking of celery for breaking bones. This process is known as *Foley* [1]. Instead of replicating the true sounds, sound designers aim to achieve an artistic effect that enhances the viewers’ experience.

Recent works [63, 66, 87, 90, 92] mainly formulated the Foley problem as a video-to-audio generation task, generating sound effects that align temporally and semantically with the video. However, those approaches restrict audio content and lack the control designers need, as they often create

sounds that diverge from real life in many different ways. Current systems only offer limited controls (*e.g.*, conditional video examples [24] or language [94]) and face issues with audio quality and synchronization.

To give users more creative control over sound design, we propose MULTIFOLEY, a video-guided Foley sound generation framework that integrates multimodal controls, using text, audio, and video conditional signals. Our model provides sound designers fine-grained control over how the audio sounds while easing the burden of synchronizing audio to videos. By jointly training across audio, video, and text modalities, MULTIFOLEY enables flexible control and various Foley applications. Users can customize audio content through text prompts – whether or not they match the visuals – while maintaining synchronization. For example, users can generate clean sound effects by removing unwanted elements like wind noise using text prompts or they can replace a cat meow with a lion roar, as shown in Fig. 1. Beyond text, our system can also accept different types of audio and audio-visual conditioning. Users can generate desired sound effects from reference audio in the sound effects (SFX) library or extend the soundtrack from a portion of a video.

One of the key challenges is that internet video soundtracks are often poorly aligned with the visual content (*e.g.*, irrelevant audio) and suffer from low quality (*e.g.*, noisy audio and limited bandwidth). To address this, we jointly train on high-quality SFX libraries (audio-text pairs) alongside internet videos, using language supervision in both cases. This approach enables our model to generate full-bandwidth audio (48kHz) that meets professional standards and enhances precise text-based customization.

MULTIFOLEY consists of a diffusion transformer, a high-quality audio autoencoder (based on [52]), a frozen video encoder for audio-video synchronization, and a novel multi-conditional training strategy enabling flexible downstream tasks like audio extension and text-driven sound design. Our key contributions are as follows:

- We present a unified framework for video-guided Foley generation that leverages multiple conditioning modalities—text, audio, and video—within a single model.
- We introduce a training approach that combines internet video datasets with low-quality audio and professional libraries via language as bridges to enable high-quality, full-bandwidth audio generation at 48kHz.
- We show that our model supports a diverse range of applications, such as text-controlled Foley generation, audio-controlled Foley generation, and Foley extension, expanding possibilities for creative audio production.
- Through extensive quantitative evaluations and human studies, we demonstrate that MultiFoley achieves better audio quality and improved cross-modal alignment, outperforming existing methods in key benchmarks.

2. Related Work

Video-to-audio generation. Video-to-audio generation has recently gained significant attention. Several works have employed auto-regressive transformer models to generate audio from visual features, *e.g.*, SpecVQGAN [43], Im2Wav [82], FoleyGen [66], and V-AURA [89]. Other approaches have introduced video-to-audio diffusion or flow matching models to address this task, *e.g.*, Diff-Foley [63], Action2Sound [11], VTA-LDM [96], and Frieren [92]. Some researchers have adapted the MaskGIT [10] framework for video-to-audio synthesis [61, 75, 87]. Other approaches use a two-stage process: first, extracting time-varying signals like sound onsets or energy curves from videos, then adapting pretrained text-to-audio diffusion models with specialized adapters for video-to-audio generation with optional text [22, 42, 45, 53, 94, 100]. V2A-Mapper [90] translates CLIP visual embeddings into CLAP space to condition audio generation via AudioLDM [56]. Recently, *concurrent unpublished* work Movie Gen Audio [77] explores generating audio conditioned on video and text. This work demonstrates that text provides complementary information for audio generation. Concurrent work MMAudio [20] also explores multimodal joint training by combining video-to-audio and text-to-audio generation across multiple datasets. In contrast, our approach stands out by providing richer multimodal controls and Foley applications, such as generating on-screen sounds with semantically different audio via text and audio-conditional Foley generation, as shown in Fig. 1.

Diffusion models. Diffusion models [23, 38, 84–86] are generative models that learn to reconstruct data by reversing a process in which data is gradually corrupted with noise. By iterative denoising, these models generate new samples starting from random noise. Latent diffusion models (LDMs) [79] perform the diffusion process in a latent space by translating data using a pretrained encoder and decoder pair. These models have proven useful for various generative applications, such as image generation and editing [9, 23, 30, 36, 67, 70, 79, 80], video generation [7, 34, 39, 83], audio generation [25, 26, 54, 57, 64, 97], and 3D generation [8, 40, 60]. They have also been applied to tasks such as semantic segmentation [2, 95], camera pose estimation [91, 99], depth estimation [46, 81], and compositional generation [6, 19, 31, 32, 59]. In our work, we leverage diffusion models for video-guided Foley sound generation, incorporating controls from multiple modalities.

Audio-visual learning. Many works have focused on learning multimodal representations from audio-visual data. Some study *semantic correspondence* between sight and sound, *e.g.*, learning common audio-visual associations [3, 5, 33, 55, 69], audio-visual sound localization [4, 41, 74], and audio-visual segmentation [58, 101]. Some investigate *temporal correspondence* between audio and video [13,

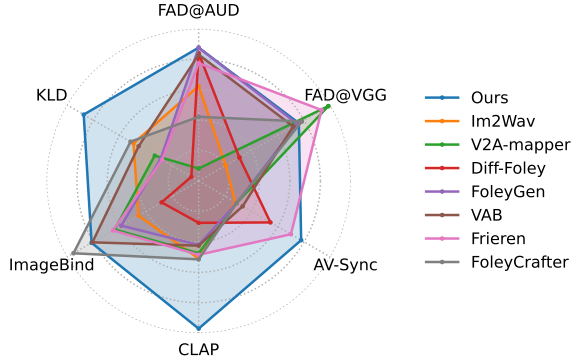


Figure 2. **Radar chart comparison for video-to-audio generation task.** Each metric is normalized for a better visualization.

[27, 44, 71, 72, 88], studying how sounds synchronize with visual events over time. Some explore *spatial correspondence* between audio and visual data, particularly how sound conveys information about the spatial environment [14–18, 28, 65, 98]. Inspired by them, we propose a joint learning approach across text, audio, and video modalities, enabling flexible control over various aspects of generated audio.

Generative models for sound design. Foley sound design has traditionally relied on Foley artists to create synchronized sound effects using everyday objects or sound libraries [1]. Recent text-to-audio models [25, 26, 64] give sound designers flexible control, allowing them to generate audio attributes directly from descriptive text. However, sound designers still need to manually synchronize the generated audio with the video. T-Foley [21] proposes to address this issue by guiding foley generation using temporal event features like the RMS of the waveform. Video-to-audio models [89, 92] generate synchronized audio but lack user control options. Du et al. [24] introduces a conditional Foley task for generating soundtracks from additional video examples. Building on the work above, our model offers artists improved tools for generating synchronized audio with versatile controls using text, audio, and video.

3. Multimodal Conditional Foley Generation

We aim to generate synchronized Foley sounds for silent videos, guided by multimodal inputs like text, audio, or video examples that define how the video should sound. First, we introduce a simple framework called MULTIFOLEY that performs video-guided Foley generation with multimodal conditions, shown in Fig. 3. Then, we show various applications our model enables for conditional Foley generation.

3.1. Generative Formulation

Given an input silent video \mathbf{v}_q , our goal is to generate a synchronized soundtrack $\hat{\mathbf{a}}$. This generation is conditioned on a text prompt \mathbf{t}_c and a reference audio-visual pair $(\mathbf{a}_c, \mathbf{v}_c)$, which can be part of \mathbf{v}_q or from different videos. We frame this problem as a conditional diffusion task. We learn a diffusion model ϵ_θ , parameterized by θ , to model the conditional

distribution $p_\theta(\mathbf{a}|\mathbf{v}_q, \mathbf{t}_c, \mathbf{a}_c, \mathbf{v}_c)$ from the audio \mathbf{a} .

We perform the diffusion process in a latent space [79]. We use a pretrained audio latent encoder \mathcal{E}_a and decoder \mathcal{D}_a that converts the waveform $\mathbf{a} \in \mathbb{R}^{T_a}$ into the compressed latent code $\mathbf{z} = \mathcal{E}_a(\mathbf{a}) \in \mathbb{R}^{T_z \times C_z}$, where T_a and T_z represent the lengths of the waveform and audio latents respectively, and C_z denotes the dimension of the latent features. During the forward diffusion process, we gradually add Gaussian noise to clean audio latents \mathbf{z}_0 to obtain noisy latent \mathbf{z}_t at different timesteps $t \in \{1, \dots, T\}$. During the reverse process, the denoiser $\epsilon_\theta(\cdot)$ computes the noise estimates $\hat{\epsilon} = \epsilon_\theta(\mathbf{z}_t, \mathbf{v}_q, \mathbf{t}_c, \mathbf{a}_c, \mathbf{v}_c, t)$ and optimize following [38, 79]:

$$\mathcal{L}_{\text{LDM}} = \mathbb{E}_{\epsilon \sim \mathcal{N}(0, \mathbf{I}), t \sim \mathcal{U}(T)} \left[\|\epsilon - \hat{\epsilon}\|_2^2 \right], \quad (1)$$

where conditions are encoded by the visual encoder \mathcal{E}_v , audio encoder \mathcal{E}_a or text encoder \mathcal{E}_t , respectively. During inference, we iteratively denoise the Gaussian noise \mathbf{z}_T with the noise estimates $\hat{\epsilon}$ to obtain the final clean latent $\hat{\mathbf{z}}_0$, which is then decoded into the waveform $\hat{\mathbf{a}} = \mathcal{D}_a(\hat{\mathbf{z}}_0)$.

3.2. Multimodal Conditioning and Training

We build our model on the Diffusion Transformer (DiT) [76] and a pretrained DAC-VAE [52] for the audio encoder and decoder, as shown in Fig. 3.

Multimodal conditioning. We use a pretrained visual encoder \mathcal{E}_v to encode the silent video $\mathbf{v}_q \in \mathbb{R}^{T_v \times 3 \times H \times W}$ into features $\mathcal{E}_v(\mathbf{v}_q) \in \mathbb{R}^{T_v \times C_v}$, where T_v is the number of video frames, H, W are video height and width, and C_v is the visual feature dimension size. These video features are interpolated to match the length of the audio latents \mathbf{z} , then concatenated along the channel dimension with noisy audio latents \mathbf{z}_t to ensure the audio-visual temporal alignment, pairing audio and video features at the same time, similar to Wang et al. [92]. The combined features are passed to the feed-forward transformer $\epsilon_\theta(\cdot)$.

For the text condition \mathbf{t}_c , we apply a frozen text encoder \mathcal{E}_t to extract embeddings $\mathcal{E}_t(\mathbf{t}_c) \in \mathbb{R}^{T_t \times C_t}$, where T_t is the tokenized text sequence length, and C_t is token’s embedding dimension. These text embeddings are then incorporated into $\epsilon_\theta(\cdot)$ via cross-attention.

To improve our model’s ability to condition on audio and video, during training we provide the model with a sample of ground truth audio-visual segments, $\mathcal{E}_a(\mathbf{a}_c)$ and $\mathcal{E}_v(\mathbf{v}_c)$, which we select randomly from the clean video. Our training loss is then only applied to the latents that were not provided as a conditioning signal. At inference, we can obtain our audio-visual conditioning, $(\mathbf{a}_c, \mathbf{v}_c)$, from a *different* video.

Learning from combined datasets. To enhance audio generation quality and text control, we train our model jointly on an audio-visual(-text) dataset, VGGSound [12], and a large proprietary licensed dataset of high-quality sound effects with text captions, termed as *HQ-SFX*. The audio in

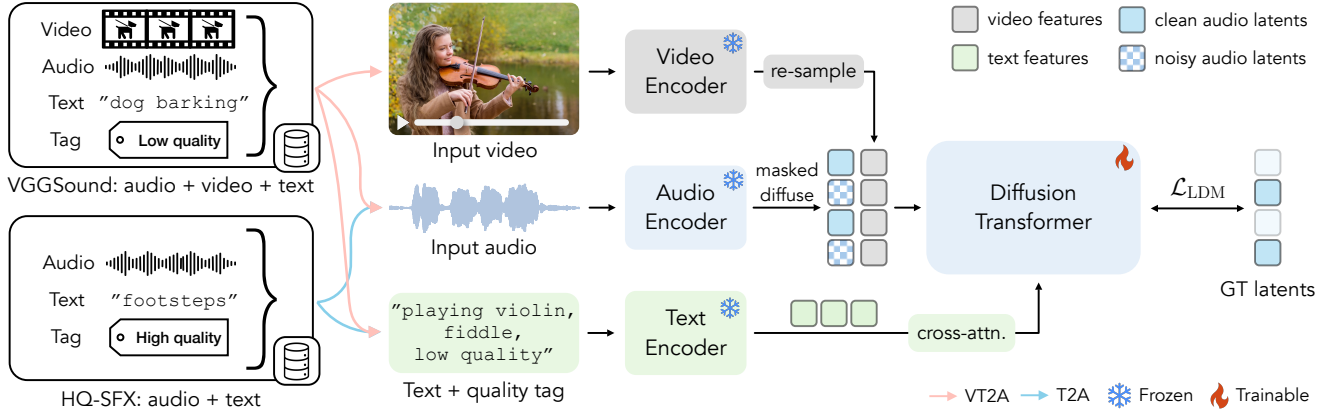


Figure 3. **Method overview.** We train our model jointly on a standard audio-video dataset *VGGSound* for VT2A generation and a high-quality audio-text dataset *HQ-SFX* for T2A generation. We encode the input audio into latents, adding noise to a portion of them. The silent video is encoded into visual features, concatenated with the audio latents along the channel dimension. The text input, including a quality tag, is encoded through a text encoder and applied via cross-attention.

these datasets is quite different, since *VGGSound* is “in-the-wild” and relatively noisy, whereas *HQ-SFX* contains professionally recorded sound effects and thus is generally higher quality and full bandwidth. To give the model control over audio quality, we assign *quality tags* to both datasets, which represent different dataset distributions: *VGGSound* examples are paired with a “low quality” tag in prompts, *e.g.*, “category, low quality”. In contrast, *HQ-SFX* examples are labeled with a “high quality” tag in their captions. During training, we randomly drop out the caption or quality tag to allow the model to disentangle audio quality from semantics and encourage the model to learn and distinguish audio quality across different data distributions.

Implementation. Our denoiser ϵ_θ consists of a 12-layer DiT. We train a DAC-VAE to map 48kHz audio waveforms into 40Hz latent features with a channel dimension C_z of 64, the encoder is frozen during the training. We use the T5-base text encoder [78] to obtain embeddings with $C_t = 768$, and we apply CAVP [63] to extract visual features at 8 FPS for 8-sec videos with $T_v = 64$ and $C_v = 512$.

The model is jointly trained for video-text-to-audio generation on *VGGSound*, containing 168K samples, and for text-to-audio generation on the *HQ-SFX* dataset, with 400K samples. We balance the datasets to ensure robust training and train the model for 600K iterations with a batch size of 128. During training, audio latents are masked with a probability of 0.25, where the masking spans 0 to 2 seconds. We also randomly drop the video, caption, and quality tag conditions during the training. We then finetune this pretrained model for 50k iterations on a curated subset of *VGGSound*, selected for high audio-visual correspondence using an ImageBind [33] score threshold of 0.3. During inference, we use the DDIM [85] method and apply a classifier-free guidance scale ranging from 3.0 to 5.0 to denoise over 100 steps. Please see more details in the supplementary material.

3.3. Foley Applications

Video-guided Foley with text control. Training audio diffusion with both video and text inputs allows our model to generate Foley sounds controlled by text. This dual conditioning setup enables flexible audio generation, where users can influence sound characteristics based on text prompts.

Our approach associates language with video cues, disentangling the semantic and temporal elements of videos. This setup allows for creative Foley applications, such as modifying a bird-chirping video to sound like a human voice or transforming a typewriter’s sound into piano notes – all while remaining synchronized with the video. Also, negative prompting provides a way to exclude unwanted audio elements by specifying them in the negative text prompts, offering flexible control over the audio output. To achieve this application, we use classifier-free guidance (CFG) [37] with negative prompting. Given a positive text prompt \mathbf{t}_{pos} for the desired sound effect and a negative prompt \mathbf{t}_{neg} representing the original expected sound or an unconditional text embedding \emptyset_t , we compute a diffusion step via:

$$\hat{\epsilon} = (\gamma + 1) \cdot \epsilon_\theta(\mathbf{z}_t, \mathbf{v}_q, \mathbf{t}_{\text{pos}}, t) - \gamma \cdot \epsilon_\theta(\mathbf{z}_t, \emptyset_v, \mathbf{t}_{\text{neg}}, t), \quad (2)$$

where γ controls the guidance strength, and \emptyset_v is the unconditional embedding for the input videos.

Video-guided Foley with audio-visual control. Our model enables generating synchronized soundtracks guided by both audio and video inputs. Our model can apply the sound characteristics (*e.g.*, rhythm and timbre) of reference audio from an SFX library to a silent video, synchronize the audio with visual events, and enable control over the generated output based on the reference audio-visual conditions. We frame this task as the video-guided audio extension problem, where we prepend the conditional audio latent $\mathbf{z}_c = \mathcal{E}_a(\mathbf{a}_c)$ and optional video feature $\mathcal{E}_v(\mathbf{v}_c)$ to the noisy audio latents \mathbf{z}_T along the sequence dimension and apply

Table 1. **Evaluation of video-to-audio generation.** Each method is evaluated on the VGGSound test set across six metrics assessing cross-modal alignment and audio quality. ImageBind and CLAP scores are reported in %. The unit of AV-Sync is seconds. DAC-VAE reconstructs the VGGSound audio and serves as an oracle baseline. The best results are in **bold**.

Method	Sampling rate (Hz)	Cross-modal alignment			Audio quality		
		ImageBind \uparrow	CLAP \uparrow	AV-Sync \downarrow	FAD@VGG \downarrow	FAD@AUD \downarrow	KLD \downarrow
Im2Wav [82]	16K	22.3	25.2	1.25	7.36	5.88	2.11
V2A-mapper [90]	16K	25.2	24.5	1.24	1.13	8.59	2.40
Diff-Foley [63]	16K	19.5	20.5	1.01	6.53	4.80	2.90
FoleyGen [66]	16K	24.4	23.4	1.24	3.01	4.62	2.48
VAB [87]	16K	27.9	23.5	1.20	3.26	4.85	2.18
Frieren [92]	16K	25.4	24.7	0.87	1.54	5.13	2.50
FoleyCrafter [100]	16K	30.2	25.3	1.24	2.74	6.89	2.07
MultiFoley (ours)	48K	28.0	34.4	0.80	2.92	4.62	1.43
DAC-VAE (VGGSound)	48K	35.4	28.2	0.62	1.21	5.91	0.28

masked denoising to generate the missing sound:

$$\hat{\epsilon} = (\gamma + 1) \cdot \epsilon_{\theta}([z_c; z_t], [v_c; v_q], t) - \gamma \cdot \epsilon_{\theta}([\varnothing_a; z_t], \varnothing_v, t), \quad (3)$$

where \varnothing_a denotes the noisy latent that matches the size of the conditional latent z_c . v_c and v_q are encoded by \mathcal{E}_v .

Video-guided Foley with quality control. We enforce quality control by incorporating quality tags in the text, enabling the generation of clean full-band (48kHz) audio. During inference, we guide the model to produce samples that align with the high-quality audio distribution while steering away from low-quality audio using CFG with a negative prompt t_{neg} ; the prompt can be “low quality” or the unconditional text embedding \varnothing_t :

$$\hat{\epsilon} = (\gamma + 1) \cdot \epsilon_{\theta}(z_t, v_q, t_c + \text{“high quality”}, t) - \gamma \cdot \epsilon_{\theta}(z_t, \varnothing_v, t_{\text{neg}}, t). \quad (4)$$

4. Experiments

In this section, we quantitatively and qualitatively evaluate our method for the tasks of video-to-audio generation and video-guided Foley generation with multimodal controls over text, reference audio, and video.

4.1. Video-to-Audio Generation

First, we evaluate the ability of our model in the video-to-audio generation task, *i.e.*, reproducing the soundtracks for silent videos, through an automatic quantitative evaluation.

Experimental setup. We evaluate this task using videos from the VGGSound test set [12]. To ensure accurate audio-visual correspondence, we apply ImageBind [33] to filter out test samples with a score below 0.3, yielding a final set of 8,702 videos, following [89, 96]. For each video, we generate 8-second audio samples. We compare our model against several video-to-audio baselines, including the autoregressive models Im2Wav [82] and FoleyGen [66], latent

diffusion models Diff-Foley [63], V2A-mapper [90] and FoleyCrafter [100], Frieren [92] based on rectified flow matching, and VAB [87] that apply MaskGIT [10] framework. We also report the performance of the reconstructed audio with our DAC-VAE on the VGGSound test set as an oracle baseline. We trim the generated audio from the baselines to 8 seconds for a fair evaluation. Our model approaches the video-to-audio generation task as video-text-to-audio (VT2A) generation, using the VGGSound category name as the text input. During the inference, we use the “low quality” tag for our model’s generation to stay within the same data distribution of VGGSound for a fair comparison and a guidance scale of 3.0 for diffusion sampling.

Evaluation metrics. Following prior work [66, 89, 92], we evaluated model performance in terms of audio quality and cross-modal alignment. We evaluated audio quality using Fréchet Audio Distance (FAD) [47] with VGGish [35] on 16kHz audio, which measures the distribution distance between generated and reference audio. For reference sets, we used the VGGSound test set as regular references and Adobe SFX Audition¹ (a professional audio library that differs from VGGSound) as high-quality references, denoting these metrics as FAD@VGG and FAD@AUD, respectively. We also use Kullback–Leibler Divergence (KLD) [51] to measure the probability distributions of class predictions by the PaSST [50] model between the ground-truth and generated samples. To assess cross-modal alignment, we use ImageBind [33] to measure the semantic correspondence between the generated audio and the input video. We also compute the CLAP score [93] to evaluate the similarity between category labels and generated audio. To measure the cross-modal (temporal) alignment between generated audio and input videos, we apply Synchformer [44] to estimate the weighted temporal offset in seconds, following [89]. The model classifies the offset from -2.0s and 2.0s (with a 0.2s

¹<https://www.adobe.com/products/audition/offers/adobeauditiondlcsfx.html>

resolution). The final AV-Sync metric is the average of the absolute offsets across all examples.

Quantitative results. We show our quantitative results in Tab. 1 and Fig. 2 and demonstrate that our method outperforms all the other methods in multiple metrics, including AV-Sync, CLAP score, FAD@AUD, and KLD. This highlights the strong overall performance of our model against baselines. Notably, our synchronization score is comparable to the DAC-VAE reconstructed results, indicating that we successfully generated synced audio for silent input videos. Additionally, our method achieves the second-best performance on the ImageBind score, reflecting strong cross-modal semantic alignment. We outperform the oracle baseline (DAC-VAE) on the CLAP score indicating that our generated clips are more semantically aligned to corresponding sound effects. The performance of DAC-VAE on FAD@VGG indicates that the VAE influences the generated data distribution, which in turn affects the FAD evaluation.

4.2. Video-guided Foley with Text Control

We conduct experiments to evaluate our model’s capability for video-guided Foley generation with text controls, focusing on synchronization and diversity of semantics.

Experimental setup. To quantitatively evaluate the semantic control of the model with text, we sampled videos of 10 categories from the VGGSound-sync [13] dataset as a test set with the ImageBind [33] filtering strategy described above. We individually provided the other 9 category names for each video as target text prompts and tasked the models with generating audio based on the given text and video. We generated 4 audio tracks for each pair. We compared our model with FoleyCrafter [100], which also supports text-based control. We also modify FoleyCrafter to disable its semantic adapter to cut the semantic signal from input videos, using this model solely as a video-onset Control-Net on the text-to-audio generation model. We evaluated four variants of our approach: 1) **w/ NegP**: using target prompts as positive prompts and the ground truth category as negative prompts for classifier-free guidance; 2) **w/o NegP**: no negative prompts are used; 3) **True category**: regular video-text-to-audio generation with true category as positive prompts (without negative prompts) to generate expected Foley sound for videos, representing the best synchronization performance that our models could achieve; 4) **T2A**: text-to-audio generation with the given target individual prompts, providing the upper bound for CLAP metrics.

Evaluation metrics. To evaluate semantic alignment between target prompts and generated audio, we use two CLAP-based [93] metrics. First, we calculate the CLAP score, defined as the average cosine similarity between the CLAP embeddings of each text prompt and generated audio pair. Additionally, we use the CLAP model as a classifier to compare target category scores against original categories (from

Table 2. **Evaluation on the Foley generation with text controls.** *NegP* denotes negative prompting. We also include two oracle baselines: one using the true category as text prompts and the other omitting video during inference. The best results are in **bold**.

Method	Variation	CLAP \uparrow		AV-Sync \downarrow
		Score	Acc	
FoleyCrafter [100]	w/o NegP	38.4	99.4	1.34
	(w/o semantic adapter) w/ NegP	35.7	99.9	1.36
FoleyCrafter [100]	w/o NegP	31.0	79.2	1.29
	w/ NegP	33.4	94.2	1.31
Ours	w/o NegP	31.4	85.5	0.81
	w/ NegP	30.9	93.2	0.93
Ours – oracle	True category	4.2	1.8	0.77
	T2A	40.3	100	1.38

videos), reporting binary classification accuracy. We also report temporal performance using the aforementioned AV-Sync metric. Although generated audio-visual examples fall outside the synchronization model’s training distribution, we found that it returns reliable scores.

Results. We present the results in Tab. 2. Our method achieves the highest synchronization performance, comparable to the oracle model (Ours–true category) that generates soundtracks with original semantics. The CLAP metrics are also reasonable, indicating that our approach effectively generates synchronized audio for input videos while allowing semantic control through text. Although FoleyCrafter [100] performs well on CLAP-based metrics, its AV-Sync score is similar to text-to-audio generation, suggesting it struggles to generate synchronized audio in this task. Additionally, the results of the CLAP accuracy highlight that the negative prompting strategy helps steer the generation away from reproducing the original semantics of the audio, enhancing semantic control through classifier-free guidance while the sync score is slightly dropped. We also show some qualitative examples for text control in Fig. 4, demonstrating that our model enables the disentangling of semantic information from the input video and maintains the temporal information.

Human studies. We conducted a human evaluation using two-alternative forced choice (2AFC) studies. We selected 10 high-quality videos from the VGGSound test set, ensuring diverse categories and clear temporal information. For each video, we used one prompt based on the original video’s semantics and another that diverged from it. We compared our method with FoleyCrafter [100].

In the user study, participants watch and listen to two video clips, each paired with an audio sample—one generated by our model and the other from the baseline. They were asked to select which audio (1) best matches the sound of “audio prompt”, (2) is best synchronized with the video, (3) which audio sounds cleaner and more high definition, and (4) overall sounds best when considering the intended

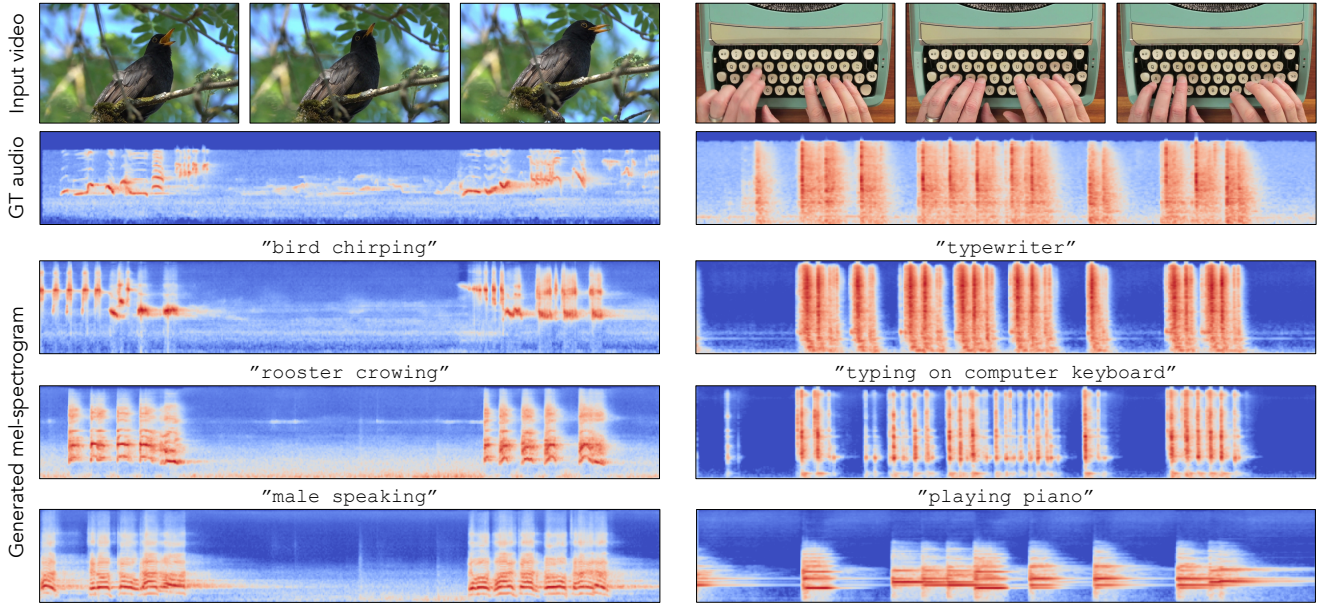


Figure 4. **Qualitative examples for Foley generation with text control.** We present generated results for two videos, each with three different text prompts, demonstrating our model’s ability to produce synchronized soundtracks with varied semantics through text control. Please refer to our [website](#) for video results.

Table 3. **Human Evaluation on the Foley generation with text control.** We show the win rates of our method against FoleyCrafter [100]. 95% confidence intervals are reported in gray. P-values are below 10^{-20} . ($N = 400$ with 20 participants)

Comparison	Win rate (%)			
	Semantic	Sync.	Quality	Overall
Ours vs FoleyCrafter	85.8 (± 3.4)	94.5 (± 2.1)	86.5 (± 3.4)	90.2 (± 2.9)

Table 4. **Evaluation on the Foley extension with different control signals.** \mathcal{V}_q means input silent videos. \mathcal{T}_c , \mathcal{A}_c and \mathcal{V}_c denote text, audio and video conditional signals respectively. The best results are in **bold**.

Eval set	Conditions				CLAP \uparrow	AV-Sync \downarrow
	\mathcal{V}_q	\mathcal{T}_c	\mathcal{A}_c	\mathcal{V}_c		
VGGSound	✓	✓			55.4	0.79
	✓		✓		59.6	0.78
	✓		✓	✓	59.8	0.77
	✓	✓	✓	✓	64.3	0.77
Greatest-Hits	✓				29.3	0.88
	✓		✓		73.8	0.94
	✓		✓	✓	74.4	0.87

audio for “audio prompt”. Additional details of the study can be found in the supplementary material.

We demonstrate the user study results in Tab. 3. As can be seen, our method outperforms the baseline in all the aspects. Human evaluators consistently rate our results as being higher in semantic alignment, audio-visual synchronization, and audio quality. It demonstrates our model’s capability in video-guided Foley generation with text control, as well as the quality of the generated audio.

4.3. Video-guided Foley with Audio-Visual Control

Our model emerges with the ability to generate video-guided Foley using reference audio-visual examples, effectively transferring sounds from conditional clips to generate synchronized soundtracks for silent videos. We evaluate this on the Foley extension task, where the first few seconds of a sounding video serve as the condition, and the model generates the remaining audio for the silent part of the video.

Experiment setup. We evaluate this task on two datasets: VGGSound-Sync [13] and Greatest-Hits [73]. We sample 1,000 8-second video examples from VGGSound-Sync as in-domain data, and 800 8-second examples from Greatest-Hits as out-of-domain data. For each video, the first 3 seconds of both audio and video (\mathbf{a}_c , \mathbf{v}_c) serve as conditional inputs, while the model generates the remaining 5 seconds of audio $\hat{\mathbf{a}}$ for the corresponding silent video segment \mathbf{v}_q . We evaluate the results using two metrics: (1) CLAP score to measure the cosine similarity between the ground-truth and generated audio, and (2) AV-Sync score to evaluate synchronization accuracy. We generate four audio clips for each video and compare our model’s performance across different combinations of multimodal conditional inputs (\mathbf{t}_c , \mathbf{a}_c , \mathbf{v}_c). Text inputs are omitted to evaluate the Greatest-Hits dataset, focusing solely on the audio and video conditions.

Results. We present the results in Tab. 4. Across both datasets, CLAP similarity improves substantially with the addition of audio conditions. When video conditions are added, we observe further improvement in the AV-Sync score, particularly for the Greatest-Hits dataset. This suggests that the model effectively leverages in-context learning to interpret information from conditional audio-visual inputs.

Table 5. **Quantitative evaluation on quality control and ablation study.** We evaluate our model with different inference settings, *i.e.*, using various quality tags and excluding text input. We also ablate the impact of the subset fine-tuning strategy. *NegP* denotes negative prompting. The best results are in **bold**.

	Variation	Inference	Quality tag	ImageBind \uparrow	CLAP \uparrow	AV-Sync \downarrow	FAD@VGG \downarrow	FAD@AUD \downarrow	KLD \downarrow
Ours	Full w/o ft.	VT2A	Low	27.3	33.8	0.81	3.00	4.39	1.47
	Full	VT2A	Low	28.0	34.4	0.80	2.92	4.62	1.43
	Full	VT2A	High	25.8	34.9	0.83	4.37	4.09	1.60
	Full	V2A (w/o text)	-	22.4	19.4	0.77	4.78	3.44	2.59

Additionally, our model demonstrates strong capability in Foley analogy tasks as shown in Fig. 1. For instance, given a reference dog bark, the model can produce synchronized Foley audio for a different dog barking video, or generate drum sounds for a basketball dribble video with a reference drum audio sample. Please see our [website](#) for examples.

4.4. Video-guided Foley with Quality Control

During our experiments, we observed that audio from VGGSound videos downloaded via YouTube is often compressed to 32kHz or lower. When resampled to 48kHz, the high-frequency content above 32kHz is missing, as shown in Fig. 5. Consequently, when training on the VGGSound dataset, models inherently generate audio with a quality aligned to 32kHz, reflecting the dataset’s distribution.

Our model incorporates high-quality text-audio data at 48kHz during training to address this limitation, paired with quality tags. This enables the generation of full-band, 48kHz audio for video-guided Foley generation. During inference, we use the “high quality” tag to guide the model toward generating audio that follows the distribution of the high-quality text-audio dataset, ensuring 48kHz output.

We provide examples in Fig. 5, demonstrating the model’s ability to control and improve audio quality. Additionally, quantitative results in Tab. 5 support this, where we observe that with the “high quality” tag, our model achieves better performance on the FAD@AUD metric, while performance on FAD@VGG decreases. This suggests the model effectively generates audio that aligns more closely with high-quality sound effect distributions.

4.5. Ablation

Subset fine-tuning. As discussed in Sec. 3.2, the VGGSound dataset includes noisy, low-quality samples where audio and video often lack alignment. To enhance model performance, we further fine-tune our pretrained model on a subset of VGGSound containing better audio-visual correspondence. We evaluate our pretrained model performance and compare it with the fine-tuned model. We present results in Tab. 5, demonstrating that fine-tuning significantly improves cross-modal alignment metrics.

No-text inference. We also perform an ablation study where we remove text conditioning and evaluate our model directly on video-to-audio generation. We report the results in Tab. 5. We observe that while performance on the

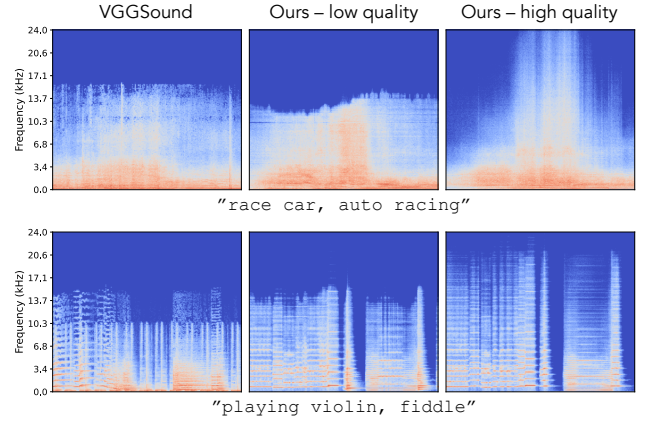


Figure 5. **Qualitative results of quality control.** We show that VGGSound audio has limited bandwidth and demonstrate our model generates full-band 48kHz audio with quality control.

semantics-related metrics drops significantly, synchronization metrics remain high, indicating that our model relies on text to drive semantic alignment but can retain temporal consistency even in the absence of any conditioning.

5. Conclusion

In this paper, we present MULTIFOLEY, a Foley system designed for video-guided Foley generation using multimodal inputs, including text, audio, and video. We evaluate our model on standard video-to-audio generation tasks, providing quantitative evidence of its effectiveness. We explore the model’s control capabilities with different conditional inputs through both quantitative and qualitative experiments, illustrating a range of Foley applications achievable with our approach. Our work is a step toward the broader goal of “user-in-the-loop” sound design. By providing easy-to-use, multimodal controls, we aim to help users create customized, synchronized high-quality audio.

Limitations and broader impacts. Our model is currently trained on a small-scale, in-the-wild audio-visual dataset, VGGSound, which constrains its capabilities. We believe a larger, high-quality Foley dataset would significantly enhance our model’s performance and broaden its applicability. Our model currently struggles with handling multiple sound events alongside text, often leading to confusion about the timing of each event. Our method can also create realistic but potentially misleading media. Responsible use is essential to prevent misuse in situations where authenticity matters.

Acknowledgements. We thank Sonal Kumar, Hugo Flores García, Xiulong Liu, Yongqi Wang, and Yiming Zhang for their valuable help and discussion, and Adolfo Hernandez Santisteban for informing the project with his sound design experience. This work was supported in part by an NSF CAREER Award #2339071.

References

- [1] Vanessa Theme Ament. *The Foley grail: The art of performing sound for film, games, and animation*. Routledge, 2014. 1, 3
- [2] Tomer Amit, Tal Shaharbany, Eliya Nachmani, and Lior Wolf. Segdiff: Image segmentation with diffusion probabilistic models. *arXiv preprint arXiv:2112.00390*, 2021. 2
- [3] Relja Arandjelovic and Andrew Zisserman. Look, listen and learn. In *Proceedings of the IEEE international conference on computer vision*, pages 609–617, 2017. 2
- [4] Relja Arandjelovic and Andrew Zisserman. Objects that sound. In *Proceedings of the European conference on computer vision (ECCV)*, pages 435–451, 2018. 2
- [5] Yuki Asano, Mandela Patrick, Christian Rupprecht, and Andrea Vedaldi. Labelling unlabelled videos from scratch with multi-modal self-supervision. *Advances in Neural Information Processing Systems*, 33:4660–4671, 2020. 2
- [6] Omer Bar-Tal, Lior Yariv, Yaron Lipman, and Tali Dekel. Multidiffusion: Fusing diffusion paths for controlled image generation. *arXiv preprint arXiv:2302.08113*, 2023. 2
- [7] Omer Bar-Tal, Hila Chefer, Omer Tov, Charles Herrmann, Roni Paiss, Shiran Zada, Ariel Ephrat, Junhwa Hur, Yuanzhen Li, Tomer Michaeli, et al. Lumiere: A space-time diffusion model for video generation. *arXiv preprint arXiv:2401.12945*, 2024. 2
- [8] Raphael Bensadoun, Tom Monnier, Yanir Kleiman, Filippos Kokkinos, Yawar Siddiqui, Mahendra Kariya, Omri Harosh, Roman Shapovalov, Benjamin Graham, Emilien Garreau, et al. Meta 3d gen. *arXiv preprint arXiv:2407.02599*, 2024. 2
- [9] Tim Brooks, Aleksander Holynski, and Alexei A Efros. Instructpix2pix: Learning to follow image editing instructions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18392–18402, 2023. 2
- [10] Huiwen Chang, Han Zhang, Lu Jiang, Ce Liu, and William T Freeman. Maskgit: Masked generative image transformer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11315–11325, 2022. 2, 5
- [11] Changan Chen, Puyuan Peng, Ami Baid, Sherry Xue, Wei-Ning Hsu, David Harwath, and Kristen Grauman. Action2sound: Ambient-aware generation of action sounds from egocentric videos. In *ECCV*, 2024. 2
- [12] Honglie Chen, Weidi Xie, Andrea Vedaldi, and Andrew Zisserman. Vgggsound: A large-scale audio-visual dataset. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 721–725. IEEE, 2020. 3, 5
- [13] Honglie Chen, Weidi Xie, Triantafyllos Afouras, Arsha Nagrani, Andrea Vedaldi, and Andrew Zisserman. Audio-visual synchronisation in the wild. *arXiv preprint arXiv:2112.04432*, 2021. 2, 6, 7
- [14] Mingfei Chen, Kun Su, and Eli Shlizerman. Be everywhere—hear everything (bee): Audio scene reconstruction by sparse audio-visual samples. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7853–7862, 2023. 3
- [15] Ziyang Chen, Xixi Hu, and Andrew Owens. Structure from silence: Learning scene structure from ambient sound. In *5th Annual Conference on Robot Learning*, 2021.
- [16] Ziyang Chen, David F Fouhey, and Andrew Owens. Sound localization by self-supervised time delay estimation. *European Conference on Computer Vision (ECCV)*, 2022.
- [17] Ziyang Chen, Shengyi Qian, and Andrew Owens. Sound localization from motion: Jointly learning sound direction and camera rotation. In *International Conference on Computer Vision (ICCV)*, 2023.
- [18] Ziyang Chen, Israel D. Gebru, Christian Richardt, Anurag Kumar, William Laney, Andrew Owens, and Alexander Richard. Real acoustic fields: An audio-visual room acoustics dataset and benchmark. In *The IEEE / CVF Computer Vision and Pattern Recognition Conference (CVPR)*, 2024. 3
- [19] Ziyang Chen, Daniel Geng, and Andrew Owens. Images that sound: Composing images and sounds on a single canvas. *Neural Information Processing Systems (NeurIPS)*, 2024. 2
- [20] Ho Kei Cheng, Masato Ishii, Akio Hayakawa, Takashi Shibuya, Alexander Schwing, and Yuki Mitsufuji. Taming multimodal joint training for high-quality video-to-audio synthesis. *arXiv preprint arXiv:2412.15322*, 2024. 2
- [21] Yoonjin Chung, Junwon Lee, and Juhan Nam. T-foley: A controllable waveform-domain diffusion model for temporal-event-guided foley sound synthesis. In *ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 6820–6824. IEEE, 2024. 3
- [22] Marco Comunità, Riccardo F Gramaccioni, Emiliano 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)*, pages 936–940. IEEE, 2024. 2
- [23] Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. *Advances in neural information processing systems*, 34:8780–8794, 2021. 2
- [24] Yuexi Du, Ziyang Chen, Justin Salamon, Bryan Russell, and Andrew Owens. Conditional generation of audio from video via foley analogies. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2426–2436, 2023. 2, 3
- [25] Zach Evans, CJ Carr, Josiah Taylor, Scott H Hawley, and Jordi Pons. Fast timing-conditioned latent audio diffusion. *arXiv preprint arXiv:2402.04825*, 2024. 2, 3
- [26] Zach Evans, Julian D Parker, CJ Carr, Zack Zukowski, Josiah Taylor, and Jordi Pons. Stable audio open. *arXiv preprint arXiv:2407.14358*, 2024. 2, 3

- [27] Chao Feng, Ziyang Chen, and Andrew Owens. Self-supervised video forensics by audio-visual anomaly detection. *Computer Vision and Pattern Recognition (CVPR)*, 2023. 3
- [28] Ruohan Gao, Changan Chen, Ziad Al-Halah, Carl Schissler, and Kristen Grauman. Visualechoes: Spatial visual representation learning through echolocation. In *European Conference on Computer Vision (ECCV)*, 2020. 3
- [29] Hugo Flores Garcia, Prem Seetharaman, Rithesh Kumar, and Bryan Pardo. Vampnet: Music generation via masked acoustic token modeling. *arXiv preprint arXiv:2307.04686*, 2023. 1
- [30] Daniel Geng and Andrew Owens. Motion guidance: Diffusion-based image editing with differentiable motion estimators. *arXiv preprint arXiv:2401.18085*, 2024. 2
- [31] Daniel Geng, Inbum Park, and Andrew Owens. Visual anagrams: Generating multi-view optical illusions with diffusion models. In *CVPR*, 2024. 2
- [32] Daniel Geng, Inbum Park, and Andrew Owens. Factorized diffusion: Perceptual illusions by noise decomposition. *arXiv:2404.11615*, 2024. 2
- [33] 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*, pages 15180–15190, 2023. 2, 4, 5, 6
- [34] Rohit Girdhar, Mannat Singh, Andrew Brown, Quentin Duval, Samaneh Azadi, Sai Saketh Rambhatla, Akbar Shah, Xi Yin, Devi Parikh, and Ishan Misra. Emu video: Factorizing text-to-video generation by explicit image conditioning. *arXiv preprint arXiv:2311.10709*, 2023. 2
- [35] Shawn Hershey, Sourish Chaudhuri, Daniel PW Ellis, Jort F Gemmeke, Aren Jansen, R Channing Moore, Manoj Plakal, Devin Platt, Rif A Saurous, Bryan Seybold, et al. Cnn architectures for large-scale audio classification. In *2017 IEEE international conference on acoustics, speech and signal processing (icassp)*, pages 131–135. IEEE, 2017. 5
- [36] Amir Hertz, Ron Mokady, Jay Tenenbaum, Kfir Aberman, Yael Pritch, and Daniel Cohen-Or. Prompt-to-prompt image editing with cross attention control. *arXiv preprint arXiv:2208.01626*, 2022. 2
- [37] Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. *arXiv preprint arXiv:2207.12598*, 2022. 4
- [38] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020. 2, 3
- [39] 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. 2
- [40] Lukas Höllein, Ang Cao, Andrew Owens, Justin Johnson, and Matthias Nießner. Text2room: Extracting textured 3d meshes from 2d text-to-image models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 7909–7920, 2023. 2
- [41] Xixi Hu, Ziyang Chen, and Andrew Owens. Mix and localize: Localizing sound sources in mixtures. *Computer Vision and Pattern Recognition (CVPR)*, 2022. 2
- [42] Zhiqi Huang, Dan Luo, Jun Wang, Huan Liao, Zhiheng Li, and Zhiyong Wu. Rhythmic foley: A framework for seamless audio-visual alignment in video-to-audio synthesis. *arXiv preprint arXiv:2409.08628*, 2024. 2
- [43] Vladimir Iashin and Esa Rahtu. Taming visually guided sound generation. In *British Machine Vision Conference (BMVC)*, 2021. 2
- [44] Vladimir Iashin, Weidi Xie, Esa Rahtu, and Andrew Zisserman. Synchformer: Efficient synchronization from sparse cues. *arXiv preprint arXiv:2401.16423*, 2024. 3, 5
- [45] 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. 2
- [46] Bingxin Ke, Anton Obukhov, Shengyu Huang, Nando Metzger, Rodrigo Caye Daudt, and Konrad Schindler. Repurposing diffusion-based image generators for monocular depth estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9492–9502, 2024. 2
- [47] Kevin Kilgour, Mauricio Zuluaga, Dominik Roblek, and Matthew Sharifi. Fr\`echet audio distance: A metric for evaluating music enhancement algorithms. *arXiv preprint arXiv:1812.08466*, 2018. 5
- [48] Diederik Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *International Conference on Learning Representation*, 2015. 1
- [49] Diederik P Kingma. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013. 1
- [50] Khaled Koutini, Jan Schlüter, Hamid Eghbal-Zadeh, and Gerhard Widmer. Efficient training of audio transformers with patchout. *arXiv preprint arXiv:2110.05069*, 2021. 5
- [51] Solomon Kullback and Richard A Leibler. On information and sufficiency. *The annals of mathematical statistics*, 22 (1):79–86, 1951. 5
- [52] Rithesh Kumar, Prem Seetharaman, Alejandro Luebs, Ishaan Kumar, and Kundan Kumar. High-fidelity audio compression with improved rvqgan. *Advances in Neural Information Processing Systems*, 36, 2024. 2, 3, 1
- [53] Junwon Lee, Jaekwon Im, Dabin Kim, and Juhan Nam. Video-foley: Two-stage video-to-sound generation via temporal event condition for foley sound. *arXiv preprint arXiv:2408.11915*, 2024. 2
- [54] Tingle Li, Renhao Wang, Po-Yao Huang, Andrew Owens, and Gopala Anumanchipalli. Self-supervised audio-visual soundscape stylization. In *European Conference on Computer Vision*, pages 20–40. Springer, 2025. 2
- [55] Yan-Bo Lin and Gedas Bertasius. Siamese vision transformers are scalable audio-visual learners. *arXiv preprint arXiv:2403.19638*, 2024. 2
- [56] 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*, 2023. 2

- [57] Haohe Liu, Qiao Tian, Yi Yuan, Xubo Liu, Xinhao Mei, Qiuqiang Kong, Yuping Wang, Wenwu Wang, Yuxuan Wang, and Mark D Plumbley. Audioldm 2: Learning holistic audio generation with self-supervised pretraining. *arXiv preprint arXiv:2308.05734*, 2023. 2
- [58] Jinxiang Liu, Yikun Liu, Fei Zhang, Chen Ju, Ya Zhang, and Yanfeng Wang. Audio-visual segmentation via unlabeled frame exploitation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024. 2
- [59] Nan Liu, Shuang Li, Yilun Du, Antonio Torralba, and Joshua B Tenenbaum. Compositional visual generation with composable diffusion models. In *European Conference on Computer Vision*, pages 423–439. Springer, 2022. 2
- [60] Ruoshi Liu, Rundi Wu, Basile Van Hoorick, Pavel Tokmakov, Sergey Zakharov, and Carl Vondrick. Zero-1-to-3: Zero-shot one image to 3d object. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 9298–9309, 2023. 2
- [61] Xiulong Liu, Kun Su, and Eli Shlizerman. Tell what you hear from what you see – video to audio generation through text. *Advances in Neural Information Processing Systems*, 2024. 2
- [62] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017. 1
- [63] Simian Luo, Chuanhao Yan, Chenxu Hu, and Hang Zhao. Diff-foley: Synchronized video-to-audio synthesis with latent diffusion models. *Advances in Neural Information Processing Systems*, 36, 2024. 1, 2, 4, 5
- [64] Navonil Majumder, Chia-Yu Hung, Deepanway Ghosal, Wei-Ning Hsu, Rada Mihalcea, and Soujanya Poria. Tango 2: Aligning diffusion-based text-to-audio generations through direct preference optimization. *arXiv preprint arXiv:2404.09956*, 2024. 2, 3
- [65] Sagnik Majumder, Ziad Al-Halah, and Kristen Grauman. Learning spatial features from audio-visual correspondence in egocentric videos. *arXiv preprint arXiv:2307.04760*, 2023. 3
- [66] Xinhao Mei, Varun Nagaraja, Gael Le Lan, Zhaoheng Ni, Ernie Chang, Yangyang Shi, and Vikas Chandra. Foley-gen: Visually-guided audio generation. *arXiv preprint arXiv:2309.10537*, 2023. 1, 2, 5
- [67] Chenlin Meng, Yutong He, Yang Song, Jiaming Song, Jiajun Wu, Jun-Yan Zhu, and Stefano Ermon. Sdedit: Guided image synthesis and editing with stochastic differential equations. *arXiv preprint arXiv:2108.01073*, 2021. 2
- [68] Daniel Morales-Brotons, Thijs Vogels, and Hadrien Hendrikx. Exponential moving average of weights in deep learning: Dynamics and benefits. *Transactions on Machine Learning Research*, 2024. 1
- [69] Pedro Morgado, Nuno Vasconcelos, and Ishan Misra. Audio-visual instance discrimination with cross-modal agreement. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 12475–12486, 2021. 2
- [70] Alex Nichol, Prafulla Dhariwal, Aditya Ramesh, Pranav Shyam, Pamela Mishkin, Bob McGrew, Ilya Sutskever, and Mark Chen. Glide: Towards photorealistic image generation and editing with text-guided diffusion models, 2021. 2
- [71] Muhammad Adi Nugroho, Sangmin Woo, Sumin Lee, and Changick Kim. Audio-visual glance network for efficient video recognition. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10150–10159, 2023. 3
- [72] Andrew Owens and Alexei A Efros. Audio-visual scene analysis with self-supervised multisensory features. In *Proceedings of the European conference on computer vision (ECCV)*, pages 631–648, 2018. 3
- [73] Andrew Owens, Phillip Isola, Josh McDermott, Antonio Torralba, Edward H Adelson, and William T Freeman. Visually indicated sounds. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2405–2413, 2016. 7
- [74] Sooyoung Park, Arda Senocak, and Joon Son Chung. Can clip help sound source localization? In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 5711–5720, 2024. 2
- [75] Santiago Pascual, Chunghsin Yeh, Ioannis Tsiamas, and Joan Serra. Masked generative video-to-audio transformers with enhanced synchronicity. *arXiv preprint arXiv:2407.10387*, 2024. 2
- [76] William Peebles and Saining Xie. Scalable diffusion models with transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4195–4205, 2023. 3
- [77] 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. 2
- [78] 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. 4
- [79] 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*, pages 10684–10695, 2022. 2, 3
- [80] Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily Denton, Seyed Kamyar Seyed Ghasemipour, Burcu Karagol Ayan, S. Sara Mahdavi, Rapha Gontijo Lopes, Tim Salimans, Jonathan Ho, David J Fleet, and Mohammad Norouzi. Photorealistic text-to-image diffusion models with deep language understanding, 2022. 2
- [81] Saurabh Saxena, Abhishek Kar, Mohammad Norouzi, and David J Fleet. Monocular depth estimation using diffusion models. *arXiv preprint arXiv:2302.14816*, 2023. 2
- [82] Roy Sheffer and Yossi Adi. I hear your true colors: Image guided audio generation. In *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 1–5. IEEE, 2023. 2, 5
- [83] Uriel Singer, Adam Polyak, Thomas Hayes, Xi Yin, Jie An, Songyang Zhang, Qiyuan Hu, Harry Yang, Oron Ashual,

- Oran Gafni, et al. Make-a-video: Text-to-video generation without text-video data. *arXiv preprint arXiv:2209.14792*, 2022. 2
- [84] Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *Proceedings of the 32nd International Conference on Machine Learning*, pages 2256–2265, Lille, France, 2015. PMLR. 2
- [85] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502*, 2020. 4
- [86] Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. *arXiv preprint arXiv:2011.13456*, 2020. 2
- [87] Kun Su, Xiulong Liu, and Eli Shlizerman. From vision to audio and beyond: A unified model for audio-visual representation and generation. In *Forty-first International Conference on Machine Learning*, 2024. 1, 2, 5
- [88] Jiatian Sun, Longxiulin Deng, Triantafyllos Afouras, Andrew Owens, and Abe Davis. Eventfulness for interactive video alignment. *ACM Transactions on Graphics (TOG)*, 42(4):1–10, 2023. 3
- [89] Ilpo Viertola, Vladimir Iashin, and Esa Rahtu. Temporally aligned audio for video with autoregression. *arXiv preprint arXiv:2409.13689*, 2024. 2, 3, 5
- [90] Heng Wang, Jianbo Ma, Santiago Pascual, Richard Cartwright, and Weidong Cai. V2a-mapper: A lightweight solution for vision-to-audio generation by connecting foundation models. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 15492–15501, 2024. 1, 2, 5
- [91] Jianyuan Wang, Christian Rupprecht, and David Novotny. Posediffusion: Solving pose estimation via diffusion-aided bundle adjustment. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9773–9783, 2023. 2
- [92] Yongqi Wang, Wenxiang Guo, Rongjie Huang, Jiawei Huang, Zehan Wang, Fuming You, Ruiqi Li, and Zhou Zhao. Frieren: Efficient video-to-audio generation with rectified flow matching. *arXiv preprint arXiv:2406.00320*, 2024. 1, 2, 3, 5
- [93] Yusong Wu, Ke Chen, Tianyu Zhang, Yuchen Hui, Taylor Berg-Kirkpatrick, and Shlomo Dubnov. 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)*, pages 1–5. IEEE, 2023. 5, 6
- [94] Zhifeng Xie, Shengye Yu, Qile He, and Mengtian Li. Sonivisionlm: Playing sound with vision language models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 26866–26875, 2024. 2
- [95] Jiarui Xu, Sifei Liu, Arash Vahdat, Wonmin Byeon, Xiaolong Wang, and Shalini De Mello. Open-vocabulary panoptic segmentation with text-to-image diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2955–2966, 2023. 2
- [96] Manjie Xu, Chenxing Li, Yong Ren, Rilin Chen, Yu Gu, Wei Liang, and Dong Yu. Video-to-audio generation with hidden alignment. *arXiv preprint arXiv:2407.07464*, 2024. 2, 5
- [97] Jinlong Xue, Yayue Deng, Yingming Gao, and Ya Li. Auf-fusion: Leveraging the power of diffusion and large language models for text-to-audio generation. *arXiv preprint arXiv:2401.01044*, 2024. 2
- [98] Karren Yang, Bryan Russell, and Justin Salamon. Telling left from right: Learning spatial correspondence of sight and sound. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9932–9941, 2020. 3
- [99] Jason Y Zhang, Amy Lin, Moneish Kumar, Tzu-Hsuan Yang, Deva Ramanan, and Shubham Tulsiani. Cameras as rays: Pose estimation via ray diffusion. In *International Conference on Learning Representations (ICLR)*, 2024. 2
- [100] Yiming Zhang, Yicheng Gu, Yanhong Zeng, Zhening Xing, Yuancheng Wang, Zhizheng Wu, and Kai Chen. Foleyrafter: Bring silent videos to life with lifelike and synchronized sounds. *arXiv preprint arXiv:2407.01494*, 2024. 2, 5, 6, 7, 1
- [101] Jinxing Zhou, Xuyang Shen, Jianyuan Wang, Jiayi Zhang, Weixuan Sun, Jing Zhang, Stan Birchfield, Dan Guo, Lingpeng Kong, Meng Wang, et al. Audio-visual segmentation with semantics. *arXiv preprint arXiv:2301.13190*, 2023. 2