

## Zero-Shot Audio-Visual Editing via Cross-Modal Delta Denoising

Yan-Bo Lin<sup>1\*</sup>   Kevin Lin<sup>2</sup>   Zhengyuan Yang<sup>2</sup>   Linjie Li<sup>2</sup>   Jianfeng Wang<sup>2</sup>  
 Chung-Ching Lin<sup>2</sup>   Xiaofei Wang<sup>2</sup>   Gedas Bertasius<sup>1</sup>   Lijuan Wang<sup>2</sup>  
<sup>1</sup>UNC Chapel Hill   <sup>2</sup>Microsoft

### Abstract

In this paper, we introduce zero-shot audio-video editing, a novel task that requires transforming original audio-visual content to align with a specified textual prompt without additional model training. To evaluate this task, we curate a benchmark dataset, AVED-Bench, designed explicitly for zero-shot audio-video editing. AVED-Bench includes 110 videos, each with a 10-second duration, spanning 11 categories from VGGSound. It offers diverse prompts and scenarios that require precise alignment between auditory and visual elements, enabling robust evaluation. We identify limitations in existing zero-shot audio and video editing methods, particularly in synchronization and coherence between modalities, which often result in inconsistent outcomes. To address these challenges, we propose AVED, a zero-shot cross-modal delta denoising framework that leverages audio-video interactions to achieve synchronized and coherent edits. AVED demonstrates superior results on both AVED-Bench and the recent OAVE dataset to validate its generalization capabilities. Our codes, data, and results are available at [https://genjib.github.io/project\\_page/AVED/index.html](https://genjib.github.io/project_page/AVED/index.html)

### 1. Introduction

Recent advancements in diffusion-based generative models [12, 20, 56, 59, 73] have demonstrated remarkable progress in image [57, 63–65, 67], video [4, 21, 24, 71, 83], music [25, 41, 42, 45, 69, 77], and audio [9, 10, 26, 46, 48, 52] generation. While these models deliver impressive quality, the development of adaptable and controllable generative models for real-world applications remains challenging. This challenge stems from the difficulty of disentangling specific attributes within diffusion models, limiting fine-grained control over generated content. To address the limitations of controllability in generative models, recent work [1, 5, 35, 37, 54, 68, 91] has focused on enhancing precision and flexibility in content creation. These models enable users to create and edit images with fine-grained control, supporting applications like photo editing and per-

\*Work done during an internship at Microsoft.

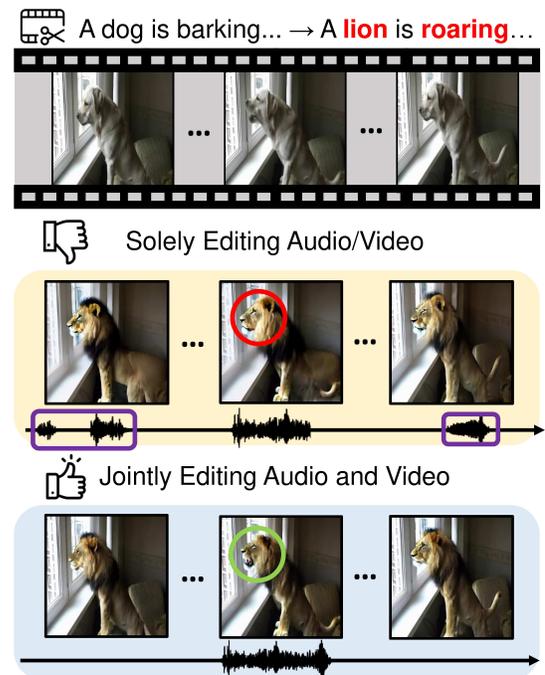


Figure 1. **Key Challenges in Joint Audio-Video Editing.** Existing methods primarily focus on zero-shot text-to-video [11, 50, 86] or text-to-audio [31, 53, 82] editing separately. Solely editing only video or only audio often leads to coherence and synchronization issues between two modalities. As highlighted in red circle, the motion or presence of sounding objects may not align with the corresponding audio. Additionally, edited content may exhibit audio artifacts along the temporal dimension (shown in the purple squares). These factors make the edited results feel less natural and cohesive. In contrast, our AVED jointly edits audio and video by leveraging cross-modal information as additional supervision to improve editing quality to alleviate synchronization issues.

sonalized content generation. While these approaches have advanced image-based editing, achieving seamless and synchronized edits across both audio and video modalities remains challenging. However, such capabilities are increasingly crucial for content creators, filmmakers, or digital artists, who require intuitive tools to edit and modify multimedia content efficiently. To illustrate the challenges of audio-video editing, consider a scenario in Figure 1 where a

Method	Venue	# Videos	Modality
DreamMotion [28]	ECCV'24	26	Video
RAVE [32]	CVPR'24	31	Video
TokenFlow [16]	ICLR'24	61	Video
AVED-Bench	N/A	110	Video+Audio

Table 1. **Existing Evaluation Sets in Video-Based Zero-Shot Editing.** Unlike prior video-only benchmarks, AVED-Bench introduces both video and audio, which is more challenging and enables a comprehensive evaluation of zero-shot audio-video editing.

dog is barking in a room. Suddenly, the dog transforms into a fierce lion with a corresponding roar, enhancing suspense and surprise. This scenario requires a model capable of not only transforming the visual appearance of the dog into a lion but also synchronously updating the audio to match the new visual context—tasks that current single-modality editing approaches [16, 28, 31, 32, 53, 70, 72] struggle to accomplish.

Efficiently and seamlessly editing real-world audio-video content without substantial computational overhead remains challenging. It typically relies on additional datasets [21, 23, 71, 82] or finetuning on pretrained text-guided models [2, 33, 40]. To reduce these costs, recent zero-shot methods have been proposed for video [1, 19, 32, 34, 61, 78, 80] and audio editing [31, 53] by leveraging pretrained text-to-image [13, 65, 91] and text-to-audio [14, 46, 48] diffusion models. However, existing approaches are limited to single modalities and lack frameworks designed for joint audio-video editing, highlighting a gap in multimodal editing models and benchmarks.

To mitigate this gap, we introduce the AVED-Bench dataset, manually curated from VGGSound [8]. Unlike video-only datasets in Table 1, AVED-Bench contains a rich diversity of natural audio-visual events. AVED-Bench consists of 110 distinct 10-second videos across 11 categories, each paired with human-annotated source and target prompts across various categories, including animals, human actions, and environmental sounds. Some example prompts in AVED-Bench, such as a dog barking, gun shooting, etc., present a unique challenge that requires precise control over audio-visual changes and synchronization.

To achieve fine-grained and synchronized editing across audio and video, we propose AVED, a zero-shot cross-modal delta denoising framework that jointly edits both modalities while maintaining temporal and structural coherence. Unlike existing methods for zero-shot video [16, 29, 32] or audio editing [47, 53], which process audio and video independently, these approaches often result in a misalignment between audio and visual content. Such limitations occur because existing methods lack a unified approach to aligning audio and video transformations holistically. To address this, built upon score distillation [18, 28, 55, 60], AVED iteratively refines the content by aligning the noise

gradients with textual prompts and patch-level audio-visual information. At each denoising step, AVED encodes audio and video into a latent space guided by textual prompts and enforces cross-modal consistency through a contrastive loss at the patch level. This ensures that edits are coherent and synchronized, maintaining both semantic meaning and temporal consistency across modalities.

We validate AVED on AVED-Bench, a benchmark dataset that we manually curated with human-annotated prompts, which consists of a wide variety of natural audio-visual events to evaluate the zero-shot audio-video editing task effectively. Additionally, we evaluate AVED on the OAVE dataset [44], which is designed for one-shot joint audio-image editing tasks. The experimental results demonstrate that our cross-modal design outperforms baselines focused solely on either video or audio editing [16, 28, 29, 32, 53].

## 2. Related Work

### 2.1. Joint Audio-visual Generation

Jointly generating audio-video content presents a challenging task compared to single-modal generation (i.e., video-to-audio [7, 51, 58, 74, 81] or audio-to-video generation [3, 30, 39, 90]). Previous methods have addressed this by introducing novel audio and video tokenizers for effective autoregressive multi-modal generation [27, 36, 66, 75, 76, 79], training diffusion models on paired audio-video data [36, 66, 79], or combining multiple single-modal diffusion models with cross-modal alignment techniques [27, 75, 76]. While these methods have shown promising results in modeling audio-video data, achieving precise control over the generated content remains challenging, particularly when editing audio-video inputs to align with specific prompts.

A recent study [44] investigates joint audio-image editing by finetuning pretrained text-to-image and text-to-audio models on a small set of text-audio-image triplets (i.e., one-shot setting) for each editing sample. Unlike audio-image editing [44] or audio-video generation [36, 66, 79], we explore a new task, zero-shot audio-video editing, which is inherently more challenging than image-only tasks due to the rich temporal information across modalities (e.g., temporal consistency and synchronization issues). Besides, the proposed AVED eliminates the need for additional finetuning of diffusion models, making it more computationally efficient than the finetuned model [44].

### 2.2. Audio/Video Editing using Diffusion Models

Recent works [1, 18, 19, 28, 32, 34, 50, 53, 55, 60, 61, 78, 80, 92] leverage pretrained text-to-image [13, 65, 91] and text-to-audio [26, 46, 48] to achieve content editing based on text prompts. Among various of these techniques [1, 18, 54], inverting source embeddings into noise vectors, such as DDIM and DDPM inversion, achieves effi-

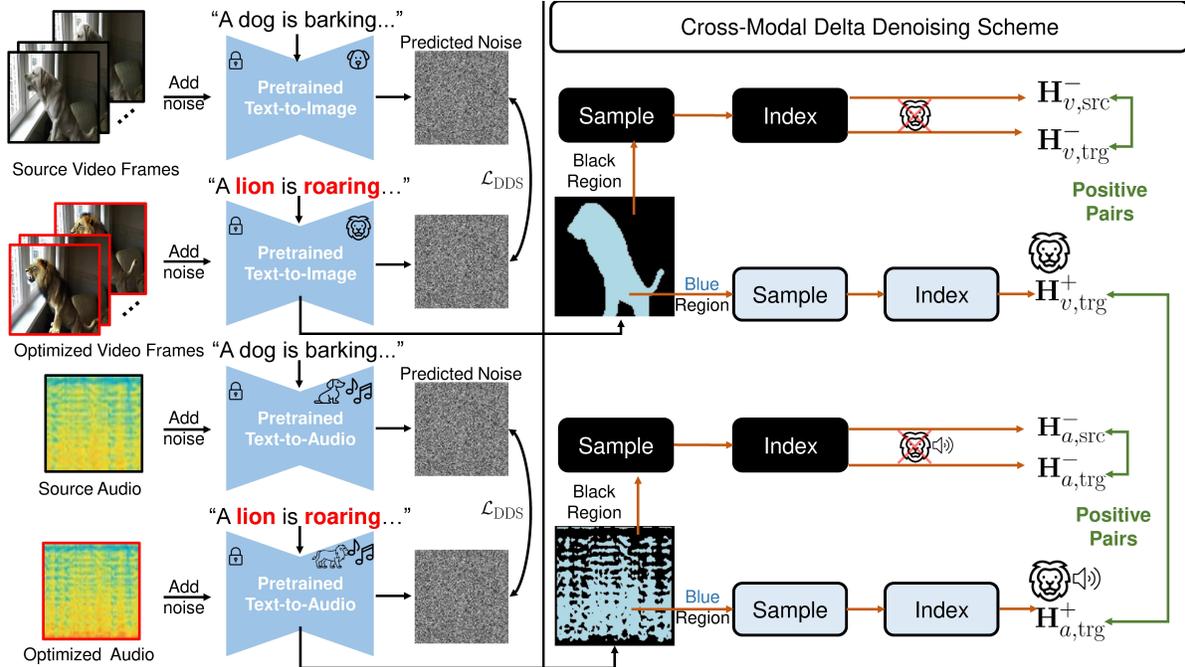


Figure 2. **Our AVED Framework.** AVED performs zero-shot audio-video editing by employing a cross-modal delta denoising score scheme to edit audio and video based on target prompts jointly. During the denoising process, relevance scores are computed between audio/image regions and target textual prompts within the cross-attention module from the diffusion model. These scores identify prompt-relevant regions (i.e., blue areas) and irrelevant patches, allowing selective editing of specific regions while preserving unaltered content. Using this region information (obtained by randomly sampling patch indices), we define positive pairs as unaltered content consistent in both the source and target branches and regions requiring edits across audio and video modalities. All other pairs are treated as negative pairs. This design enables synchronized, high-fidelity edits aligned with target prompts, maintaining coherence across audio and video.

cient zero-shot editing for visual or audio data based on text prompts. Text-based audio editing models [31, 53] leverage DDIM/DDPM inversion and pretrained diffusion models [14, 46, 48] to achieve zero-shot editing without model finetuning or additional training datasets [82]. In video editing, recent works [11, 15, 16, 28, 32, 43, 50, 61, 84, 88, 89, 92] have extended diffusion-based image editing methods [18, 19, 54, 55] to the video domain by improving the temporal consistency of edits in both zero-shot and few-shot settings. In particular, these works [11, 43, 61] inflate spatial self-attention layers to spatiotemporal layers to enhance video frame coherence by leveraging the effectiveness of inversion techniques. Another branch of work [16, 32, 86] seeks to maintain the structure of source videos by incorporating structural cues, such as optical flow [16, 86], or depth maps [29, 32].

Beyond inversion methods, Score Distillation Sampling (SDS) [60] refine images and videos by calculating gradients from pretrained diffusion models based on target prompts. While SDS enables selective editing of prompt-relevant areas, it can suffer from over-smoothing and over-saturation. To mitigate this, Delta Denoising Score (DDS) [18] introduces a reference branch with source and target prompts for SDS and demonstrates promising edit-

ing results on different tasks [28, 55]. However, these approaches [16, 18, 31, 32, 53, 55], solely editing audio and video, still suffer from the synchronization and coherence issues across two modalities. To address these limitations, we propose AVED with a cross-modal delta denoising scheme that jointly leverages audio and video information. This scheme provides cross-modal supervision to yield synchronized and high-fidelity edits in both audio and video.

### 3. Method

This section defines the notations used for zero-shot audio-video editing and provides an overview of the Delta Denoising Score (DDS) [18]. We then introduce the cross-modal denoising scheme in AVED, specifically designed to facilitate zero-shot joint audio-video editing. Our approach leverages cross-modal supervision to improve the coherence and quality of audio-visual edits.

#### 3.1. Preliminaries

**Video, Audio, and Text Inputs.** Let  $\mathbf{z}^v \in \mathbb{R}^{F \times d_v \times W_v \times H_v}$  and  $\mathbf{z}^a \in \mathbb{R}^{d_a \times W_a \times H_a}$  represent the latent features of video and audio, respectively, encoded by a pretrained VAE [38], where  $F$  denotes the number of video frames. Here,  $W_v$  and  $H_v$  represent the width and height of video frames, while  $W_a$  and  $H_a$  denote

the dimensions of the audio spectrogram. The channel dimensions for video and audio embeddings are given by  $d_v$  and  $d_a$ , respectively. Text prompts, denoted as  $y_{src}$  and  $y_{trg}$ , specify the source (optional) and target descriptions of the audio-video content. Following RAVE [32], we shuffle the temporal order of video frames and process them as a spatial grid, allowing a pretrained text-to-image diffusion model to perform video-level editing. We represent the video latent in a grid format of  $n_g \times n_g$  as  $\mathbf{z}^g \in \mathbb{R}^{M \times d_v \times W_g \times H_g}$ , where  $M = F/n_g \times n_g$ ,  $W_g = W_v/n_g$ , and  $H_g = H_v/n_g$ .

**Delta Denoising Score (DDS).** The Delta Denoising Score (DDS) [18] extends the Score Distillation Sampling (SDS) [60] by introducing a structured comparison between two distinct textual inputs and the same image inputs, referred to as the source and target branches, for precise editing. Here, the source branch represents the original, unedited content, while the target branch aligns with the desired edits specified by the target prompt. This approach enables more precise control over modifications by refining the difference between these branches. Specifically, the gradient computation in DDS starts with noise prediction in a text-conditioned diffusion model using classifier-free guidance (CFG) [22], which is formulated as follows:

$$\epsilon_\phi^\omega(\mathbf{z}_t, y, t) = (1 + \omega)\epsilon_\phi(\mathbf{z}_t, y, t) - \omega\epsilon_\phi(\mathbf{z}_t, \emptyset, t), \quad (1)$$

where  $\omega$  is the guidance parameter,  $\epsilon_\phi$  denotes the noise prediction network, and  $\emptyset$  is the null-text prompt. Here,  $y$  represents either the source or target text prompt, and  $\mathbf{z}_t$  can be an audio or image latent at a timestamp  $t \sim \mathcal{U}(0, 1)$  sampled from a uniform distribution. DDS computes gradients based on both target and reference inputs, defined by the following objective:

$$\mathcal{L}_{\text{DDS}}(\theta; y_{trg}) = \|\epsilon_\phi^\omega(\mathbf{z}_t(\theta), y_{trg}, t) - \epsilon_\phi^\omega(\mathbf{z}_t, y_{src}, t)\|^2, \quad (2)$$

where  $\mathbf{z}_t(\theta)$  represents the target latent, parameterized by  $\theta$ .  $\mathbf{z}_t(\theta)$  is iteratively updated in the direction of the gradient  $\nabla_\theta \mathcal{L}_{\text{DDS}}$  to adjust toward the target prompt precisely.

### 3.2. Cross-Modal Delta Denoising Scheme

We introduce a cross-modal delta denoising scheme to improve the quality and coherence of audio-video editing by leveraging interactions between modalities during each DDS process. Unlike single-modality approaches, our method incorporates complementary information from both audio and video, ensuring synchronized and contextually consistent edits. We first identify **prompt-relevant regions** in audio and video to achieve this. This is done by computing similarity scores within the cross-attention modules of the diffusion layers. These scores indicate how well different patches in the audio and video streams align with the prompts. By applying thresholding, we classify relevant

patches (aligned with the prompt) and irrelevant patches (background or unchanged content). Once identified, we sample pairs from the source (i.e.,  $\epsilon_\phi^\omega(\mathbf{z}_t, y_{src}, t)$ ) and target (i.e.,  $\epsilon_\phi^\omega(\mathbf{z}_t(\theta), y_{trg}, t)$ ) prompts. These sampled pairs serve as the foundation for a contrastive loss, which enforces alignment between corresponding regions across modalities, improving synchronization and overall coherence.

**Prompt-Relevant Patches.** To identify regions in the audio and video relevant to prompts, we leverage intermediate representations in the cross-modal attention layers of pretrained diffusion models. By computing the similarity between audio/video features (queries) and textual prompt features (keys), we highlight areas aligned with prompts.

Let  $\mathbf{Q}_a \in \mathbb{R}^{n_q \times d}$  and  $\mathbf{Q}_v \in \mathbb{R}^{M \times n_q \times d}$  be the audio and video query features, while  $\mathbf{K}_a \in \mathbb{R}^{n_k \times d}$  and  $\mathbf{K}_v \in \mathbb{R}^{M \times n_k \times d}$  represent their corresponding key features derived from the target prompt. The similarity scores  $\mathbf{S}$  are computed as:

$$\mathbf{S}_i^a = \max_j (\mathbf{Q}_a \mathbf{K}_a^\top)_{i,j}, \quad \mathbf{S}_i^v = \max_j (\mathbf{Q}_v \mathbf{K}_v^\top)_{i,j}. \quad (3)$$

Here,  $\mathbf{S}^a \in \mathbb{R}^{n_q}$  and  $\mathbf{S}^v \in \mathbb{R}^{M \times n_q}$  represent the prompt relevance scores for audio and video patches. To focus on the most relevant regions, we apply max pooling across the text dimension  $j$ . We then normalize the scores using min-max normalization to ensure values range between 0 and 1 for consistency. The final normalized scores,  $\hat{\mathbf{S}}_{\text{trg}}^a$  and  $\hat{\mathbf{S}}_{\text{trg}}^v$ , will be thresholded to distinguish relevant and irrelevant patches for further processing.

**Contrastive Loss for Denoising.** Unlike typical contrastive loss, which typically maximizes the similarity between features in a batch (e.g., different audio-video samples), we leverage contrastive loss on relevant and irrelevant patches within the same audio/video instance, which gradually transforms over timesteps by the DDS process, to improve coherence in audio-video editing. For example, when editing a video where a dog transforms into a lion, we aim to align audio and video regions such as the lion’s roar and fur as positive pairs, while ensuring that the background or unrelated objects remain unchanged. To preserve context, irrelevant patches with the same spatial location in both the source and target branches of each modality (e.g., background) are also treated as positive pairs. All other combinations are considered negative pairs to achieve precise and contextually consistent edits.

To achieve this, we define  $\mathcal{I}_a^+$  and  $\mathcal{I}_v^+$  as indices for patches in the target audio and video branches aligned with a textual prompt, where  $\hat{\mathbf{S}}_{\text{trg}}^a > \tau_a$  or  $\hat{\mathbf{S}}_{\text{trg}}^v > \tau_v$  with given threshold  $\tau_a$  and  $\tau_v$ . These relevant patches correspond to areas that need editing, such as the sound and appearance of the lion when it transforms from a dog. Similarly,  $\mathcal{I}_a^-$  and  $\mathcal{I}_v^-$  represent indices for patches in the target branch that are not related to the prompt and do not require mod-

ification, such as background sounds or static visual elements where  $\hat{\mathbf{S}}_{\text{trg}}^a < \tau_a$  or  $\hat{\mathbf{S}}_{\text{trg}}^v < \tau_v$ . We then extract audio embeddings  $\mathbf{h}_a \in \mathbb{R}^{n_a^a \times d_a}$  and video embeddings  $\mathbf{h}_v \in \mathbb{R}^{M \times n_v^v \times d_v}$  from the hidden states of cross-modal attention layers to capture rich information for precise editing. To achieve more diverse views and robust editing, we randomly sample embeddings for relevant regions from the target branch:

$$\mathbf{H}_{a,\text{trg}}^+ = \{\mathbf{h}_{a,\text{trg},i} \mid i \in \mathcal{I}_a^+\}, \quad \mathbf{H}_{v,\text{trg}}^+ = \{\mathbf{h}_{v,\text{trg},i} \mid i \in \mathcal{I}_v^+\}. \quad (4)$$

Here,  $\mathbf{H}_a^+$  and  $\mathbf{H}_v^+$  represent sets of relevant audio and video embeddings sampled from the target branch. Similarly, for irrelevant regions, we sample embeddings directly from both the source and target branches:

$$\begin{aligned} \mathbf{H}_{a,\text{src}}^- &= \{\mathbf{h}_{a,\text{src},i} \mid i \in \mathcal{I}_a^-\}, & \mathbf{H}_{a,\text{trg}}^- &= \{\mathbf{h}_{a,\text{trg},i} \mid i \in \mathcal{I}_a^-\}, \\ \mathbf{H}_{v,\text{src}}^- &= \{\mathbf{h}_{v,\text{src},i} \mid i \in \mathcal{I}_v^-\}, & \mathbf{H}_{v,\text{trg}}^- &= \{\mathbf{h}_{v,\text{trg},i} \mid i \in \mathcal{I}_v^-\}. \end{aligned} \quad (5)$$

Here,  $\mathbf{H}_a^-$  and  $\mathbf{H}_v^-$  represent sets of irrelevant audio and video embeddings sampled from both branches. We then apply a contrastive loss that encourages high similarity between positive pairs, including relevant patches and irrelevant patches matched by the same index across source and target branches within the same modality while discouraging similarity with unrelated negative pairs. The standard contrastive loss is described as follows:

$$\mathcal{L}_c(\mathbf{F}_x, \mathbf{F}_y) = -\frac{1}{N} \sum_{i=1}^N \log \frac{\exp(\text{sim}(\mathbf{F}_x^i, \mathbf{F}_y^i)/\alpha)}{\sum_{j=1}^N \exp(\text{sim}(\mathbf{F}_x^i, \mathbf{F}_y^j)/\alpha)}, \quad (6)$$

where  $\text{sim}(\cdot, \cdot)$  denotes cosine similarity,  $N$  is the mini-batch size, and  $\alpha$  is a temperature parameter. Our final cross-modal contrastive loss combines alignments as follows:

$$\begin{aligned} \mathcal{L}_{\text{cmds}} &= \frac{1}{M} \sum_{i=1}^M \left( \mathcal{L}_c([\mathbf{H}_{a,\text{trg}}^+, \mathbf{H}_{v,\text{src}}^-[i]], [\mathbf{H}_{v,\text{trg}}^+[i], \mathbf{H}_{a,\text{src}}^-]) \right. \\ &\quad \left. + \mathcal{L}_c([\mathbf{H}_{v,\text{trg}}^+[i], \mathbf{H}_{a,\text{src}}^-], [\mathbf{H}_{a,\text{trg}}^+, \mathbf{H}_{v,\text{src}}^-[i]]) \right). \end{aligned} \quad (7)$$

The proposed  $\mathcal{L}_{\text{cmds}}$  not only leverages cross-modal information to improve editing quality, but also preserves unedited content by considering unrelated patches across branches in the same modality. For instance, when converting a barking dog into a roaring lion, the generated lion’s roar ( $\mathbf{H}_{a,\text{trg}}^+$ ) should be distinct from non-dog sounds in the source audio ( $\mathbf{H}_{a,\text{src}}^-$ ), and the lion’s appearance ( $\mathbf{H}_{v,\text{trg}}^+$ ) should differ from regions in the original video that lack lion-related features ( $\mathbf{H}_{v,\text{src}}^-$ ). If  $d_v$  and  $d_a$  differ in some layers, we simply interpolate to match the dimensions. The final objective integrates this contrastive loss with DDS, formulated as  $\mathcal{L}_{\text{cmds}} + \mathcal{L}_{\text{DDS}}(\theta; y_{\text{trg}})$ .<sup>1</sup>

<sup>1</sup>Implementation details are provided in the supplementary material.

## 4. Experimental Setup

### 4.1. Downstream Datasets

- **AvED-Bench**<sup>2</sup>, our newly curated dataset, consists of 110 10-second videos in 11 distinct categories from VGGSound [8], covering diverse scenes such as animals, objects, and environmental sounds. Each video is paired with annotated source and target prompts that specify audio-visual events and object categories. This dataset provides a comprehensive benchmark for evaluating zero-shot audio-video editing capabilities.
- **OAVE** [44] contains 44 categories, each with 10 images from a clip and separate audio and visual annotations. It includes 25 prompt templates to modify either the sounding object or the environmental context for evaluating editing performance.

### 4.2. Evaluation Metrics

Following previous work [16, 18, 32, 53–55], we evaluate our generated audio-video samples using several metrics: **CLIP-F** evaluates the consistency of edited frames by calculating the similarity of frame-based CLIP embeddings [62]. **CLIP-T** evaluates the alignment between the target prompt and the edited video by computing the CLIP similarity [62] between each video frame and the prompt, then averaging these similarity scores across all frames. **DINO** emphasizes the preservation of the overall structure between the source and target frames by computing cosine similarity with self-supervised DINO-ViT embeddings [6]. **Obj** utilizes Grounding-DINO [49] to detect and assess the presence and likelihood of target objects specified in the prompt, evaluating how accurately these objects are generated. **CLAP** measures the cosine similarity between audio and target prompt using the CLAP model [85], indicating the fidelity of the edited sound. **LPAPS** evaluates perceptual distance in source and target audio in CLAP feature space [85] to provide an assessment of faithfulness/consistency between edited and source sound. **AV-Align** [87] metric examines alignment between audio cues and visual changes, low-level coherence in audio-visual transitions. **IB** leverages ImageBind [17] embeddings to assess audio-visual similarity to evaluate high-level audio-visual coherence.

### 4.3. Human Evaluation

We conduct a human evaluation by asking subjects to select their preferred edited audio-video samples according to their alignment with the target prompt. Specifically, given a source (unedited) video and a pair of edited audio-video samples, human raters are asked to select their preferred sample based on the following question: *Which video do you think has the better editing quality overall?* For each question, subjects can choose one of the two methods or a

<sup>2</sup>Full annotations are included in the supplementary material.

Video Model	Audio Model	Video-Only				Audio-Only		Joint AV	
		CLIP-F $\uparrow$	CLIP-T $\uparrow$	Obj. $\uparrow$	DINO $\uparrow$	CLAP $\uparrow$	LPAPS $\downarrow$	IB. $\uparrow$	AV-Align $\uparrow$
ControlVideo [92]	SDEdit [54]	0.883	0.255	0.176	0.892	0.190	6.93	0.21	0.29
TokenFlow [16]	SDEdit [54]	0.876	0.252	0.173	0.924	0.190	6.93	0.19	0.26
RAVE [32]	SDEdit [54]	0.885	0.251	0.170	0.881	0.190	6.93	0.18	0.29
ControlVideo [92]	ZEUS [53]	0.883	0.255	0.176	0.892	0.211	6.41	0.21	0.30
TokenFlow [16]	ZEUS [53]	0.876	0.252	0.173	0.924	0.211	6.41	0.20	0.27
RAVE [32]	ZEUS [53]	0.885	0.251	0.170	0.881	0.211	6.41	0.18	0.31
RAVE [32] $\rightarrow$ Diff-Foley [51]		0.885	0.251	0.170	0.881	0.191	7.33	0.16	0.35
Delta Denoising Score (DDS) [18]		0.890	0.250	0.175	0.921	0.210	5.93	0.20	0.33
AVED (Ours)		<b>0.903</b>	<b>0.260</b>	<b>0.180</b>	<b>0.956</b>	<b>0.226</b>	<b>5.55</b>	<b>0.23</b>	<b>0.42</b>

Table 2. **Comparison to the State-of-the-Art Zero-Shot Video and Audio Editing Models.** We compare our AVED with baselines for zero-shot video [16, 18, 32, 92] or audio [18, 53, 54] editing on AVED-Bench. Our evaluation metrics evaluate diverse aspects, including video-only, audio-only, and joint audio-video editing quality.

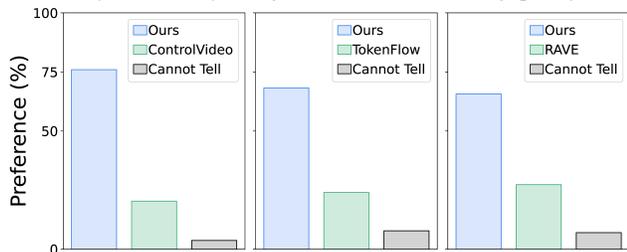


Figure 3. **Human Evaluation.** Human raters evaluate edited audio and video quality based on alignment with the target text prompt. We report the average human preference rate for each method. All samples are presented in a random order to ensure unbiased assessment.

third option, “Cannot tell.” Each subject evaluates five randomly selected video pairs, with one sample always from AVED. To prevent bias, the methods remain unknown to the raters. We compare our AVED method against competing approaches [16, 32, 92]. Results are reported as the average human preference rate for each method. Our human study is conducted with 300 subjects on Amazon Mechanical Turk.

#### 4.4. Baselines

We compare our model to recent baselines in zero-shot audio or video editing, evaluating each modality independently due to the lack of joint audio-visual approaches. For video editing, we assess: (i) RAVE [32], which uses pre-trained text-to-image diffusion models and noise-shuffling for temporally consistent edits, (ii) TokenFlow [16], which maintains feature consistency through inter-frame correspondences for high-quality edits without extra training, and (iii) ControlVideo [92], which adapts ControlNet [91] for training-free text-to-video generation using depth maps and human poses. For audio editing, we evaluate: (i) ZEUS [53], which edits via DDPM inversion and reversion to edit sound, and (ii) SDEdit [54], which blends denoising processes with initial noise to edit audio outputs. We implement a sequential baseline, RAVE [32]  $\rightarrow$  Diff-Foley [51],

Method	Video-Only		Audio-Only		Joint AV
	CLIP-T $\uparrow$	DINO $\uparrow$	CLAP $\uparrow$	LPAPS $\downarrow$	IB. $\uparrow$
ControlVideo [92]	0.261	0.892	0.245	5.91	0.28
TokenFlow [16]	0.263	0.901	0.245	5.91	0.28
RAVE [32]	<b>0.268</b>	0.891	0.245	5.91	0.29
DDS	0.260	0.930	0.240	5.58	0.28
AVED	0.265	<b>0.959</b>	<b>0.250</b>	<b>5.15</b>	<b>0.32</b>

Table 3. **Comparison on the OAVE Dataset.** We evaluate AVED on the OAVE dataset alongside leading zero-shot audio and video editing models [16, 18, 32, 92]. The comparison includes video-only, audio-only, and joint audio-video editing metrics. Since OAVE emphasizes audio-image editing, metrics like CLIP-F and AV-Align may be less applicable here.

where RAVE edits video first, followed by Diff-Foley for audio generation. We also implement DDS [18], which applies the same methodology separately to audio and video using image grids and shuffling. This evaluation allows us to compare the performance of AVED with leading methods in both domains.

## 5. Results and Analysis

### 5.1. Comparison with the State-of-the-Art

In Table 2, we present a detailed comparison of AVED with the leading zero-shot video and audio editing models [16, 18, 32, 53, 54, 92] and the sequential baseline. We evaluate performance on video-only, audio-only, and joint audio-video metrics. For **video-only metrics**, AVED consistently outperforms baseline models, achieving a notable increase in DINO scores (e.g., **0.956** vs. **0.921** for DDS [18]), which highlights AVED’s strength in preserving visual coherence and structural fidelity between source and edited video frames. Furthermore, AVED achieves the highest CLIP-F, CLIP-T, and Obj. scores (**0.903**, **0.260**, and **0.180**, respectively), demonstrating its ability to maintain frame consistency and align edits precisely with the target prompt.

Configuration	DINO $\uparrow$	LPAPS $\downarrow$	AV-Align $\uparrow$
Baseline	0.921	5.93	0.33
+ Audio-Only	0.921	5.60	0.36
+ Video-Only	0.937	5.60	0.38
AVED (+AV)	<b>0.956</b>	<b>5.55</b>	<b>0.42</b>

Table 4. **Cross-Modal vs. Single-Modal Delta Denoising Schemes.** We study AVED with single-modal delta denoising schemes (“Audio-Only” and “Video-Only”) and the baseline (i.e., AVED without Eq. 7) across key metrics about faithfulness, consistency, and audio-visual alignment on AVED-Bench.

For **audio-only metrics**, AVED achieves the highest CLAP score (**0.226**) and the lowest LPAPS score (**5.55**). Compared to DDS [18], ZEUS [53], and SDEdit [54], AVED shows a significant improvement in LPAPS scores (i.e., **5.55** vs. **5.93**, **6.41**, and **6.93**) indicating superior perceptual consistency between the source and edited audio. We note that AVED also outperforms the sequential baseline, achieving higher CLAP (**0.226** vs. **0.191**) and lower LPAPS (**5.55** vs. **7.33**), demonstrating that the sequential approach leads to degraded audio quality due to potential imperfect results from edited videos.

For **joint audio-video metrics**, AVED demonstrates a remarkable improvement, achieving over a **20%** relative improvement in AV-Align scores compared to baselines. This result validates AVED’s ability to synchronize visual and audio edits accurately, aligning actions and sounds seamlessly. Furthermore, in the ImageBind (IB) score, AVED also presents the highest score, suggesting a better high-level audio-video alignment. Compared to the sequential baseline, AVED also achieves a higher IB (**0.23** vs. **0.16**) and AV-Align (**0.42** vs. **0.35**), indicating that the sequential approach accumulates misalignment errors from imperfect video edits and degraded audio quality.

In Table 3, we evaluate AVED on the OAVE dataset, a recent benchmark for one-shot audio-image editing. Since this task primarily focuses on audio-image editing, metrics such as CLIP-F and AV-Align are less applicable. The results show that AVED outperforms baselines across key metrics, achieving the highest DINO score (**0.959** vs. **0.930**) and the lowest LPAPS score (**5.15** vs. **5.58**), demonstrating its ability to maintain coherence and structural fidelity within audio-visual content. Furthermore, AVED achieves a notable increase in the IB score (**0.32** vs. **0.29**), indicating a stronger alignment of audio-visual semantics.

The results in both AVED-Bench and OAVE [44] validate the effectiveness of our cross-modal delta denoising scheme in preserving structural consistency (DINO, LPAPS), maintaining perceptual fidelity (CLIP-T, CLAP), and ensuring aligned audio-visual content (AV-Align, IB). These findings underline the necessity of joint audio-video

Configuration	DINO $\uparrow$	LPAPS $\downarrow$	AV-Align $\uparrow$
Baseline	0.921	5.93	0.33
Random A/V	0.930	5.83	0.35
Random A+V	0.902	6.32	0.28
AVED	<b>0.956</b>	<b>5.55</b>	<b>0.42</b>

Table 5. **Selecting Positive and Negative Pairs.** We investigate how AVED determines positive and negative pairs without prior knowledge of prompt-relevant patches. Similar to the setup in AVED, we randomly sample indices without identifying prompt-relevant patches. In the same modality, patches with the same index across different branches are treated as positive pairs (“Random A/V”). For cross-modal settings (Random A+V), we assume that corresponding audio and video patches are positive pairs.

editing, as sequential and single-modality approaches introduce misalignment and degrade perceptual quality.

**Human Evaluation.** Figure 3 presents the results of our human study to assess the overall quality of the edited audio video based on alignment with the target prompt. We report human preference rates, which indicate the percentage of raters who preferred AVED over each baseline method. Each comparison is conducted between AVED and the baseline methods, including ControlVideo [92], TokenFlow [16], and RAVE [32]. As shown in Figure 3, approximately **75%** of participants preferred AVED over ControlVideo, and over **60%** favored AVED compared to TokenFlow and RAVE, highlighting AVED’s superior editing quality.

## 5.2. Ablation Studies

**Impact of Cross-Modal vs. Single-Modal Delta Denoising Scheme.** In Table 4, we analyze the impact of our proposed cross-modal delta denoising scheme compared to a single-modal denoising approach on AVED-Bench. The single-modal denoising approaches (i.e., “Audio-Only” and “Video-Only”) indicate that only one modality’s features (either audio or video) are used in Eq. 7. Starting with the baseline (i.e., DDS [18] baseline), we see improvements using single-modal schemes.

In the audio-only denoising scheme, we observe an LPAPS reduction from **5.93** (Baseline) to **5.60**, which demonstrates greater consistency between edited and source sounds, resulting in fewer artifacts and a notable improvement in audio-visual alignment as shown by the AV-Align metric increase from **0.33** to **0.36**. A similar trend is observed in the video-only denoising scheme, where the DINO score increases from **0.921** to **0.937**, indicating enhanced visual structure preservation of the source content. Significant improvements are observed when both audio and visual information are integrated within the cross-modal denoising scheme. Compared to the baseline, the cross-modal approach increases the DINO score from **0.921**

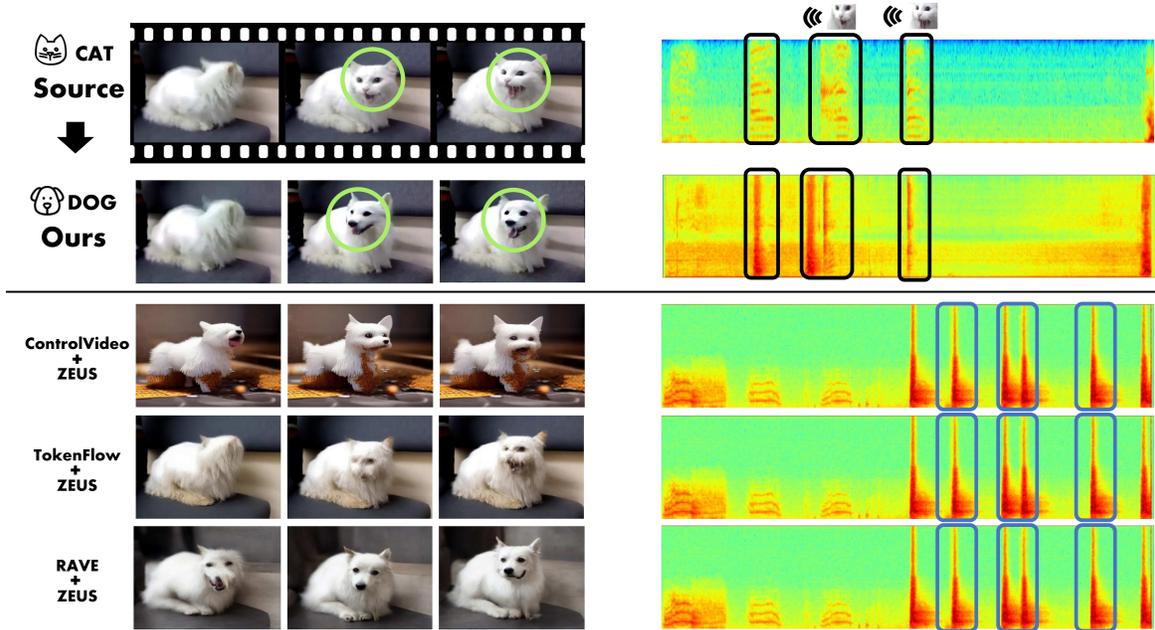


Figure 4. **Qualitative Zero-Shot Audio-Video Editing Results.** We present qualitative results of audio-video editing for a video depicting a transition from "Cat" to "Dog." AVED is compared with video models, including ControlVideo [92], TokenFlow [16], and RAVE [32], along with the audio model ZEUS [53]. The green circles highlight well-aligned motion matches in the video frames, while the black rectangles emphasize precise audio matching. The blue rectangles indicate audio artifacts in the competing models, leading to the misalignment between video actions and audio output.

to **0.956** and reduces LPAPS from **5.93** to **5.55**, thus significantly boosting the AV-Align metric from **0.33** to **0.42**. These results show the effectiveness of cross-modal design in leading synchronized and coherent audio-video edits.

**Impact of Selecting Positive and Negative Pairs.** Next, we study patch selection strategies in AVED without prompt-relevant information in Table 5. The "Random A/V" configuration, which samples positive pairs within the same modality across source and target branches, slightly improves the baseline (**0.930** DINO, **5.83** LPAPS, **0.35** AV-Align). Since most patches are unrelated to the editing patches, even random selection provides some advantages. Instead, the "Random A+V" configuration, assuming all intra-branch audio-video pairs are positive, degrades performance (**0.902** DINO, **6.32** LPAPS, **0.28** AV-Align) since the objects or sounds we aim to edit are usually limited to specific regions. Our prompt-relevant selection yields the best results (**0.956** DINO, **5.55** LPAPS, **0.42** AV-Align), confirming its importance for audio-video editing.

### 5.3. Qualitative Results

Figure 4 presents qualitative results of zero-shot audio-video editing for a "Cat" to "Dog" transition. AVED is compared with video editing methods, including ControlVideo [92], TokenFlow [16], RAVE [32], and the audio model ZEUS [53]. The green circles highlight well-aligned motion in the video frames to demonstrate that AVED accurately transforms the visual appearance and motion, poten-

tially producing sound. Besides, black rectangles emphasize precise audio matching, demonstrating temporal consistency with the source sound. In contrast, competing models exhibit misalignment, as indicated by blue rectangles showing random audio artifacts that disrupt synchronization. Overall, this qualitative comparison highlights the effectiveness of AVED's cross-modal delta denoising scheme in achieving synchronized audio-video edits.

## 6. Conclusions

In this paper, we introduce AVED, a zero-shot audio-video editing framework developed to address the novel task of synchronized audio-visual editing. AVED proposes a cross-modal delta denoising scheme that enables synchronized and coherent edits by integrating interactions between audio and video modalities. To support this task, we curate a benchmark dataset, AVED-Bench, which features diverse and challenging audio-visual editing scenarios paired with human-annotated prompts. We evaluate AVED on AVED-Bench and the OAVE dataset, demonstrating its superior performance compared to existing single-modality and joint editing baselines. AVED consistently achieves strong coherence between original and edited content in both modalities, as well as low-level coherence in audio-visual transitions, validating the effectiveness of our cross-modal approach for producing synchronized, high-quality audio-video editing.

## Acknowledgments

This work was supported by Laboratory for Analytic Sciences via NC State University, ONR Award N00014-23-1-2356, NIH Award R01HD11107402, and Sony Focused Research award.

## References

- [1] Omri Avrahami, Dani Lischinski, and Ohad Fried. Blended diffusion for text-driven editing of natural images. In *CVPR*, 2022. 1, 2
- [2] Omer Bar-Tal, Dolev Ofri-Amar, Rafail Fridman, Yoni Kasten, and Tali Dekel. Text2live: Text-driven layered image and video editing. In *ECCV*, 2022. 2
- [3] Burak Can Biner, Farrin Marouf Sofian, Umur Berkay Karakaş, Duygu Ceylan, Erkut Erdem, and Aykut Erdem. Sonicdiffusion: Audio-driven image generation and editing with pretrained diffusion models. *arXiv Preprint*, 2024. 2
- [4] Andreas Blattmann, Robin Rombach, Patrick Esser, and Björn Ommer. Videoldm: Text-to-video generation with latent diffusion models. In *ICLR*, 2023. 1
- [5] Tim Brooks, Aleksander Holynski, and Alexei A Efros. Instructpix2pix: Learning to follow image editing instructions. In *CVPR*, 2023. 1
- [6] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *ICCV*, 2021. 5
- [7] Gehui Chen, Guan'an Wang, Xiaowen Huang, and Jitao Sang. Semantically consistent video-to-audio generation using multimodal language large model. In *arXiv Preprint*, 2024. 2
- [8] Honglie Chen, Weidi Xie, Andrea Vedaldi, and Andrew Zisserman. Vggsound: A large-scale audio-visual dataset. In *ICASSP*, 2020. 2, 5
- [9] Ziyang Chen, Daniel Geng, and Andrew Owens. Images that sound: Composing images and sounds on a single canvas. In *NeurIPS*, 2024. 1
- [10] Ziyang Chen, Prem Seetharaman, Bryan Russell, Oriol Nieto, David Bourgin, Andrew Owens, and Justin Salamon. Video-guided foley sound generation with multimodal controls. In *CVPR*, 2025. 1
- [11] Nathaniel Cohen, Vladimir Kulikov, Matan Kleiner, Inbar Huberman-Spiegelglas, and Tomer Michaeli. Slicedit: Zero-shot video editing with text-to-image diffusion models using spatio-temporal slices. In *ICML*, 2024. 1, 3
- [12] Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. In *NeurIPS*, 2021. 1
- [13] Patrick Esser, Sumith Kulal, Andreas Blattmann, Rahim Entezari, Jonas Müller, Harry Saini, Yam Levi, Dominik Lorenz, Axel Sauer, Frederic Boesel, et al. Scaling rectified flow transformers for high-resolution image synthesis. In *ICML*, 2024. 2
- [14] Zach Evans, CJ Carr, Josiah Taylor, Scott H Hawley, and Jordi Pons. Fast timing-conditioned latent audio diffusion. In *ICML*, 2024. 2, 3
- [15] Xiang Fan, Anand Bhattad, and Ranjay Krishna. Videoshop: Localized semantic video editing with noise-extrapolated diffusion inversion. In *ECCV*, 2024. 3
- [16] Michal Geyer, Omer Bar-Tal, Shai Bagon, and Tali Dekel. Tokenflow: Consistent diffusion features for consistent video editing. In *ICLR*, 2023. 2, 3, 5, 6, 7, 8
- [17] 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 *CVPR*, 2023. 5
- [18] Amir Hertz, Kfir Aberman, and Daniel Cohen-Or. Delta denoising score. In *ICCV*, 2023. 2, 3, 4, 5, 6, 7
- [19] Amir Hertz, Ron Mokady, Jay Tenenbaum, Kfir Aberman, and Dani Lischinski. Prompt-to-prompt image editing with cross attention control. In *ICLR*, 2023. 2, 3
- [20] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *NeurIPS*, 2020. 1
- [21] Jonathan Ho, Chitwan Saharia, William Chan, Tim Salimans, David J. Fleet, and Mohammad Norouzi. Imagen video: High definition video generation with diffusion models. *arXiv Preprint*, 2022. 1, 2
- [22] Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. *arXiv Preprint*, 2022. 4
- [23] Jonathan Ho, Tim Salimans, Alexey Gritsenko, William Chan, Mohammad Norouzi, and David J. Fleet. Video diffusion models. In *NeurIPS*, 2022. 2
- [24] Wenyi Hong, Ming Ding, Wendi Chen, Woonhyuk Baek Zhang, Zhuoyi Yang, Xiaojie Xu, Junyuan Wang, Chang Zhou, Hongxia Yang, and Jie Tang. Cogvideo: Large-scale pretraining for text-to-video generation via transformers. In *ICLR*, 2023. 1
- [25] Qingqing Huang, Daniel S. Park, Tao Wang, Timo I. Denk, Andy Ly, Nanxin Chen, Zhengdong Zhang, Zhishuai Zhang, Jiahui Yu, Christian Frank, et al. Noise2music: Text-conditioned music generation with diffusion models. *arXiv Preprint*, 2023. 1
- [26] Rongjie Huang, Jiawei Huang, Dongchao Yang, Yi Ren, Luping Liu, Mingze Li, Zhenhui Ye, Jinglin Liu, Xiang Yin, and Zhou Zhao. Make-an-audio: Text-to-audio generation with prompt-enhanced diffusion models. In *ICML*, 2023. 1, 2
- [27] Masato Ishii, Akio Hayakawa, Takashi Shibuya, and Yuki Mitsufuji. A simple but strong baseline for sounding video generation: Effective adaptation of audio and video diffusion models for joint generation. *arXiv Preprint*, 2024. 2
- [28] Hyeonho Jeong, Jinho Chang, Geon Yeong Park, and Jong Chul Ye. Dreammotion: Space-time self-similar score distillation for zero-shot video editing. In *ECCV*, 2024. 2, 3
- [29] Hyeonho Jeong and Jong Chul Ye. Ground-a-video: Zero-shot grounded video editing using text-to-image diffusion models. In *ICLR*, 2024. 2, 3
- [30] Yujin Jeong, Wonjeong Ryoo, Seunghyun Lee, Dabin Seo, Wonmin Byeon, Sangpil Kim, and Jinkyu Kim. The power of sound (tpos): Audio reactive video generation with stable diffusion. In *NeurIPS*, 2023. 2
- [31] Yuhang Jia, Yang Chen, Jinghua Zhao, Shiwan Zhao, Wenjia Zeng, Yong Chen, and Yong Qin. Audioeditor: A training-free diffusion-based audio editing framework. *arXiv Preprint*, 2024. 1, 2, 3
- [32] Ozgur Kara, Bariscan Kurtkaya, Hidir Yesiltepe, James M Rehg, and Pinar Yanardag. Rave: Randomized noise shuffling for fast and consistent video editing with diffusion mod-

- els. In *CVPR*, 2024. 2, 3, 4, 5, 6, 7, 8
- [33] Yoni Kasten, Dolev Ofri, Oliver Wang, and Tali Dekel. Layered neural atlases for consistent video editing. *ACM TOG*, 2021. 2
- [34] Levon Khachatryan, Andranik Movsisyan, Vahram Tadevosyan, Roberto Henschel, Zhangyang Wang, Shant Navasardyan, and Humphrey Shi. Text2video-zero: Text-to-image diffusion models are zero-shot video generators. In *ICCV*, 2023. 2
- [35] Gwanghyun Kim, Taesung Kwon, and Jong Chul Ye. Diffusionclip: Text-guided diffusion models for robust image manipulation. In *CVPR*, 2022. 1
- [36] Gwanghyun Kim, Alonso Martinez, Yu-Chuan Su, Brendan Jou, José Lezama, Agrim Gupta, Lijun Yu, Lu Jiang, Aren Jansen, Jacob Walker, et al. A versatile diffusion transformer with mixture of noise levels for audiovisual generation. In *NeurIPS*, 2024. 2
- [37] Jang-Hyun Kim and Jong Chul Ye. Ilvr: Conditioning method for denoising diffusion probabilistic models. In *ICCV*, 2021. 1
- [38] Diederik P Kingma. Auto-encoding variational bayes. *arXiv Preprint*, 2013. 3
- [39] Seung Hyun Lee, Gyeongrok Oh, Wonmin Byeon, Chanyoung Kim, Won Jeong Ryoo, Sang Ho Yoon, Hyunjun Cho, Jihyun Bae, Jinkyu Kim, and Sangpil Kim. Sound-guided semantic video generation. In *ECCV*, 2022. 2
- [40] Yao-Chih Lee, Ji-Ze Genevieve Jang, Yi-Ting Chen, Elizabeth Qiu, and Jia-Bin Huang. Shape-aware text-driven layered video editing. In *CVPR*, 2023. 2
- [41] Peike Li, Boyu Chen, Yao Yao, Yikai Wang, Allen Wang, and Alex Wang. Jen-1: Text-guided universal music generation with omnidirectional diffusion models. In *CAI*, 2024. 1
- [42] Ruiqi Li, Siqi Zheng, Xize Cheng, Ziang Zhang, Shengpeng Ji, and Zhou Zhao. Muvi: Video-to-music generation with semantic alignment and rhythmic synchronization. *arXiv Preprint*, 2024. 1
- [43] Xirui Li, Chao Ma, Xiaokang Yang, and Ming-Hsuan Yang. Vidtome: Video token merging for zero-shot video editing. In *CVPR*, 2024. 3
- [44] Susan Liang, Chao Huang, Yapeng Tian, Anurag Kumar, and Chenliang Xu. Language-guided joint audio-visual editing via one-shot adaptation. In *ACCV*, 2024. 2, 5, 7
- [45] Yan-Bo Lin, Yu Tian, Linjie Yang, Gedas Bertasius, and Heng Wang. Vmas: Video-to-music generation via semantic alignment in web music videos. In *WACV*, 2025. 1
- [46] 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. In *ICML*, 2023. 1, 2, 3
- [47] Huadai Liu, Jialei Wang, Rongjie Huang, Yang Liu, Jiayang Xu, and Zhou Zhao. Medic: Zero-shot music editing with disentangled inversion control. *arXiv Preprint*, 2024. 2
- [48] Haohe Liu, Yi Yuan, Xubo Liu, Xinhao Mei, Qiuqiang Kong, Qiao Tian, Yuping Wang, Wenwu Wang, Yuxuan Wang, and Mark D Plumbley. Audioldm 2: Learning holistic audio generation with self-supervised pretraining. *TASLP*, 2024. 1, 2, 3
- [49] Shilong Liu, Zhaoyang Zeng, Tianhe Ren, Feng Li, Hao Zhang, Jie Yang, Qing Jiang, Chunyuan Li, Jianwei Yang, Hang Su, et al. Grounding dino: Marrying dino with grounded pre-training for open-set object detection. In *ECCV*, 2024. 5
- [50] Shaoteng Liu, Yuechen Zhang, Wenbo Li, Zhe Lin, and Jiaya Jia. Video-p2p: Video editing with cross-attention control. In *CVPR*, 2024. 1, 2, 3
- [51] Simian Luo, Chuanhao Yan, Chenxu Hu, and Hang Zhao. Diff-foley: Synchronized video-to-audio synthesis with latent diffusion models. In *NeurIPS*, 2023. 2, 6
- [52] 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. In *ACM MM*, 2024. 1
- [53] Hila Manor and Tomer Michaeli. Zero-shot unsupervised and text-based audio editing using ddpm inversion. In *ICML*, 2024. 1, 2, 3, 5, 6, 7, 8
- [54] 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. In *ICLR*, 2022. 1, 2, 3, 6, 7
- [55] Hyelin Nam, Gihyun Kwon, Geon Yeong Park, and Jong Chul Ye. Contrastive denoising score for text-guided latent diffusion image editing. In *CVPR*, 2024. 2, 3, 5
- [56] Alex Nichol and Prafulla Dhariwal. Improved denoising diffusion probabilistic models. In *ICML*, 2021. 1
- [57] 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. *arXiv Preprint*, 2021. 1
- [58] Santiago Pascual, Chunghsin Yeh, Ioannis Tsiamas, and Joan Serra. Masked generative video-to-audio transformers with enhanced synchronicity. In *ECCV*, 2024. 2
- [59] 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*, 2024. 1
- [60] Ben Poole, Ajay Jain, Jonathan T Barron, and Ben Mildenhall. Dreamfusion: Text-to-3d using 2d diffusion. In *ICLR*, 2023. 2, 3, 4
- [61] Chenyang Qi, Xiaodong Cun, Yong Zhang, Chenyang Lei, Xintao Wang, Ying Shan, and Qifeng Chen. Fatezero: Fusing attentions for zero-shot text-based video editing. In *ICCV*, 2023. 2, 3
- [62] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *ICML*, 2021. 5
- [63] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with clip latents. *arXiv Preprint*, 2022. 1
- [64] Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. Zero-shot text-to-image generation. In *ICML*, 2021.
- [65] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image syn-

- thesis with latent diffusion models. In *CVPR*, 2022. 1, 2
- [66] Ludan Ruan, Y Ma, Huan Yang, Huiguo He, Bei Liu, Jianlong Fu, Nicholas Jing Yuan, Qin Jin, and Baining Guo. Mm-diffusion: Learning multi-modal diffusion models for joint audio and video generation. 2023 ieee. In *CVPR*, 2022. 2
- [67] 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, et al. Photorealistic text-to-image diffusion models with deep language understanding. *arXiv Preprint*, 2022. 1
- [68] Chitwan Saharia, Jonathan Ho, William Chan, Tim Salimans, David J Fleet, and Mohammad Norouzi. Palette: Image-to-image diffusion models. In *ACM SIGGRAPH*, 2022. 1
- [69] Flavio Schneider, Ojasv Kamal, Zhijing Jin, and Bernhard Schölkopf. Moûsai: Text-to-music generation with long-context latent diffusion. *arXiv Preprint*, 2023. 1
- [70] Shelly Sheynin, Adam Polyak, Uriel Singer, Yuval Kirstain, Amit Zohar, Oron Ashual, Devi Parikh, and Yaniv Taigman. Emu edit: Precise image editing via recognition and generation tasks. In *CVPR*, 2024. 2
- [71] Uri Singer, Adam Polyak, Ethan Fetaya, Jonathan Berant, Yaniv Hoshen, and Ronen Shalev. Make-a-video: Text-to-video generation without text-video data. In *ICLR*, 2023. 1, 2
- [72] Uriel Singer, Amit Zohar, Yuval Kirstain, Shelly Sheynin, Adam Polyak, Devi Parikh, and Yaniv Taigman. Video editing via factorized diffusion distillation. In *ECCV*, 2024. 2
- [73] Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. In *ICLR*, 2021. 1
- [74] Kun Su, Xiulong Liu, and Eli Shlizerman. Audeo: Audio generation for a silent performance video. In *NeurIPS*, 2020. 2
- [75] Zineng Tang, Ziyi Yang, Mahmoud Khademi, Yang Liu, Chenguang Zhu, and Mohit Bansal. Codi-2: In-context interleaved and interactive any-to-any generation. In *CVPR*, 2024. 2
- [76] Zineng Tang, Ziyi Yang, Chenguang Zhu, Michael Zeng, and Mohit Bansal. Any-to-any generation via composable diffusion. In *NeurIPS*, 2024. 2
- [77] Zeyue Tian, Zhaoyang Liu, Ruibin Yuan, Jiahao Pan, Qifeng Liu, Xu Tan, Qifeng Chen, Wei Xue, and Yike Guo. Vid-muse: A simple video-to-music generation framework with long-short-term modeling. In *CVPR*, 2025. 1
- [78] Narek Tumanyan, Michal Geyer, Shai Bagon, and Tali Dekel. Plug-and-play diffusion features for text-driven image-to-image translation. In *CVPR*, 2023. 2
- [79] Kai Wang, Shijian Deng, Jing Shi, Dimitrios Hatzinakos, and Yapeng Tian. Av-dit: Efficient audio-visual diffusion transformer for joint audio and video generation. *arXiv Preprint*, 2024. 2
- [80] Wen Wang, Yan Jiang, Kangyang Xie, Zide Liu, Hao Chen, Yue Cao, Xinlong Wang, and Chunhua Shen. Zero-shot video editing using off-the-shelf image diffusion models. *arXiv Preprint*, 2023. 2
- [81] Yongqi Wang, Wenxiang Guo, Rongjie Huang, Jiawei Huang, Zehan Wang, Fuming You, Ruiqi Li, and Zhou Zhao. Frieren: Efficient video-to-audio generation network with rectified flow matching. In *NeurIPS*, 2024. 2
- [82] Yuancheng Wang, Zeqian Ju, Xu Tan, Lei He, Zhizheng Wu, Jiang Bian, et al. Audit: Audio editing by following instructions with latent diffusion models. In *NeurIPS*, 2023. 1, 2, 3
- [83] Chenfei Wu, Jianlong Bao, Dongdong Chen, Weining Zhang, Luowei Zhao, Lu Yuan, Dong Zhang, and Fang Wen. Nuwa: Visual synthesis pre-training for neural visual world creation. In *ECCV*, 2022. 1
- [84] Jay Zhangjie Wu, Yixiao Ge, Xintao Wang, Stan Weixian Lei, Yuchao Gu, Yufei Shi, Wynne Hsu, Ying Shan, Xiaoju Qie, and Mike Zheng Shou. Tune-a-video: One-shot tuning of image diffusion models for text-to-video generation. In *ICCV*, 2023. 3
- [85] 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. 5
- [86] Shuai Yang, Yifan Zhou, Ziwei Liu, and Chen Change Loy. Rerender a video: Zero-shot text-guided video-to-video translation. In *SIGGRAPH Asia*, 2023. 1, 3
- [87] Guy Yariv, Itai Gat, Sagie Benaim, Lior Wolf, Idan Schwartz, and Yossi Adi. Diverse and aligned audio-to-video generation via text-to-video model adaptation. In *AAAI*, 2024. 5
- [88] Jaehong Yoon, Shoubin Yu, and Mohit Bansal. Raccoon: A versatile instructional video editing framework with auto-generated narratives. *arXiv Preprint*, 2024. 3
- [89] Shoubin Yu, Difan Liu, Ziqiao Ma, Yicong Hong, Yang Zhou, Hao Tan, Joyce Chai, and Mohit Bansal. Veggie: Instructional editing and reasoning video concepts with grounded generation. *arXiv Preprint*, 2025. 3
- [90] Lin Zhang, Shentong Mo, Yijing Zhang, and Pedro Morgado. Audio-synchronized visual animation. In *ECCV*, 2024. 2
- [91] Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. Adding conditional control to text-to-image diffusion models. In *ICCV*, 2023. 1, 2, 6
- [92] Yabo Zhang, Yuxiang Wei, Dongsheng Jiang, Xiaopeng Zhang, Wangmeng Zuo, and Qi Tian. Controlvideo: Training-free controllable text-to-video generation. In *ICLR*, 2024. 2, 3, 6, 7, 8