

# AV-Edit: Multimodal Generative Sound Effect Editing via Audio-Visual Semantic Joint Control

Xinyue Guo<sup>1</sup>, Xiaoran Yang<sup>2</sup>, Lipan Zhang<sup>1</sup>, Jianxuan Yang<sup>1\*</sup>, Zhao Wang<sup>3</sup>, Jian Luan<sup>1</sup>

<sup>1</sup>MiLM Plus, Xiaomi Inc., China

<sup>2</sup>School of Electronic Information, Wuhan University, Wuhan, China

<sup>3</sup>Mi Ecosystem, Xiaomi Inc., China

guoxinyue1@xiaomi.com, yxr888@whu.edu.cn, {zhanglipan,yangjianxuan,wangzhao3,luanjian}@xiaomi.com

## Abstract

Sound effect editing—modifying audio by adding, removing, or replacing elements—remains constrained by existing approaches that rely solely on low-level signal processing or coarse text prompts, often resulting in limited flexibility and suboptimal audio quality. To address this, we propose AV-Edit, a generative sound effect editing framework that enables fine-grained editing of existing audio tracks in videos by jointly leveraging visual, audio, and text semantics. Specifically, the proposed method employs a specially designed contrastive audio-visual masking autoencoder (CAV-MAE-Edit) for multimodal pre-training, learning aligned cross-modal representations. These representations are then used to train an editorial Multimodal Diffusion Transformer (MM-DiT) capable of removing visually irrelevant sounds and generating missing audio elements consistent with video content through a correlation-based feature gating training strategy. Furthermore, we construct a dedicated video-based sound editing dataset as an evaluation benchmark. Experiments demonstrate that the proposed AV-Edit generates high-quality audio with precise modifications based on visual content, achieving state-of-the-art performance in the field of sound effect editing and exhibiting strong competitiveness in the domain of audio generation.

## Extended version and code —

<https://github.com/V2AResearch/AV-Edit>

## Introduction

Recent advances in diffusion-based generative models have driven substantial progress in video editing, enabling diverse and controllable modifications while preserving the creative potential of generative synthesis (Feng et al. 2024; Yang et al. 2025). Meanwhile, video-to-audio (V2A) generation (Luo et al. 2023a; Xu et al. 2024; Gramaccioni et al. 2024) has attracted increasing attention, with research focusing on temporal synchronization and semantic alignment between generated audio and video. Despite these developments, sound effect editing—the task of modifying existing audio by adding, removing, or replacing elements—remains significantly less mature than video editing, particularly

\*corresponding author

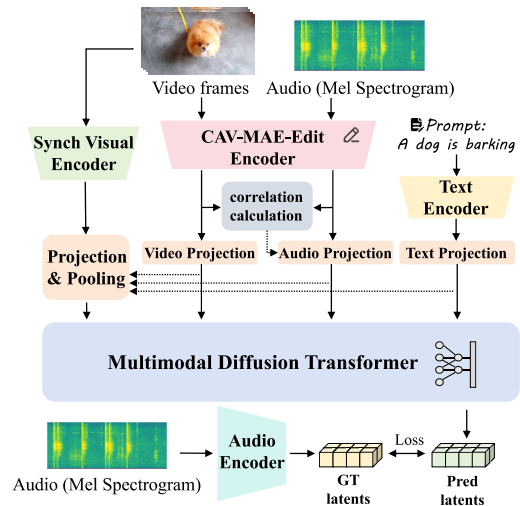


Figure 1: Overview of the AV-Edit framework. The pre-trained CAV-MAE-Edit encoder extracts joint audio-visual features, which are then fed into a multimodal diffusion model to generate the edited audio.

when precise synchronization with visual cues is required. Real-world cases—such as background noise unrelated to the scene contaminating the soundtrack of a live-action video, or objects being added or removed during visual editing without corresponding audio adjustments—underscore the need for seamless audio editing to ensure coherent multimedia experiences.

Traditional sound effect editing methods rooted in signal processing manipulate waveforms, frequency spectra, or time-frequency features using operations such as cropping, splicing, and adaptive filtering. These techniques, while effective for basic modifications, rely heavily on manual parameter tuning and offer limited semantic control. Recent deep learning approaches attempt to overcome these limitations by employing diffusion models to edit audio conditioned on textual prompts (Wang et al. 2023; Manor and Michaeli 2024), where the audio is encoded into a latent space and regenerated under prompt guidance. However, these methods rely on manual adjustments and prior knowledge of signal characteristics, struggling to dynamically syn-

chronize audio edits with visual content while preserving original audio fidelity and audio–visual alignment.

To address these challenges, we introduce AV-Edit, a multimodal generative framework for fine-grained and visually aligned sound effect editing task. AV-Edit is designed to process audio, visual, and textual inputs, leveraging these multimodal cues within a generative model to achieve precise semantic alignment and temporal synchronization between the edited audio and the accompanying video. As illustrated in Figure 1, the framework consists of four primary components: a joint audio–visual encoder (CAV-MAE-Edit), a text encoder (Radford et al. 2021), a visual synchronization encoder (Iashin et al. 2024), and a multimodal diffusion transformer (MM-DiT) (Peebles and Xie 2023). Among these, CAV-MAE-Edit is trained using a combination of contrastive pre-training and masked self-encoding. It learns joint audio–visual representations by projecting both modalities into a shared feature space, and employs mixing-based coding method with cross-attention to visual semantics applied during audio encoding/decoding, producing audio embeddings that are explicitly aligned with the visual content. The MM-DiT then leverages these embeddings, together with an audio–visual correlation-based feature gating training strategy, to guide the audio generation process and perform the desired editing operations.

For evaluation, we construct VGG-Edit, a benchmark dataset derived from the VGGSound (Chen et al. 2020) test set. We select 450 clips and systematically apply three types of edits—sound effect addition, deletion, and replacement—to create standardized evaluation scenarios for video-based sound effect editing. Both objective and subjective evaluations demonstrate that AV-Edit achieves state-of-the-art performance, producing high-quality, visually consistent audio while remaining competitive across multiple metrics in both sound effect editing and audio generation fields.

In summary, our main contributions are as follows:

- A new generative paradigm for sound effect editing. AV-Edit jointly leverages audio, visual, and textual semantics to enable fine-grained, visually synchronized sound effect editing in videos.
- A contrastive audio–visual encoder tailored for sound effect editing. CAV-MAE-Edit integrates contrastive pre-training and masked self-encoding, to learn semantically relevant joint audio-visual representations.
- A benchmark for video-driven sound effect editing to facilitate fair comparison. VGG-Edit is a high-quality benchmark dataset featuring systematically constructed addition, deletion, and replacement scenarios.

## Related Work

### Audio Editing

Research on audio editing spans both traditional signal processing and modern generative methods. Early techniques, such as waveform editing tools (Derry 2012) and the WSOLA algorithm (Verhelst and Roelands 1993), enable operations like cutting, splicing, time-stretching, and pitch-shifting. While suitable for basic modifications, these ap-

proaches rely heavily on manual parameter tuning and cannot synthesize content beyond the original recording. With the advent of deep learning, generative models have introduced a new paradigm for audio editing. Instruction-driven systems (Wang et al. 2023; Vyas et al. 2023; Han et al. 2024) train task-specific generative models to add, replace, or mix audio segments, but require extensive annotated datasets. Other works (Liu et al. 2023; Manor and Michaeli 2024) employ Denoising Diffusion Probabilistic Models (DDPM) and Denoising Diffusion Implicit Models (DDIM) to guide edits from latent noise using text prompts; however, their reliance on precise textual descriptions limits applicability to more flexible or context-driven scenarios. Beyond text-conditioned methods, localized spectrogram-based editing has also been explored. For example, AudioMorphix (Liang et al. 2025a) supports spectrogram modifications but requires target and reference audio, while linguistically guided audio-visual editing (Liang et al. 2024) depends on hand-crafted feature mappings, constraining scalability. Overall, current approaches struggle to achieve fine-grained, context-aware edits that dynamically adapt to video.

### Audio-Visual Pre-training

Video-based sound effect editing requires robust joint representations of audio and video to ensure semantic alignment and temporal synchronization. Foundational works like CLIP (Radford et al. 2021) and CLAP (Wu\* et al. 2023) align text with images and audio, respectively, laying the groundwork for cross-modal alignment. Beyond text-based alignment, several studies focus on video–audio representation learning. For example, Synchformer (Iashin et al. 2024) and DiffAV (Luo et al. 2023b) adopt contrastive learning to obtain temporally and semantically aligned audio–visual features. Other approaches exploit feature fusion to enhance cross-modal understanding. UAVM (Gong et al. 2022a) and VALOR (Liu et al. 2024) integrate complementary information by extracting features from each modality and sharing parameters across specific network layers, thereby promoting modality interaction. In parallel, masked signal modeling serves as a representation learning approach, where a large portion of the input data—such as images or audio—is masked out. This compels the encoder to learn rich semantic representations of the underlying signal in order to reconstruct the masked modality. Recent works (Gong et al. 2022b; Georgescu et al. 2023; Araujo et al. 2025) combine masked autoencoding with contrastive learning to jointly capture modality-specific features and cross-modal relationships, enabling the development of robust audio–visual models suited for generative editing tasks.

### Multimodal Audio Generation

Text-driven approaches (Hung et al. 2024; Lee et al. 2024) synthesize audio from textual descriptions, while video-to-audio approaches (Zhang et al. 2024; Viertola, Iashin, and Rahtu 2025) generate sound effects directly from visual inputs. Some methods, such as SSV2A (Guo et al. 2024), locally perceive multimodal sound sources from a scene with visual detection and cross-modality translation, address V2A generation at the sound-source level. Others pursue

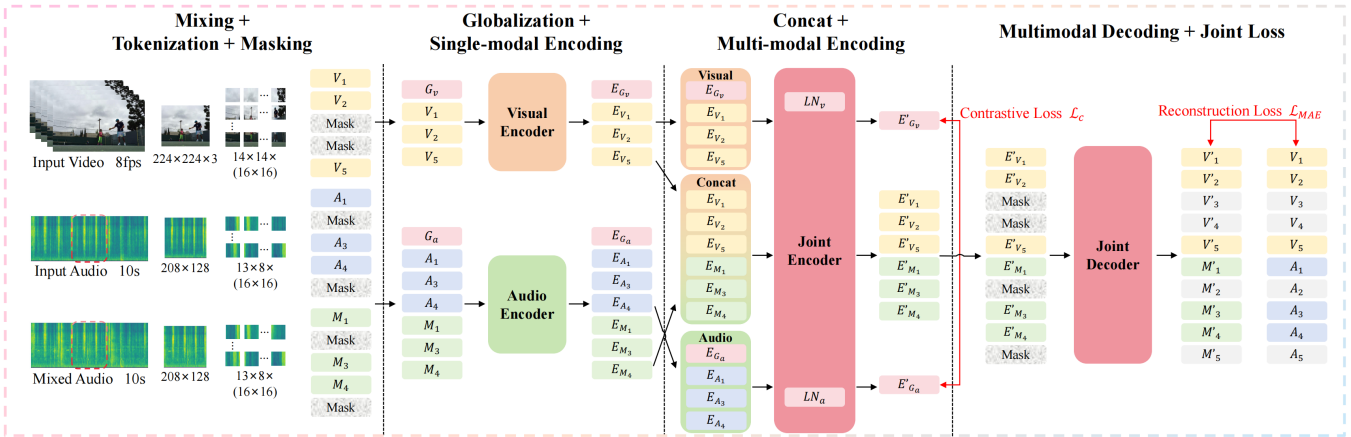


Figure 2: Overview of CAV-MAE-Edit network. The single-modal encoders encode the visual and audio inputs separately. The multi-modal encoder and decoder process the joint embeddings of vision and audio.

fully joint modeling across text, video, and audio: MMAudio (Cheng et al. 2025) and MultiFoley (Chen et al. 2025) integrates all three modalities into Transformer, enable flexible control of audio generation across modalities. From an architectural perspective, autoregressive strategies (Viertola, Iashin, and Rahtu 2025), latent diffusion models (LDMs) and diffusion transformers (DiTs)(Luo et al. 2023b; Cheng et al. 2025), and rectified flow matching(Wang et al. 2024b) have all been explored. While these methods advance multimodal generation, most focus on generating audio from silent videos. Editing sound effects in videos that already contain original audio—requiring both fine-grained control and preservation of temporal and semantic consistency—remains a largely unsolved challenge.

## Methodology

The implementation of AV-Edit consists of two core stages: the pre-training stage of the audio-visual co-encoder CAV-MAE-Edit, and the training stage of the generative sound effect editing model implemented through modal-correlation feature gating and the MMDiT architecture. The overall structure of this framework is shown in Figure 1.

### Editorial Audio-Visual Co-Encoder

To facilitate signal reconstruction through joint semantic learning of visual and auditory inputs while acquiring unified audio-visual representations, we adopt the Contrastive Audio-Visual Masked Autoencoder (CAV-MAE) (Gong et al. 2022b) framework for audio-visual pre-training. Furthermore, departing from conventional approaches, we innovatively introduce spectrogram segmenting and audio mixing strategies: the former can enhance the granularity of audio-visual pairs, while the latter enables the model to extract visually relevant information from the mixed audio based on visual semantics and encode the required information in an editing manner. Figure 2 illustrates the overall framework of our proposed method.

We shall commence with a review of the fundamental principles of CAV-MAE. For a given audio-visual pair, CAV-

MAE initially crops the visual input to a fixed size and extracts the Mel spectrogram of the corresponding audio signal. Subsequently, the multi-modal signals denoted as  $\{v_i, a_i\}$  undergo tokenization to generate corresponding token sequences  $\{V_i, A_i\}$ , which are then projected into the  $\mathbb{R}^{768}$  space via the modality-specific linear projection layers. Meanwhile, the model incorporates 2-D positional embeddings ( $p_v, p_a$ ) and modality-type embeddings ( $m_v, m_a$ ) into the aforementioned patch tokens. Ultimately, 75% of the content for each modality is masked. The input to the encoder is formulated as follows:

$$\begin{aligned} V_i^{\text{unmask}} &= \text{Mask}_{0.75}(\text{Proj}_v(\text{Patch}_v(v_i)) + p_v + m_v), \\ A_i^{\text{unmask}} &= \text{Mask}_{0.75}(\text{Proj}_a(\text{Patch}_a(a_i)) + p_a + m_a). \end{aligned} \quad (1)$$

The sequences of visual and audio tokens are then forwarded through single-modal encoders  $E_v(\cdot)$  and  $E_a(\cdot)$ , respectively, for learning modality-specific representations  $E_{V_i}$  and  $E_{A_i}$ . All single-modal encoders employ the Vision Transformer (ViT) (Dosovitskiy et al. 2021) architecture without sharing weights. Subsequently,  $E_{V_i}$ ,  $E_{A_i}$ , and  $\text{Concat}(E_{V_i}, E_{A_i})$  are fed into a joint audio-visual encoder, yielding  $E'_{V'_i}$ ,  $E'_{A'_i}$ , and  $\text{Concat}(E'_{V'_i}, E'_{A'_i})$ . These three streams incorporate modality-specific layer normalization while sharing weights of other network layers in this multimodal encoder. The model leverages the single-modal outputs  $E'_{V'_i}$  and  $E'_{A'_i}$  for contrastive learning, and adopts the multimodal output  $\text{Concat}(E'_{V'_i}, E'_{A'_i})$  for the reconstruction task. The contrastive loss is defined as:

$$\mathcal{L}_c = -\frac{1}{N} \sum_{i=1}^N \log \left( \frac{\exp(s_{i,i}/\tau)}{\sum_{k \neq i} \exp(s_{i,k}/\tau) + \exp(s_{i,i}/\tau)} \right), \quad (2)$$

where  $s_{i,j} = \left\| \left\| E'_{V'_i} \right\| \right\|^T \left\| E'_{A'_j} \right\|$ , and  $\tau$  is the temperature parameter of the similarity distribution. On the other hand, the reconstruction loss is defined as the mean squared error between the original patches ( $V_i, A_i$ ) and the reconstructed patches ( $V'_i, A'_i$ ), where the latter are the outputs

of  $\text{Concat}(E_{V_i''}, E_{A_i''})$  after the multimodal decoder:

$$\mathcal{L}_{mae} = \frac{1}{N} \sum_{i=1}^N \left( \frac{\sum (V_i^{\text{mask}} - \text{norm}(V_i^{\text{mask}}))^2}{|V_i^{\text{mask}}|} + \frac{\sum (A_i^{\text{mask}} - \text{norm}(A_i^{\text{mask}}))^2}{|A_i^{\text{mask}}|} \right), \quad (3)$$

where  $N$  denotes the batch size, and  $|V_i^{\text{mask}}|$  as well as  $|A_i^{\text{mask}}|$  represent the number of masked visual and audio patches, respectively. Finally, CAV-MAE weights and aggregates the contrastive loss and reconstruction loss to yield the final loss function:

$$\mathcal{L}_{\text{CAV-MAE}} = \lambda_c \mathcal{L}_c + \lambda_{mae} \mathcal{L}_{mae}. \quad (4)$$

With respect to our proposed model, to acquire the target audio-visual co-encoder, we introduce the following improvements within the framework of CAV-MAE.

- Segment the spectrogram to increase the fine-grained nature of audio-visual pairs.** CAV-MAE maps the entire audio clip to a random frame of the video, failing to achieve accurate alignment between audio-visual events. Consequently, we preform frame-wise segmentation on audio-visual pairs, realizing manual temporally alignment on one hand, and ensure the semantic relevance of the two by narrowing the time range on the other. Specifically, frames are extracted from the video at a sampling rate of 8 FPS, where each video frame is paired with a 0.125s audio segment cropped from the full Mel spectrogram at the corresponding temporal position. This approach enables fine-grained segmentation of the entire audio while preserving a certain level of semantic information, thereby achieving accurate alignment with the corresponding visual information. Let  $T$  denote the total number of video frames, and  $L$  represent the total length of the Mel spectrogram. The temporal center of the audio segment corresponding to the  $i$ -th frame is given by  $i * L/T$ . Given that the duration of the audio segment for each frame is fixed as  $l$ , the start and end positions of the audio segment corresponding to the  $i$ -th frame within the Mel spectrogram are defined as  $l_{start} = i * L/T - l/2$  and  $l_{end} = i * L/T + l/2$ , respectively. This method enables the encoder to be trained with more precisely aligned audio-visual pairs.
- Mixing to compel the model to learn audio-visual related semantics.** To encode features of audio segments that are relevant to visual semantics and achieve the desired editing effect, we mix the target audio with irrelevant audio, as detailed in Appendix B1. As shown in Figure 2, we mask the same positions of the original and mixed audio, concatenate their outputs, and feed the result into an audio encoder. Subsequently, the mixed audio component within the audio embeddings is concatenated with the visual embeddings and input to the joint blocks of the multimodal encoder, where inter-modal information interaction is performed through Attention and MLP layers to derive semantically relevant audio-video embeddings. At the same time, the original audio embeddings, together with the duplicated visual embeddings,

are encoded through the single-modal blocks within the multimodal encoder. This entire process can be represented as:

$$\begin{aligned} (e'_V, e'_M) &= E_j(\text{concat}(e_V, e_M)); LN1_{av}; LN2_{av}), \\ e''_V &= E_v(e_V; LN1_v; LN2_v), \\ e'_A &= E_a(e_A; LN1_a; LN2_a). \end{aligned} \quad (5)$$

We input  $(e'_V, e'_M)$  into the joint decoder to reconstruct data of each modality. For the reconstruction loss, we calculate it using the decoded audio-visual patches and the original audio-visual patches.

- Adding modal related global tokens to focus on global information.** Referring to (Araujo et al. 2025), we similarly introduce the global tokens  $G_v$  and  $G_a$ , which continuously aggregate information during the single-modal encoding and multimodal encoding phases, focusing on the global representation of the respective modality. These encoded global tokens  $E'_{G_v}$  and  $E'_{G_a}$  are also ultimately used to calculate the contrastive loss.

## Generative Sound Effect Editing

With the aforementioned editorial audio-visual co-encoder in place, we next employ a multimodal diffusion model for generative sound effect editing. AV-Edit, as a generative sound effect editing model, can be divided into two parts: feature extraction and diffusion-based generation, which will be elaborated in detail below.

**Feature Extraction** Audio and visual semantic embeddings  $F_a$  and  $F_v$  incorporating both temporal and spatial dimensions are extracted via the pre-trained co-encoder CAV-MAE-Edit. These embeddings are flattened along the temporal dimension and pooled across the spatial dimension. Subsequently, a correlation-based feature gating strategy is adopted to compute the cosine similarity between audio and visual embeddings frame-by-frame, thereby determining the degree of semantic alignment between the current image and audio features. This mechanism enables the determination of whether to feed the audio features of the current frame into the generative network. The proposed algorithm retains audio embeddings whose similarity scores exceed a predefined threshold  $r$ , while substituting semantically irrelevant audio features with empty features. The textual feature  $F_t$  is extracted using CLIP, and the visual synchronization feature  $F_{sync}$  is obtained through Synchformer to facilitate the generation of video-synchronized audio. Thereafter, the visual-audio-textual semantic features are projected onto the pre-defined dimensions by their respective projection layers; on the one hand, they serve as inputs to the diffusion model, and on the other hand, they undergo pooling across time and are summed with the temporal embedding to form global control information  $g_c$ . Furthermore, the visual synchronization features, after projection, are added to  $g_c$  to form the temporal synchronization information  $f_c$ , which is incorporated into the generation process. During the training process of AV-Edit, all these encoders are frozen, and the projection layer will be trained from scratch.

Method			Metrics			
Segmenting	Global Token	Mixing	Audio MAE Loss	Visual MAE Loss	Constrastive Loss	Total Loss
×	×	×	2.219	1.335	1.484	3.569
✓	×	×	0.450	1.097	8.552	1.632
✓	✓	×	<b>0.400</b>	0.942	<b>0.807</b>	1.350
✓	✓	✓	0.661	<b>0.619</b>	1.701	<b>1.297</b>

Table 1: Comparison of the losses for CAV-MAE-Edit under different methods.

**Diffusion Generation** Aiming to fully leverage multi-modal information, and thus, similar to MMAudio, we adopt the MM-DiT block from (Esser et al. 2024). The detailed structure of the generative model component in AV-Edit is provided in Appendix A1. Within the diffusion model, initially, the target audio is encoded into latent space as  $x$  by a pre-trained VAE encoder, followed by forward diffusion. Correspondingly, the VAE decoder leverages the backward denoising process to obtain the output  $\hat{x}$ , namely the high-dimensional representation retrieved from the latent space. During the forward diffusion, Gaussian noise is added to the latent  $x_0$  of the clean audio to obtain noisy latent  $x_t$  at different time steps ( $t \in \{1, \dots, T\}$ ). In step  $t$ ,  $x_t = \sqrt{\alpha_t} \cdot x_{t-1} + \sqrt{1 - \alpha_t} \cdot \epsilon_t$ , where  $\epsilon_t \sim \mathcal{N}(0, \mathbf{I})$  is Gaussian noise and  $\alpha_t$  is a scale parameter. Correspondingly, the inverse process is an iterative denoising of pure noise  $x_T$  to recover the original data  $x_0$ . It is well-known that DiT adopts Transformer as the backbone network of the diffusion model to modeling the noise prediction function  $\epsilon_\theta$ . Furthermore, the multimodal diffusion model we use incorporates control conditions into the noise prediction process. Specifically, it concatenates the queries, keys, and values of the visual-audio-textual embeddings with those of the latent  $x_t$ , and then fuses features from these modalities through joint attention. The predicted latent can be expressed as  $\epsilon_\theta(x_t, t, (f_a, f_v, f_t))$  and the network is trained with the goal:

$$\mathcal{L}(\theta) = \mathbb{E}_{x_0, \epsilon, t} [\|\epsilon - \epsilon_\theta(x_t, t, (f_a, f_v, f_t))\|^2], \quad (6)$$

where  $\epsilon \sim \mathcal{N}(0, \mathbf{I})$  and  $t \sim \text{Uniform}(1, T)$ . Detailed implementation specifications for the remaining modules of the generative network are provided in Appendix A2. Furthermore, the Classifier-Free Guidance (CFG) technique is incorporated during the training process. Specifically, in the training phase, the visual embedding  $f_v$ , audio embedding  $f_a$ , and text embedding  $f_t$  are randomly masked with learnable empty embeddings  $\emptyset_v$ ,  $\emptyset_a$ , and  $\emptyset_t$  respectively, with a masking probability of 10%. The outputs of the model under conditional and unconditional settings is balanced by a guidance scale  $s$ , and the adjusted predicted latent variable is formulated as follows:

$$\begin{aligned} \hat{\epsilon}_\theta(x_t, t, (f_a, f_v, f_t)) &= s \cdot \epsilon_\theta(x_t, t, (f_a, f_v, f_t)) \\ &+ (1 - s) \epsilon_\theta(x_t, t, (\emptyset_a, \emptyset_v, \emptyset_t)). \end{aligned} \quad (7)$$

Classifier-Free Guidance serves as an effective approach to enhancing generation performance and editing capabilities, while enabling a balanced trade-off between the quality and diversity of the final output.

## Experiments

In this section, we perform comprehensive experiments and present various evaluation results.

### Audio-Visual Co-Encoder CAV-MAE-Edit

The proposed CAV-MAE-Edit undergoes pre-training to acquire joint audio-visual representations and the signal reconstruction ability corresponding to each modality.

**Dataset** We utilize the complete VGGSound dataset for training the proposed encoder. VGGSound contains 200K 10-second videos spanning 309 categories, with each video containing a semantically explicit audio-visual pair. For each video, we extract video frames at a sampling rate of 8 FPS and the complete audio at a sampling rate of 16 kHz. Eventually, 80 frames of video and a 128-dimensional Mel spectrogram with a length of 1024 are obtained.

**Implementation Details** We set the length of the Mel spectrogram segment corresponding to each video frame to 208, and slice the audio to obtain fine-grained visual-audio pairs. The video frame and audio clip are then segmented with a patch size of 16, resulting in  $14 \times 14$  blocks of visual embeddings and  $13 \times 8$  blocks of audio embeddings. Audio mixing is performed with a probability of 0.5, while audio and visual embeddings are masked with a probability of 0.5, resulting in three sets of data: mixed-masked audio, masked clean audio and masked visual embeddings. These three streams are fed into their respective encoders. Training is conducted on four NVIDIA H800 GPUs with a learning rate of  $1e-4$  and a batch size of 160 until convergence, which takes approximately 150 epochs. Furthermore, the weights parameters of the loss function in Equation 4 are set to  $\lambda_c = 0.01$  and  $\lambda_{mae} = 1$ , respectively. The variation of various loss values during training is illustrated in Figure B of Appendix B2.

**Ablation** To explore the effectiveness of the improvements proposed in this work for the audio-visual co-encoder with respect to CAV-MAE, we conducted ablation experiments under equivalent conditions. Table 1 shows the impacts of spectrogram segmenting, adding global tokens, and audio mixing on the model. It can be observed that spectrogram segmentation significantly reduces the Audio MAE loss, while the incorporation of global tokens enhances the model’s capability in contrastive learning. Although the introduction of audio mixing increases the audio MSE loss, it enables the model to perform audio encoding in an edit-oriented manner while minimizing the overall model loss.

Method	Distribution matching		Audio quality	Semantic align	Temporal align	Efficiency	
	KL <sub>PANN<sub>s</sub></sub> ↓	KL <sub>PaSST</sub> ↓	IS ↑	IB ↑	DeSync ↓	Params	Time(s)
Seeing&Hearing	2.92	2.93	6.19	<b>33.85</b>	1.193	415M	30.18
FoleyCrafter	2.30	2.25	15.67	25.75	1.222	1.22B	3.71
V-AURA*	2.44	2.08	10.29	27.64	0.813	695M	33.63
Frieren*	3.36	3.59	8.37	17.31	1.259	159M	2.68
V2A-Mapper* $\diamond$	2.69	2.55	12.47	22.55	1.225	229M	-
AudioX	1.99	1.86	16.82	25.86	1.237	1.10B	13.00
MMAudio-S	<b>1.63</b>	1.60	14.53	28.23	0.505	157M	2.56
MMAudio-L	<b>1.63</b>	1.46	18.13	<u>32.73</u>	<u>0.465</u>	1.03B	3.01
AV-Edit-S (ours)	1.71	1.61	20.38	28.08	0.480	178M	2.69
AV-Edit-L (ours)	<u>1.67</u>	<b>1.41</b>	<u>22.48</u>	31.68	<b>0.441</b>	1.15B	3.18
AV-Edit w/o Audio	1.73	<u>1.45</u>	<b>23.08</b>	30.73	0.481	1.15B	3.16

Table 2: Audio Generation results on the VGGSound test set. The number of parameters excludes all pre-trained encoders. All samples were generated on an H800 GPU. \*: does not use text input during testing.  $\diamond$ : evaluated using generation samples obtained directly from the authors. The best result for each indicator is bolded and the second best result is underlined.

### Audio Generation Part of AV-Edit

After completing the pre-training of the encoder, the audio is next generatively edited using the model AV-Edit.

**Experimental Setup** In this part, we use the datasets VGGSound, Audiocaps (Kim et al. 2019), Wavcaps (Mei et al. 2024), and Clotho (Drossos, Lipping, and Virtanen 2020) to train the model. Audio, visual, and textual features are extracted from these datasets, and for modalities lacking in the datasets, we set their features to learnable null tokens such as  $\varnothing_a$ ,  $\varnothing_v$  and  $\varnothing_t$ . For the aforementioned audio-visual similarity threshold  $r$ , we set it to 0.3. The detailed process of threshold selection is provided in Appendix A3. The structure of the diffusion component, the same as in MMAudio, is divided into two types: small model and large model. The small model contains 4 MM-DiT blocks and 8 Single-Modal DiT blocks, generating 16 kHz audio encoded as 20-dimensional latent, while the large model contains 7 MM-DiT blocks and 14 Single-Modal DiT blocks, generating 44.1 kHz audio encoded as 40-dimensional latent. Using 25 inference steps and setting the classifier-free guidance strength to 4.5, the model is trained on four NVIDIA H800 GPUs for about 43 epochs.

**Evaluation Metrics** The generated audio is evaluated using distribution matching, audio quality, semantic alignment, and temporal alignment. Distribution matching is measured using the Kullback-Leibler (KL) distance of the features extracted by the PANNs (Kong et al. 2020) and PaSST (Koutini et al. 2022) models. Audio quality is evaluated using the Inception Score, while the semantic alignment score IB and the temporal alignment score DeSync are calculated using ImageBind (Girdhar et al. 2023) and Synchformer, respectively. Appendix C1 details how these evaluation indicators were calculated.

**Quantitative Results** We evaluate the generative performance of the model using the test set of VGGSound, and the comparison models contained representative V2A models Seeing and Hearing (Xing et al. 2024), FoleyCrafter (Zhang et al. 2024), V-AURA (Viertola, Iashin, and Rahtu 2025),

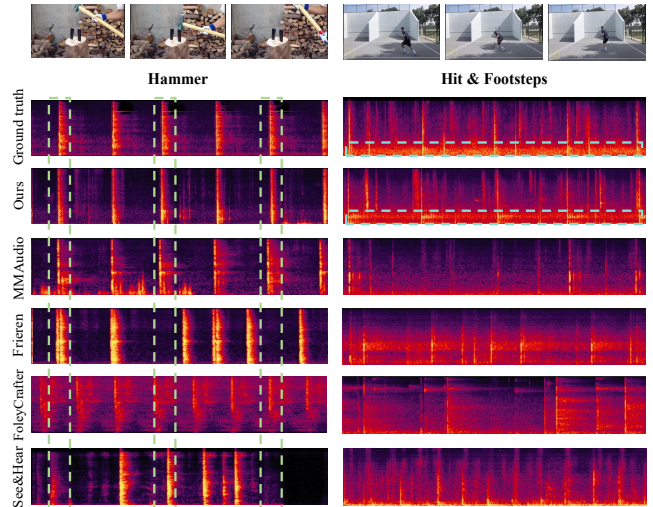


Figure 3: The spectrograms of generated audios.

Frieren (Wang et al. 2024b), V2A-Mapper (Wang et al. 2024a), AudioX (Tian et al. 2025), and MMAudio. Unlike other models, the input of our model AV-Edit includes not only visual-textual features but also raw audio features, which, however, interferes with audio generation to some extent. The above models are used to generate audios with 8 seconds for testing, and the results are shown in Table 2. Our model is overwhelmingly superior in terms of generated audio quality, with top performance in distribution matching and temporal alignment, and semantic alignment second only to Seeing&Hearing and MMAudio. It is worth noting that, as a generative sound effect editing model, our model has already outperformed existing state-of-the-art models in most key metrics and ranked the best when evaluated solely from the perspective of audio generation without the input of audio features. For a better comparison, the visualization results are shown in Figure 3 and analyzed in Appendix C2.

## Sound effect Editing Based on Video Content

Lastly, we explore the editing performance of our proposed framework in detail, focusing on the addition, removal, and replacement of original audio tracks in videos.

**Dataset VGG-Edit** To evaluate the model’s capability of video-based sound effect editing, we design a dedicated dataset called VGG-Edit based on the test set of VGGSound. To be specific, we manually select 450 visually meaningful videos from approximately 15,000 test videos. For 150 of these videos, we add sound effects irrelevant to the visual semantics to their audio tracks; for another 150 videos, we remove the sound effects relevant to the visual semantics from their audio tracks; and for the remaining 150 videos, we replace the original visually semantic-relevant sound effects in their audio tracks with sound effects that are irrelevant.

**Objective Results** We select the V2A models AudioX and DeepSound (Liang et al. 2025b) that support audio track input as comparative methods, and conduct tests on the proposed sound effect editing dataset, aiming to perform operations of adding, deleting, and replacing the audio tracks of the input videos. For AudioX, we input the videos with the original audio tracks and the specified textual descriptions. DeepSound handles this slightly differently: in the add and replace tasks, the videos with the original audio tracks and the specified textual descriptions are fed into step 1 of the model, while in the delete task, we skip the first step of the model and feed the aforementioned inputs into the second step of the model to separate the sound effects that are not related to the visual semantics. After extracting the features of the output audio and target audio using PANNs, we use the KL distance and Fréchet (Fréchet 1906) distance to evaluate the objective metrics of the editing performance of the aforementioned models, and the results are shown in Table 3.

Method	Addition		Removal		Replacement	
	FD↓	KL↓	FD↓	KL↓	FD↓	KL↓
VGG-Edit	131.63	5.86	121.07	5.74	135.96	6.12
AudioX w/ T	47.76	2.60	51.71	2.86	79.64	3.66
DeepSound w/ T	40.69	2.53	75.16	3.63	126.10	5.91
AV-Edit w/o T	36.05	1.56	27.56	0.99	48.24	2.13
AV-Edit w/ T	<b>31.04</b>	<b>1.39</b>	<b>25.24</b>	<b>0.95</b>	<b>38.00</b>	<b>1.54</b>

Table 3: The objective results of various sound effect editing methods on the VGG-Edit dataset.

Our method outperforms others across all metrics, enabling purposeful sound effect addition, removal, and replacement to approximate the target audio. Compared to AudioX and DeepSound, it achieves significant improvements in removal and replacement tasks, with FD and KL distances more aligned with the target. Additionally, textual instructions enhance sound effect processing effectiveness.

**Subjective Results** In terms of subjective evaluation, we use the metrics proposed by the audio editing model AudioMorphix: fidelity (F), perceptual quality (PQ), consistency (C), regional specificity (RS), and instruction adherence (IA). The specific meanings and scoring methods of

Method	F↑	PQ↑	C↑	RS↑	IA↑
AudioX w/ T	42.2	49.1	49.7	52.5	58.8
DeepSound w/ T	47.3	39.8	46.9	48.0	36.7
AV-Edit w/o T	<b>64.4</b>	58.8	63.2	63.3	68.0
AV-Edit w/ T	62.4	<b>66.2</b>	<b>64.2</b>	<b>64.7</b>	<b>70.8</b>

Table 4: The subjective results of various sound effect editing methods on the VGG-Edit dataset.

these metrics are given in Appendix D1. Table 4 demonstrates the superiority of our method AV-Edit, which can well refer to the features of the original video’s audio track, perform sound effect addition or deletion operations based on visual semantics. The textual feature input has a guiding effect on the model, which improves the region specificity and instruction adherence, etc. of the edited audio with a slight reduction in fidelity. Figure 4 shows the visualization results of sound effect editing performed by AV-Edit and analyzed in Appendix D2. For example, it can be seen from the middle part of Figure 4 that the edited audio removes the sound effect associated with the eliminated image from the original track, and the retained portion is enhanced. At last, More ablation experiments can be seen in Appendix E.

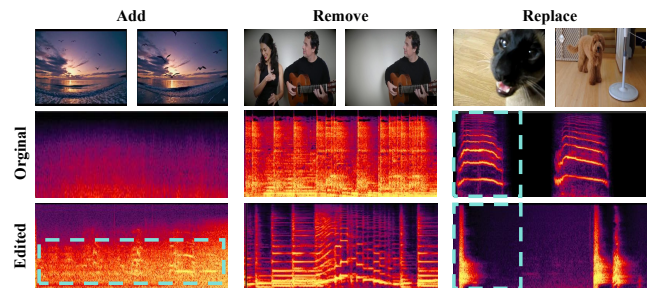


Figure 4: Examples of three editing manipulations: add, remove and replace.

## Conclusion

In this paper, we propose AV-Edit, a groundbreaking generative sound effect editing framework for fine-grained editing of existing audio tracks in videos. With a specially designed audio-visual co-encoder and a multimodal training strategy that introduces correlation-based feature gating, our approach achieves superior sound editing performance, removes visually irrelevant sounds, and generates missing sound effect elements. For evaluation, we construct VGG-Edit, a benchmark dataset for video-based sound effect editing. Through comprehensive experiments, we demonstrate that AV-Edit outperforms existing audio generation models and, simultaneously, fills the gap in video-based sound effect editing models. In addition, our approach also has a limitation - it does not preserve the original audio completely without distortion - and we hope to address this pain point in future work.

## Acknowledgements

This work was supported by the MiLM Plus, Xiaomi Inc.

## References

- Araujo, E.; Rouditchenko, A.; Gong, Y.; Bhati, S.; Thomas, S.; Kingsbury, B.; Karlinsky, L.; Feris, R.; Glass, J. R.; and Kuehne, H. 2025. CAV-MAE Sync: Improving Contrastive Audio-Visual Mask Autoencoders via Fine-Grained Alignment. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 18794–18803.
- Chen, H.; Xie, W.; Vedaldi, A.; and Zisserman, A. 2020. Vg-gsound: A large-scale audio-visual dataset. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 721–725. IEEE.
- Chen, Z.; Seetharaman, P.; Russell, B.; Nieto, O.; Bourgin, D.; Owens, A.; and Salamon, J. 2025. Video-guided foley sound generation with multimodal controls. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 18770–18781.
- Cheng, H. K.; Ishii, M.; Hayakawa, A.; Shibuya, T.; Schwing, A.; and Mitsufuji, Y. 2025. MMAudio: Taming Multimodal Joint Training for High-Quality Video-to-Audio Synthesis. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 28901–28911.
- Derry, R. 2012. *PC audio editing with Adobe Audition 2.0: Broadcast, desktop and CD audio production*. Routledge.
- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; Uszkoreit, J.; and Houlsby, N. 2021. An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. In *International Conference on Learning Representations*.
- Drossos, K.; Lipping, S.; and Virtanen, T. 2020. Clotho: An audio captioning dataset. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 736–740. IEEE.
- Esser, P.; Kulal, S.; Blattmann, A.; Entezari, R.; Müller, J.; Saini, H.; Levi, Y.; Lorenz, D.; Sauer, A.; Boesel, F.; et al. 2024. Scaling rectified flow transformers for high-resolution image synthesis. In *Forty-first international conference on machine learning*.
- Feng, R.; Weng, W.; Wang, Y.; Yuan, Y.; Bao, J.; Luo, C.; Chen, Z.; and Guo, B. 2024. CCEdit: Creative and Controllable Video Editing via Diffusion Models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 6712–6722.
- Fréchet, M. 1906. Sur quelques points du calcul fonctionnel. *Georgescu, M.-I.; Fonseca, E.; Ionescu, R. T.; Lucic, M.; Schmid, C.; and Arnab, A.* 2023. Audiovisual masked autoencoders. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 16144–16154.
- 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.
- Gong, Y.; Liu, A. H.; Rouditchenko, A.; and Glass, J. 2022a. Uavm: Towards unifying audio and visual models. *IEEE Signal Processing Letters*, 29: 2437–2441.
- Gong, Y.; Rouditchenko, A.; Liu, A. H.; Harwath, D.; Karlinsky, L.; Kuehne, H.; and Glass, J. 2022b. Contrastive audio-visual masked autoencoder. *arXiv preprint arXiv:2210.07839*.
- Gramaccioni, R. F.; Marinoni, C.; Postolache, E.; Comunità, M.; Cosmo, L.; Reiss, J. D.; and Comminiello, D. 2024. Stable-v2a: Synthesis of synchronized sound effects with temporal and semantic controls. *CoRR*.
- Guo, W.; Wang, H.; Ma, J.; and Cai, W. 2024. Gotta Hear Them All: Sound Source Aware Vision to Audio Generation. *arXiv preprint arXiv:2411.15447*.
- Han, B.; Dai, J.; Hao, W.; He, X.; Guo, D.; Chen, J.; Wang, Y.; Qian, Y.; and Song, X. 2024. InstructME: An Instruction Guided Music Edit Framework with Latent Diffusion Models. In Larson, K., ed., *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence, IJCAI-24*, 5835–5843. International Joint Conferences on Artificial Intelligence Organization. Main Track.
- Hung, C.-Y.; Majumder, N.; Kong, Z.; Mehrish, A.; Zadeh, A.; Li, C.; Valle, R.; Catanzaro, B.; and Poria, S. 2024. TangoFlux: Super Fast and Faithful Text to Audio Generation with Flow Matching and Clap-Ranked Preference Optimization. *arXiv:2412.21037*.
- Iashin, V.; Xie, W.; Rahtu, E.; and Zisserman, A. 2024. Synchformer: Efficient synchronization from sparse cues. In *ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 5325–5329. IEEE.
- Kim, C. D.; Kim, B.; Lee, H.; and Kim, G. 2019. Audio-caps: Generating captions for audios in the wild. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, 119–132.
- Kong, Q.; Cao, Y.; Iqbal, T.; Wang, Y.; Wang, W.; and Plumbley, M. D. 2020. Panns: Large-scale pretrained audio neural networks for audio pattern recognition. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 28: 2880–2894.
- Koutini, K.; Schlüter, J.; Eghbal-zadeh, H.; and Widmer, G. 2022. Efficient Training of Audio Transformers with Patchout. In *Interspeech 2022, 23rd Annual Conference of the International Speech Communication Association, Incheon, Korea, 18-22 September 2022*, 2753–2757. ISCA.
- Lee, S.-g.; Kong, Z.; Goel, A.; Kim, S.; Valle, R.; and Catanzaro, B. 2024. ETTA: Elucidating the Design Space of Text-to-Audio Models. *arXiv preprint arXiv:2412.19351*.
- Liang, J.; Chen, Y.; Yuan, Y.; Jia, D.; Zhuang, X.; Chen, Z.; Wang, Y.; and Wang, Y. 2025a. AudioMorphix: Training-free audio editing with diffusion probabilistic models. *arXiv preprint arXiv:2505.16076*.
- Liang, S.; Huang, C.; Tian, Y.; Kumar, A.; and Xu, C. 2024. Language-Guided Joint Audio-Visual Editing via One-Shot

- Adaptation. In *Proceedings of the Asian Conference on Computer Vision (ACCV)*, 1011–1027.
- Liang, Y.; Chen, Z.; Ding, C.; and Di, X. 2025b. DeepSound-V1: Start to Think Step-by-Step in the Audio Generation from Videos. *arXiv preprint arXiv:2503.22208*.
- Liu, H.; Chen, Z.; Yuan, Y.; Mei, X.; Liu, X.; Mandic, D.; Wang, W.; and Plumbley, M. D. 2023. AudioLDM: Text-to-Audio Generation with Latent Diffusion Models. *Proceedings of the International Conference on Machine Learning*, 21450–21474.
- Liu, J.; Chen, S.; He, X.; Guo, L.; Zhu, X.; Wang, W.; and Tang, J. 2024. Valor: Vision-audio-language omni-perception pretraining model and dataset. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- Luo, S.; Yan, C.; Hu, C.; and Zhao, H. 2023a. Diff-Foley: Synchronized Video-to-Audio Synthesis with Latent Diffusion Models. In Oh, A.; Naumann, T.; Globerson, A.; Saenko, K.; Hardt, M.; and Levine, S., eds., *Advances in Neural Information Processing Systems*, volume 36, 48855–48876. Curran Associates, Inc.
- Luo, S.; Yan, C.; Hu, C.; and Zhao, H. 2023b. Diff-foley: Synchronized video-to-audio synthesis with latent diffusion models. *Advances in Neural Information Processing Systems*, 36: 48855–48876.
- Manor, H.; and Michaeli, T. 2024. Zero-Shot Unsupervised and Text-Based Audio Editing Using DDPM Inversion. In Salakhutdinov, R.; Kolter, Z.; Heller, K.; Weller, A.; Oliver, N.; Scarlett, J.; and Berkenkamp, F., eds., *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, 34603–34629. PMLR.
- Mei, X.; Meng, C.; Liu, H.; Kong, Q.; Ko, T.; Zhao, C.; Plumbley, M. D.; Zou, Y.; and Wang, W. 2024. Wavcaps: A chatgpt-assisted weakly-labelled audio captioning dataset for audio-language multimodal research. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 32: 3339–3354.
- Peebles, W.; and Xie, S. 2023. Scalable diffusion models with transformers. In *Proceedings of the IEEE/CVF international conference on computer vision*, 4195–4205.
- 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.
- Tian, Z.; Jin, Y.; Liu, Z.; Yuan, R.; Tan, X.; Chen, Q.; Xue, W.; and Guo, Y. 2025. Audiox: Diffusion transformer for anything-to-audio generation. *arXiv preprint arXiv:2503.10522*.
- Verhelst, W.; and Roelands, M. 1993. An overlap-add technique based on waveform similarity (WSOLA) for high quality time-scale modification of speech. In *1993 IEEE International Conference on Acoustics, Speech, and Signal Processing*, volume 2, 554–557. IEEE.
- Viertola, I.; Iashin, V.; and Rahtu, E. 2025. Temporally aligned audio for video with autoregression. In *ICASSP 2025-2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 1–5. IEEE.
- Vyas, A.; Shi, B.; Le, M.; Tjandra, A.; Wu, Y.-C.; Guo, B.; Zhang, J.; Zhang, X.; Adkins, R.; Ngan, W.; et al. 2023. Audiobox: Unified audio generation with natural language prompts. *arXiv preprint arXiv:2312.15821*.
- Wang, H.; Ma, J.; Pascual, S.; Cartwright, R.; and Cai, W. 2024a. 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.
- Wang, Y.; Guo, W.; Huang, R.; Huang, J.; Wang, Z.; You, F.; Li, R.; and Zhao, Z. 2024b. Frieren: Efficient video-to-audio generation network with rectified flow matching. *Advances in Neural Information Processing Systems*, 37: 128118–128138.
- Wang, Y.; Ju, Z.; Tan, X.; He, L.; Wu, Z.; Bian, J.; et al. 2023. Audit: Audio editing by following instructions with latent diffusion models. *Advances in Neural Information Processing Systems*, 36: 71340–71357.
- 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 *IEEE International Conference on Acoustics, Speech and Signal Processing, ICASSP*.
- 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.
- Xu, M.; Li, C.; Tu, X.; Ren, Y.; Chen, R.; Gu, Y.; Liang, W.; and Yu, D. 2024. Video-to-audio generation with hidden alignment. *arXiv preprint arXiv:2407.07464*.
- Yang, X.; Zhu, L.; Fan, H.; and Yang, Y. 2025. Videograin: Modulating space-time attention for multi-grained video editing. In *The Thirteenth International Conference on Learning Representations*.
- Zhang, Y.; Gu, Y.; Zeng, Y.; Xing, Z.; Wang, Y.; Wu, Z.; and Chen, K. 2024. Foleyrafter: Bring silent videos to life with lifelike and synchronized sounds. *arXiv preprint arXiv:2407.01494*.