

Audio-sync Video Instance Editing with Granularity-Aware Mask Refiner

Haojie Zheng^{1,2†} Shuchen Weng^{2,3†} Jingqi Liu^{1,2} Siqi Yang⁵
Boxin Shi^{3,4‡} Xinlong Wang²

¹School of Software and Microelectronics, Peking University ²Beijing Academy of Artificial Intelligence

³State Key Lab of Multimedia Info. Processing, School of Computer Science, Peking University

⁴Nat'l Eng. Research Ctr. of Visual Tech., School of Computer Science, Peking University

⁵Institute for Artificial Intelligence, Peking University

{suimu, liujingqi}@stu.pku.edu.cn {scweng, wangxinlong}@baai.ac.cn

{yousiki, shiboxin}@pku.edu.cn

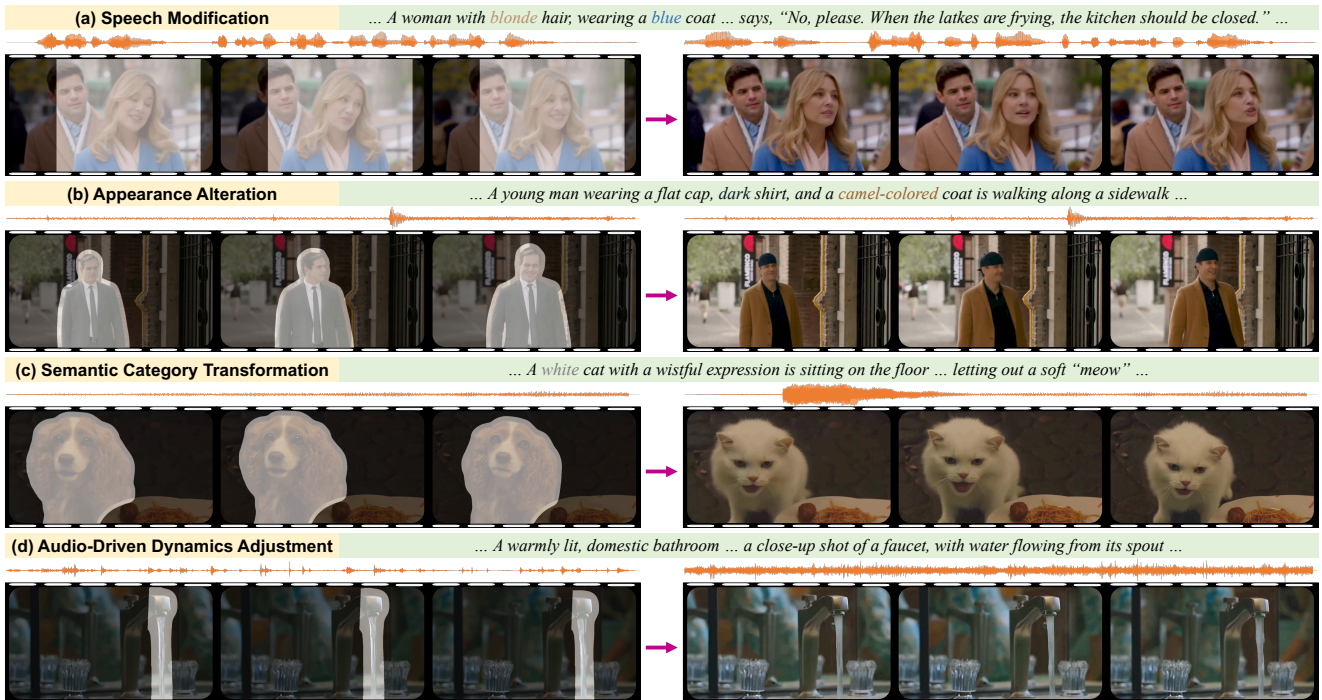


Figure 1. AVI-Edit effectively edits audio-sync video instance based on a coarse instance mask to indicate the target instance and a text description to specify the edit direction. We demonstrate this capability across four representative scenarios: (a) modifying the speech of the target woman while preserving her appearance; (b) altering the appearance of the male character while preserving his original speech; (c) changing the dog into a cat while transforming its vocalization; and (d) adjusting the dynamics of the water flow purely via audio cues. Generated videos and accompanying audios are available in the supplementary material.

Abstract

Recent advancements in video generation highlight that realistic audio-visual synchronization is crucial for engaging content creation. However, existing video editing methods largely overlook audio-visual synchronization and lack the fine-grained spatial and temporal controllability required for precise instance-level edits. In this paper, we propose AVI-Edit, a framework for audio-sync video instance editing. We propose a granularity-aware mask refiner that it-

eratively refines coarse user-provided masks into precise instance-level regions. We further design a self-feedback audio agent to curate high-quality audio guidance, providing fine-grained temporal control. To facilitate this task, we additionally construct a large-scale dataset with instance-centric correspondence and comprehensive annotations. Extensive experiments demonstrate that AVI-Edit outperforms state-of-the-art methods in visual quality, condition following, and audio-visual synchronization. Project page: <https://hjzheng.net/projects/AVI-Edit/>.

† Equal contribution.

‡ Corresponding author.

1. Introduction

Video editing methods provide users with a powerful approach to content creation (*e.g.*, modifying the appearance of a specific person). Recent commercial models (*e.g.*, Sora-2 [3] and Veo3 [4]) demonstrate that realistic audio is crucial for creating engaging experiences. This advancement naturally leads to an urgent need for methods that perform instance-level edits on videos while preserving this crucial audio-visual synchronization.

However, the vast majority of existing video editing models [9, 15, 16, 65, 70] focus exclusively on the visual features, breaking the original audio-visual synchronization. To introduce audio cues, AvED [41] introduces a cross-modal contrastive scheme to address the synchronization, but it mainly focuses on scene-level alignment. In contrast, while Object-AVEdit [21] achieves object-specific control, its inversion-regeneration paradigm inherently lacks temporal controllability (*e.g.*, specifying the precise timing of an event). While recent work [64] demonstrates that audio is a fine-grained temporal guidance for video generation, this approach still remains largely unexplored for audio-sync video editing.

In this paper, we propose **AVI-Edit**, a framework for **Audio-sync Video Instance Editing**. As illustrated in Fig. 1, given an instance mask to indicate the target instance and a text description to specify the edit direction, AVI-Edit seamlessly modifies the target instance and its accompanying audio, while preserving the original background and non-target audio components. This approach thereby enables a range of applications, including modifying the speech of an active speaker (Fig. 1 (a)), altering the appearance of the person (Fig. 1 (b)), changing the semantic category of target instance (Fig. 1 (c)), or adjusting the dynamics of the subject purely with audio cues (Fig. 1 (d)).

We build our framework upon the pretrained Wan2.2-5B [60] to leverage its strong generative priors in generating temporally consistent and visually appealing videos. To achieve precise spatial control at the instance level, we propose the granularity-aware mask refiner with the similar diffusion transformer architecture of the video backbone, which receives a precision factor and references video and audio tokens to estimate precise edited regions. We further design the self-feedback audio agent to run a separate-generate-remix-rework cycle, which finally produces the curated audio tokens to provide explicit temporal guidance and control the event timing of edited videos.

To facilitate model training and evaluation, we further present AVISET, split into 71k training, 1k validating, and 1k testing videos with accompanying audios. All clips are carefully filtered to ensure each contains one primary and non-silent sounding instance. We annotate all splits with instance masks and scene-level text descriptions, while the testing set additionally includes paired original–edited text

instructions for fair evaluation.

Our contributions are summarized as follows:

- We propose an audio-sync video instance editing framework with fine-grained spatial and temporal control, along with a tailored dataset for training and evaluation.
- We propose the granularity-aware mask refiner to estimate precise instance mask, guided by the precision factor for precise spatial control over instance regions.
- We design the self-feedback audio agent as a separate-generate-remix-rework pipeline, curating accompanying audios to provide explicit temporal control for videos.

2. Related work

2.1. Visual-only Video Generation and Editing

With the rapid development of video foundation models [33, 60, 71], video generation models demonstrate the great potential for animation creation [26], advertising [7], cinematic visual effects [73, 74], and game development [11]. This has motivated researchers to explore high-freedom editing techniques. Early editing methods [30, 39, 47, 66] primarily adapt pre-trained text-to-image models [34, 51]. However, as these backbones lack dynamic video priors, they inevitably struggle with temporal consistency. To address this, recent works [9, 29, 65, 70] propose dedicated video editing architectures. These methods achieve higher-quality and more consistent edits by incorporating specialized mechanisms (*e.g.*, dual-branch encoders [9, 70] and auxiliary conditions [29, 65]) for precise spatial control. However, these methods are fundamentally visual-only. They entirely neglect the accompanying audio cues, thereby breaking the original audio-visual synchronization that is crucial for an engaging video experience.

2.2. Audio-guided Image Generation and Editing

Audio conveys rich semantic and affective cues that can serve as a powerful conditioning signal. Early audio-guided works focus on specific domains (*e.g.*, music [35] and speech [45]). With the emergence of diffusion models, these works begin to project audio into a shared embedding space to align audio semantics with visual content. Lee *et al.* [38] propose the audio attention and sentence attention to represent the rich audio characteristics. AudioToken [72] introduces an audio embedder to transform audio into text tokens used as conditioning for the pretrained text-to-image model [51]. However, constrained by the static nature of images, these methods reduce audio to a holistic semantic feature. They inherently overlook the temporal cues and the frame-by-frame synchronization natural to the visual world.

2.3. Audio-sync Video Generation and Editing

Modeling the natural synchronization between audio and video is a long-standing challenge. Early approaches

achieve initial success for audio-sync video generation [22, 28, 36] based on GANs [19, 31]. Recent generation models are built upon diffusion architecture [43, 61], which use multi-modal cross-attention mechanisms to naturally synthesize synchronized audios. In contrast, editing models extract fundamental audio signals (*e.g.*, magnitude modulation [55] and distance guidance [37]) to modify the global semantics of videos. Towards frame-by-frame temporal synchronization, recent works [21, 41] invert the pretrained generation models and edit the latent code in a training-free manner. Despite this, these approaches still struggle to achieve full control (*e.g.*, instance-level editing with free-form temporal controllability). Therefore, we propose AVI-Edit to bridge this gap by introducing the granularity-aware mask refiner and the self-feedback audio agent.

3. Dataset

While existing datasets [13, 23, 36] provide large-scale audio-visual correspondence, they typically lack the instance-level focus. To facilitate model training and evaluation, we construct AVISET, a new dataset providing instance-centric correspondence and comprehensive annotations, tailored for the audio-sync video instance editing.

Data sources. We construct our initial data pool from two main sources: (i) 5,529 hours of publicly available videos, including MovieBench [67] (69h), Condensed Movies [6] (1,270h), Short-Films-20K [24] (3,582h), and VGGSound [13] (608h); and (ii) 5,607 hours of open-access online videos (*e.g.*, from YouTube). All collected videos retain their accompanying audios. **Video filtering.** Following prior work [10, 14, 63], we first use PySceneDetect [12] to segment the original videos into single-shot clips. We then use RAFT [56] to compute optical flow to discard static or low-motion clips. The aesthetic predictor [53] is further adopted to measure the visual appeal and filter out clips with low visual quality.

Audio filtering. We first use Audiobox-aesthetics [57] to assess the quality of the accompanying audios, discarding low-quality clips. Next, we adopt Qwen-Omni [68] to classify all remaining clips into human speech or non-speech event clips for content-specific filtering. For human speech clips, we apply TalkNet [8] to detect the active speaker’s bounding box, and adopt Scribe [2] to identify speech segments of each speaker, allowing us to discard clips with multiple simultaneous speakers. For non-speech event clips, we further use Qwen-Omni [68] to filter out clips containing multiple sounding instances, and then predict the semantic category label (*e.g.*, dog barking) for the single instance in the retained clips.

Instance masking. Using the speaker’s bounding box or the event’s semantic category, we employ Grounded-SAM-2 [49] to generate instance masks for each clip. We further discard samples with small or empty mask regions, thereby

filtering out cases where the sounding instance is not clearly visible or is off-screen (*e.g.*, a voice-over). After this stage, each video in the pool contains only one primary sounding instance, paired with its corresponding instance masks. All retained clips (for both human speech and non-speech event paths) are then pooled together, ensuring the final dataset is mixed for a general-purpose generation model.

Text annotations. We adopt Qwen-VL [5] to generate a scene-level text description for each retained clip. For the testing set, we provide additional paired original–edited text instructions. Specifically, we prompt the LLM [69] with the clip’s original text description and its predicted semantic category to generate a plausible editing instruction. Finally, all generated instruction pairs are manually verified by volunteers to ensure their naturalness.

Data statistics. The AVISET includes 71k training, 1k validating, and 1k testing clips, totaling over 197 hours. Each clip is approximately 10 seconds, presented at 720P resolution (24 FPS). We annotate each clip with an instance mask and a scene-level text description, while the testing set additionally includes paired original–edited text instructions.

4. Methodology

AVI-Edit consists of three main components. We first introduce the audio-sync video backbone (Sec. 4.1) with its architecture and formulation. We then detail the granularity-aware mask refiner (Sec. 4.2), which provides precise spatial control over the edited instance. Finally, we describe the self-feedback audio agent (Sec. 4.3), which generates fine-grained temporal guidance for the entire editing process.

4.1. Audio-Sync Video Backbone

As shown in Fig. 2 (a), the audio-sync video backbone generates audio-sync video editing results based on a curated audio a , an instance mask m , and a text description y .

Video compression. To enable editing in the compact latent space, we encode the original clip x with a pretrained VAE encoder \mathcal{E} into a clean latent code as $z = \mathcal{E}(x)$. A corresponding VAE decoder \mathcal{D} is used to reconstruct the video clip from the latent code as $x = \mathcal{D}(z)$.

Flow matching formulation. We train our video diffusion transformer v_θ using the flow matching objective [42], and define the probability path \hat{z}_t at timestep t as a linear combination of the latent code z and Gaussian noise $\epsilon \sim \mathcal{N}(0, \mathbf{I})$:

$$\hat{z}_t = tz + (1 - t)\epsilon, \quad (1)$$

where $t \in [0, 1]$ is the timestep, with $\hat{z}_0 = \epsilon$ and $\hat{z}_1 = z$.

Background preservation. To avoid producing undesirable alterations and inconsistencies in background regions, we adopt the instance mask m to indicate the edited regions, and downsample it to compose the interpolated \hat{z}_t with the clean input video tokens z :

$$z_t = \hat{z}_t \odot \hat{m} + z \odot (1 - \hat{m}), \quad (2)$$

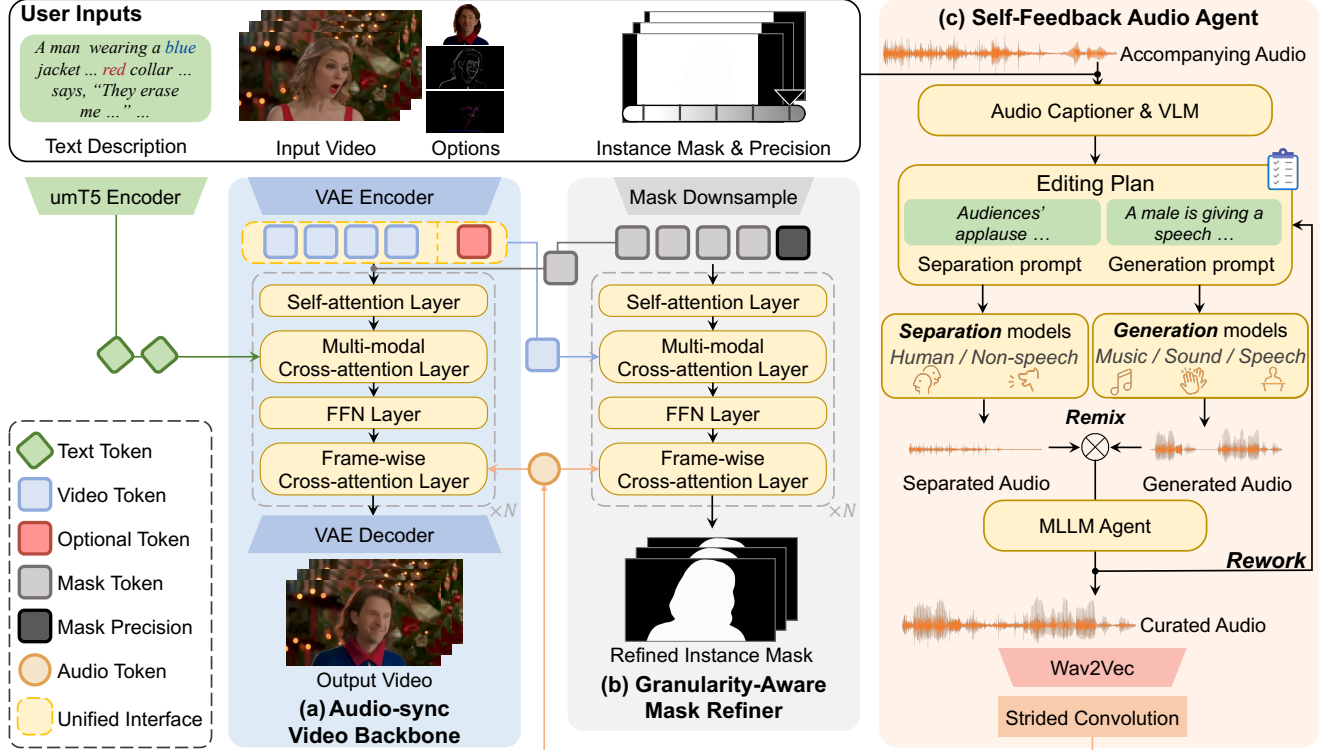


Figure 2. Illustration of our AVI-Edit framework. Given multi-modal user inputs, AVI-Edit separately encodes them into latent tokens, and generates audio-sync video instance editing results with the following components: (a) The audio-sync video backbone (Sec. 4.1) is built upon the pretrained video diffusion transformer, which utilizes a frame-wise cross-attention layer to understand audio cues and generates editing results based on the masks, text, and optional contexts. (b) The granularity-aware mask refiner (Sec. 4.2) features a similar architecture and refines the coarse instance mask into a precise one, which utilizes visual semantics by replacing text tokens with video tokens, and is guided by the precision factor and audio cues. (c) The self-feedback audio agent (Sec. 4.3) is designed with a separate–generate–remix–rework pipeline that integrates off-the-shelf audio components to robustly handle diverse scenarios, which provides the resulting curated audio tokens as the temporal guidance for the video and mask components.

where \hat{m} is the downsampled instance binary mask.

Audio-sync diffusion transformer. To leverage priors of the pretrained video generation model, we build the video backbone upon the Wan2.2 architecture [60]. We first concatenate the downsampled instance mask \hat{m} with the noisy latent code z_t as video tokens. After that, these tokens are sequentially processed by a self-attention layer to interact with long-range context, a multi-modal cross-attention layer to understand the text description, and a feed-forward network to extract features for N Diffusion Transformer (DiT) blocks. To enable audio-sync video editing, we additionally introduce a frame-wise cross-attention to each DiT block, which receives curated audio tokens a and guides the temporal alignment with the video latent codes.

Optional control contexts. To enable video editing results precisely satisfying the user’s intention, we design a unified interface to incorporate optional control contexts c (e.g., scribble, pose, and reference image). These conditions are encoded to context tokens using the pretrained VAE encoder \mathcal{E} , and then injected through element-wise addition for scribble and pose, and concatenation for the reference

image. This provides fine-grained controllability.

Training and inference process. During training, we fine-tune the video diffusion transformer v_θ to estimate the velocity field $v_t = z - \epsilon$, by minimizing the standard flow matching objective:

$$\mathcal{L}_{\text{fm}} = \mathbb{E}_{\epsilon, t, z, \hat{m}, y, a, c} [\|v_\theta(z_t, t, \hat{m}, y, a, c) - v_t\|^2]. \quad (3)$$

During inference, we transform the Gaussian noise ϵ into the edited clean latent code z_1 by solving:

$$dz_t/dt = v_\theta(z_t, t, \hat{m}, y, a, c), \quad (4)$$

from $t = 0$ to $t = 1$ using a numerical solver.

4.2. Granularity-Aware Mask Refiner

The instance mask allows users to specify the target instance to be edited. However, user-provided instance masks are typically inaccurate (e.g., bounding box). As presented in Fig. 2 (b), we propose the Granularity-Aware Mask Refiner (GAMR) to refine these coarse instance mask, thereby achieving precise instance-level video editing.

Mask granularity. To quantitatively describe the granularity of instance mask, we first introduce a precision factor $p \in [0, P]$, which measures the degree of alignment between the target instance region and the provided mask. Conceptually, $p = P$ corresponds to maximal degradation (e.g., a bounding box), while $p = 0$ denotes the precise instance contour m_{gt} .

Diffusion-based refinement. We implement the GAMR with a similar architecture of video diffusion transformer, designed to predict the accurate instance mask. To understand the granularity of the instance mask, the precision factor p is linearly encoded and injected into GAMR through the AdaLN and Gate in each diffusion transformer block. The curated audio tokens are also incorporated via frame-wise cross-attention, enabling the GAMR to infer instance masks aligned with event timing. In contrast, the text tokens in the multi-modal cross-attention layers are replaced with video tokens, which enables the GAMR to further reason about instance masks based on the visual semantics.

Training strategy. To simulate the inaccuracy of user-provided masks, we sample input videos x and their corresponding precise instance masks m_{gt} from our AVISSET, then degrade each mask to an arbitrary granularity level p to construct training pairs. Specifically, the degradation is implemented using a set of predefined Gaussian blur kernels:

$$m_p = \text{GaussianBlur}(m_{\text{gt}}, k_p, \sigma_p), \quad (5)$$

where the kernel size k_p and standard deviation σ_p are determined by the precision factor p . We introduce a mask refinement objective [40] that adopts focus loss to measure the discrepancy between the downsampled generated mask \hat{m} and ground-truth mask \hat{m}_{gt} to train the GAMR:

$$\begin{aligned} \mathcal{L}_{\text{mask}} = & -\alpha \hat{m}_{\text{gt}}(1 - \hat{m})^\gamma \log(\hat{m}) \\ & - (1 - \alpha)(1 - \hat{m}_{\text{gt}}) \hat{m}^\gamma \log(1 - \hat{m}), \end{aligned} \quad (6)$$

where $\alpha = 0.25$ and $\gamma = 2.0$ are hyperparameters.

Inference process. The GAMR estimates the refined mask \hat{m} at each step of the ODE solving process. Specifically, this iterative refinement is precision-aware for each inference step. At the first step ($k = 0$), we provide the downsampled user-provided mask \hat{m}_p as \hat{m}_p^0 and its precision factor p (e.g., $p = P$ for a bounding box). For all subsequent steps ($k > 0$), we iteratively feed the estimated mask \hat{m}_p^{k-1} from the previous step as the new input mask \hat{m}_p^k . The precision factor p for these steps is adjusted according to a predefined degradation schedule. At each step k , the concurrently refined mask \hat{m}_p^k is then used as the definitive mask \hat{m} for the video generation backbone (Eq. (4)).

4.3. Self-Feedback Audio Agent

To provide the temporal guidance for audio-sync video editing, we design a dedicated audio agent to curate the

original audio a_{orig} . As illustrated in Fig. 2 (c), the audio agent is guided by the text description y and the associated visual context. It is designed with a separate-generate-remix-rework pipeline to integrate a collection of off-the-shelf audio processing components, thereby enabling robust handling of diverse scenarios (e.g., human speech and non-speech events).

Audio separation and generation. To obtain a comprehensive understanding of the original audio a_{orig} , we employ a captioner [68] to translate the audio into a detailed text description and summarize its semantic content c_{sem} (e.g., sound events and environment atmosphere). Next, we leverage a VLM [5] to reason about the detailed editing plan:

$$(c^{\text{sep}}, c^{\text{gen}}) = \text{VLM}([x, m_p, c_{\text{sem}}, y]), \quad (7)$$

where c^{sep} and c^{gen} are text descriptions for the audio component to be remained from the separation (e.g., applause) and to be generated (e.g., male speech), respectively. Guided by this editing plan, the agent parses these descriptions to select the most suitable separation (T_s^{sep}) and generation (T_g^{gen}) models among a collection of off-the-shelf models, and applies them to obtain the remaining (a^{sep}) and generated (a^{gen}) components:

$$a^{\text{sep}} = T_s^{\text{sep}}(a_{\text{orig}}, c^{\text{sep}}), \quad a^{\text{gen}} = T_g^{\text{gen}}(c^{\text{gen}}). \quad (8)$$

Model collection. As human speech is best handled by specialized models, we define separation component as a collection of domain-specific models (i.e., human speech [27] and non-speech events [17]) as $T^{\text{sep}} = \{T_s^{\text{sep}}\}_{s \in \mathcal{S}}$, where the domain set is $\mathcal{S} = \{\text{speech}, \text{non-speech}\}$. As audio can also be decomposed into distinct tracks [58], we define our generation components as $T^{\text{gen}} = \{T_g^{\text{gen}}\}_{g \in \mathcal{G}}$, where the track set is $\mathcal{G} = \{\text{speech}, \text{music}, \text{sound}\}$. These are implemented with text-to-speech, text-to-music, and text-to-sound models from ElevenLabs [1].

Remix and rework judgment. After obtaining the separated component a^{sep} and generated component a^{gen} , we remix them to produce the curated audio a [50]. To prevent the degraded editing quality, we employ a multimodal large language model (MLLM) [68] to evaluate the overall perceptual quality q of the remixed curated audio (e.g., whether it sounds natural and realistic). Only remixed audios with quality score exceeding a predefined threshold τ are accepted. When the judging MLLM rejects a remixed result, it generates an improvement instruction ($\hat{c}^{\text{sep}}, \hat{c}^{\text{gen}}$), providing feedback to both the separation model T^{sep} (e.g., female speech remains audible) and the generation model T^{gen} (e.g., synthesized male speech volume is insufficient), enabling an iterative refinement loop. The rework loop continues until the remixed audio’s quality score surpasses τ or a maximum number of iterations is reached.



Figure 3. Qualitative comparison results with state-of-the-art methods for audio-sync video instance editing.

5. Experiment

Training details. We initialize our video backbone and GAMR with pre-trained weights from Wan2.2-5B [60], freezing the spatial-temporal VAE and fine-tuning the diffusion Transformer to minimize the joint flow matching objective (Eq. (3)) and mask refiner objective (Eq. (6)):

$$\mathcal{L} = \mathcal{L}_{\text{fm}} + \lambda \mathcal{L}_{\text{mask}}, \quad (9)$$

where the hyperparameter $\lambda = 1.0$ in our experiments. We train our AVI-Edit framework for 160k steps on 8 NVIDIA A800 GPUs, using a spatial resolution of 720p. We use the Adam optimizer [32] with a learning rate of 2×10^{-5} .

Evaluation Datasets. In addition to our proposed AVISET for the general audio-sync video editing task, we further evaluate our AVI-Edit framework on the publicly available AvED-Bench [41] to demonstrate its superior performance and generalization capability. To enable effective instance-level evaluation on AvED-Bench, we annotate instance masks according to the procedure described in Sec. 3. For evaluation, we randomly sample 100 video clips from each dataset’s test split.

5.1. Comparison with State-of-the-art Methods

As audio-sync video instance editing is an emerging research task, the comparison methods are still under development. We compare our framework with two recent state-of-

the-art methods, AvED [41] and Ovi [43]. We re-implement AvED [41] strictly following its original settings, due to its code being unavailable. As Ovi [43] is a generation model, we adapt it for editing using a common zero-shot inpainting strategy, where only the mask region is replaced with Gaussian noise. Furthermore, we implement a sequential baseline (VACE-Foley), where VACE [29] edits the video first, followed by Hunyuan-Foley [54] to generate the accompanying audio. All baselines are fine-tuned on the AVISET.

Quantitative comparisons. We quantitatively evaluate performance across three main aspects: (i) Visual quality is examined using Fréchet Video Distance (FVD) [59] and Inception Score (IS) [52]. (ii) Frame consistency (FC) is measured by calculating similarity between consecutive frames using CLIP [48]. (iii) Alignment is assessed via text-video (TC) with CLIP-XL [62], audio-video (AC) with Image-Bind [25], and lip motion synchronization (*i.e.*, Sync-C and Sync-D) [18]. Notably, since AvED-Bench [41] does not contain human talking videos, the lip motion synchronization metrics are not applicable for it. Since the baseline methods are not designed to handle coarse masks, we evaluate all methods using ground-truth instance masks and a text description to ensure a fair comparison. As presented in Tab. 1, our framework demonstrates significant advantages on both datasets, outperforming state-of-the-art methods on most quantitative metrics. Metrics details are pro-

Table 1. Quantitative experiment results of comparison with state-of-the-art methods. \uparrow (\downarrow) means higher (lower) is better. Throughout the paper, best performances are highlighted in **bold**.

Method	AVISSET							AvED-Bench				
	FVD \downarrow	IS \uparrow	FC (%) \uparrow	TC (%) \uparrow	AC (%) \uparrow	Sync-C \uparrow	Sync-D \downarrow	FVD \downarrow	IS \uparrow	FC (%) \uparrow	TC (%) \uparrow	AC (%) \uparrow
AvED	364.69	1.104	95.03	23.69	23.31	1.69	11.80	413.82	1.118	94.89	24.59	20.45
Ovi	407.08	1.122	96.47	25.83	26.68	4.00	9.12	504.14	1.122	95.65	25.21	21.49
VACE-Foley	391.64	1.115	96.60	25.92	26.45	1.72	10.37	405.76	1.108	95.74	25.17	21.46
AVI-Edit (Ours)	312.89	1.127	96.65	26.16	26.93	4.12	9.19	349.31	1.125	95.82	25.30	21.64

Table 2. Percentage (%) of user preference study results.

Method	AVISSET			AvED-Bench		
	AVS	TA	OP	AVS	TA	OP
AvED	2.40	3.20	1.60	3.60	4.80	4.00
Ovi	36.00	36.80	38.40	31.60	31.20	32.00
VACE-Foley	12.40	17.20	14.80	19.20	21.60	22.80
AVI-Edit (Ours)	49.20	42.80	45.20	45.60	42.40	41.20

Table 3. Percentage (%) of audio quality study results.

Rating	AF	RP	TAC
Failed	2.24	5.36	3.76
Borderline	6.32	9.12	7.60
Acceptable	8.48	19.68	14.96
Perfect	82.96	65.84	73.68

vided in the supplementary material.

Qualitative comparisons. We present qualitative comparisons with the aforementioned baselines [29, 41, 43, 54] to simulate a more realistic use case. For this evaluation, all methods are provided with the coarse mask and text description. As shown in Fig. 3, the left and right columns evaluate the case of human speech and non-speech events, respectively. As a result, AvED [41] struggles to maintain temporal consistency, resulting in noticeable flickering artifacts (Fig. 3 left and right). Ovi [43] suffers from visual inconsistencies (e.g., Fig. 3 right), while VACE-Foley [29, 54] fails in synthesizing the target speech (Fig. 3 left). In contrast, our AVI-Edit framework produces visually appealing video instance editing results with realistic accompanying audios.

User preference study. We additionally conduct a user study to evaluate human preference. Given ground-truth instance masks and text instructions, AVI-Edit and relevant methods produce their edit results. Participants are asked to evaluate these results from three aspects: (i) **Audio-Visual Synchronization (AVS)**: Participants are asked to select the video clip that is the best synchronized with the accompanying audio. (ii) **Text Alignment (TA)**: Participants are asked to select the video clip that best matches the input text instruction. (iii) **Overall Preference (OP)**: Participants are asked to select the video clip they prefer overall. We randomly selected 10 samples from each dataset, and recruited 25 volunteers to provide independent evaluations across all three aspects. As shown in Tab. 2, our model achieves the highest preference scores in all aspects.

Audio quality study. To assess the quality of audios pro-

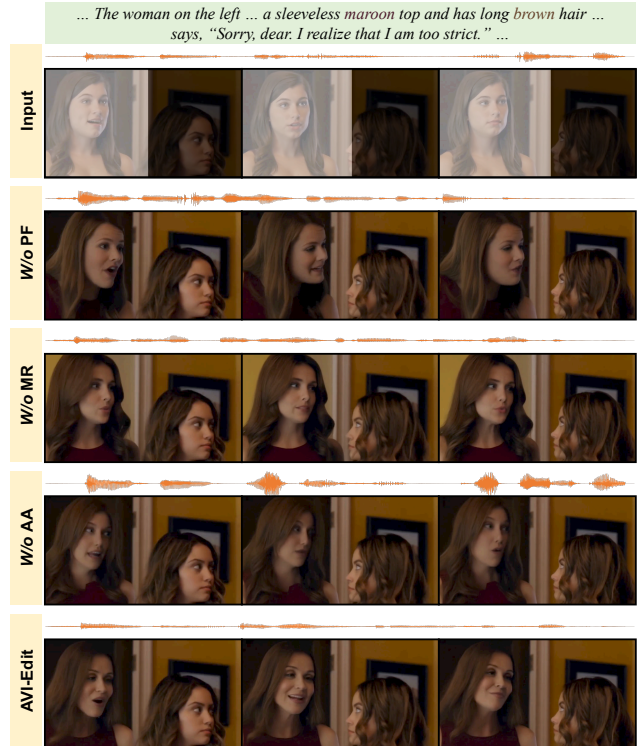


Figure 4. Visual results of ablation study with baseline variants.

duced by our self-feedback audio agent, we conduct an additional user study. We run our agent on 50 random clips from AVISSET, and recruit 25 volunteers to evaluate these results on the following three aspects: (i) **Audio Fidelity (AF)**: to evaluate the realism and clarity of the resulting audio; (ii) **Remaining Preservation (RP)**: to assess whether the non-target audio components are naturally preserved; and (iii) **Text-Audio Consistency (TAC)**: to measure the semantic alignment between the result audio and the user-provided text description. Volunteers rate each sample using a 4-point scale: “Failed”, “Borderline”, “Acceptable”, or “Perfect”. As shown in Tab. 3, our agent performs strongly, with over 91% (AF), 85% (RP), and 88% (TAC) of ratings being “Acceptable” or “Perfect”.

5.2. Ablation Study

We conduct ablation studies to evaluate the impact of key components of AVI-Edit. During ablation, the instance mask is randomly degraded into a coarse one to evaluate the

Table 4. Quantitative experiment results of ablation studies. \uparrow (\downarrow) means higher (lower) is better.

Method	AVISSET							AvED-Bench				
	FVD \downarrow	IS \uparrow	FC (%) \uparrow	TC (%) \uparrow	AC (%) \uparrow	Sync-C \uparrow	Sync-D \downarrow	FVD \downarrow	IS \uparrow	FC (%) \uparrow	TC (%) \uparrow	AC (%) \uparrow
w/o PF	354.43	1.119	96.49	26.07	26.50	4.12	9.43	490.92	1.118	95.47	25.06	21.51
w/o MR	372.44	1.107	96.32	25.68	26.38	4.07	9.36	539.83	1.103	95.29	24.96	21.45
w/o AA	342.75	1.114	96.54	25.84	25.97	3.83	9.61	445.56	1.105	95.36	25.13	21.22
AVI-Edit	335.32	1.121	96.63	26.13	26.77	4.18	9.27	402.74	1.122	95.58	25.17	21.63

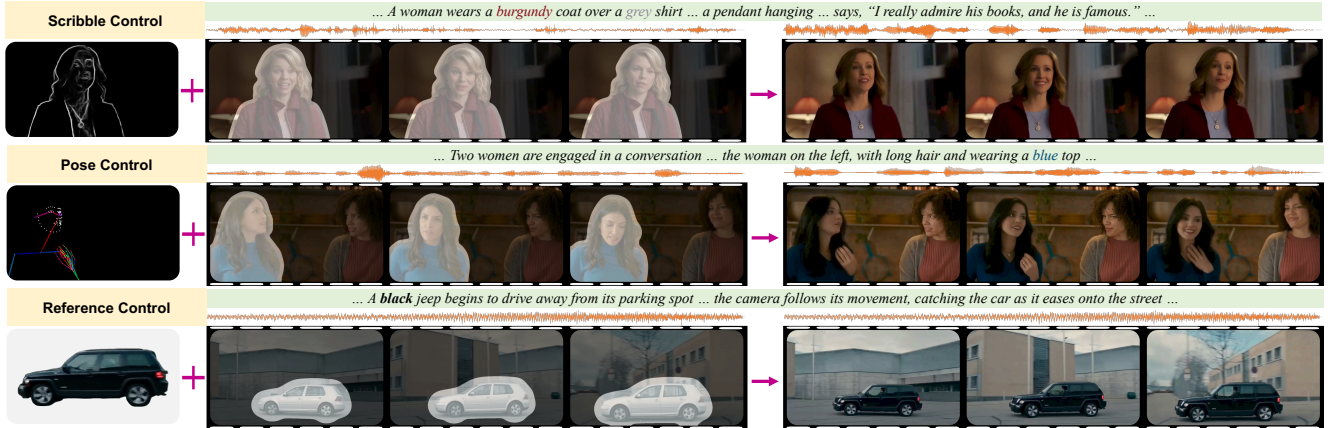


Figure 5. Versatile applications of AVI-Edit demonstrating its diverse controllability.

robustness of the mask refinement. The evaluation scores and editing results of the ablation study are presented in Tab. 4 and Fig. 4, respectively.

W/o Precision Factor (PF). This variant discards the precision factor that indicates the mask granularity, causing the granularity-aware mask refiner to lose its granularity guidance. This variant results in an imprecise estimation of the woman’s head regions within the mask (Fig. 4, second row).

W/o Mask Refiner (MR). This variant removes the granularity-aware mask refiner, forcing the model to rely solely on the initial coarse mask. As a result, this variant struggles to faithfully preserve background regions, altering the wall behind the head (Fig. 4, third row).

W/o Audio Agent (AA). We replace our self-feedback audio agent with a general audio editing model [44]. This leads to noisier audio guidance and degraded audio-visual synchronization (Tab. 4, Audio-C and Sync-C scores).

5.3. Application

As illustrated in Sec. 4.1, our video backbone is equipped with a unified interface that supports diverse, *optional* control contexts beyond the primary text and audio guidance. We demonstrate this additional capability in Fig. 5, where users can draw the scene with the scribble, manipulate the instance’s gesture with the pose, and customize the instance’s appearance with the reference image. This controllability allows users to satisfy highly specific intentions. Due to space limitations, we provide further demonstrations in the supplementary material (*e.g.*, instance insertion and removal, long video editing, and audio-sync generation).

6. Conclusion

In this paper, we present AVI-Edit, a framework designed for audio-sync video instance editing. We build the video generation backbone upon the pretrained text-to-video model [60], and equip it with frame-wise attention to receive audio guidance and a condition interface to edit the video content. We then introduce the precision factor to indicate the mask granularity, and adopt a similar architecture to build the granularity-aware mask refiner, which effectively refines the potentially coarse instance mask for accurate instance-level editing. We further propose the self-feedback audio agent, designed with a separate-generate-remix-rework pipeline to robustly handle scenarios of human speech and non-speech events. Extensive experiments demonstrate that AVI-Edit achieves state-of-the-art performance with diverse application scenarios.

Limitation. Since the target mask indicates an instance to be edited, editing videos with multiple target instances requires running AVI-Edit sequentially, processing each instance one-by-one. We leave the exploration of simultaneous multi-instance editing in the future work.

Acknowledgements

This work is supported by National Natural Science Foundation of China (Grant No. 62136001) and Beijing Major Science and Technology Project (Grant No. Z251100008125009). We thank all the insightful reviewers for the helpful suggestions, and the colleagues at Beijing Academy of Artificial Intelligence for their support.

7. Appendix

7.1. Additional Application Scenarios

As illustrated in Sec. 5.3, we additionally present application scenarios in Fig. 6 to demonstrate the versatility of our framework.

Instance insertion. By specifying target regions within a user-provided video, AVI-Edit can generate the corresponding instance and its accompanying audio according to the text description, seamlessly composited with the background regions (Fig. 6 first row, inserting a car with its accompanying audio into an empty field).

Instance removal. When an unwanted instance is specified, AVI-Edit removes the target instance and its accompanying audio from all frames based on the text description, resulting in a clean background (Fig. 6 second row, removing a plane flying over the sea and its accompanying audio).

Audio-sync video generation. In the extreme case where the instance mask is expanded to cover the full video frame, AVI-Edit discards all video content and accompanying audios, becoming an audio-sync generation model according to the text description (Fig. 6 third row, generating an audio-sync video from the provided text description).

Long-duration video editing. To enable AVI-Edit to edit long-duration videos, we incorporate a conditioning strategy during training to unmask k randomly selected frames. During inference, we generate the initial segment from noise. Subsequent segments are then generated using the last k frames of the preceding clip as visual guidance, enabling the iterative editing of arbitrary-length videos (Fig. 6 last row, consistently editing an instance spanning 241 frames).

7.2. Precision Factor Details

As illustrated in Sec. 4.2, as user-provided instance masks are typically inaccurate (*e.g.*, bounding box), we introduce a precision factor p to indicate the mask granularity.

In real applications, users typically provide the maximal precision factor $p = P$ for bounding boxes, or manually specify an intermediate p value to provide precise control for the target instance region based on the provided mask. To simulate the inaccuracy of user-provided masks, we randomly degrade each precise instance mask to an arbitrary granularity level p during training (Eq. (5)).

We further provide the architecture details of the Granularity-Aware Mask Refiner (GAMR) in Fig. 7, which demonstrates the approach of precision factor injection. Specifically, the precision factor p is linearly encoded and added to the timestep embedding to produce the modulation

Table 5. The IoU scores of different degradation schedules.

Schedule	Linear	Constant	Instant
IoU (%) [†]	71.87	68.01	76.23

parameters:

$$f^p = \text{Linear}(p) + \text{Linear}(t), \quad (10)$$

$$(\gamma, \beta, \alpha) = \text{Linear}(f^p), \quad (11)$$

where t is the timestep embedding and (γ, β, α) are estimated parameters for modulation. These parameters are then used to modulate the adaptive layer normalization (AdaLN) [46] followed by a self- or cross-attention:

$$h' = \text{Attention}(h \odot (\gamma + 1) + \beta), \quad (12)$$

where h is the mask feature before modulation. The resulting feature is then further modulated via the gating mechanism [46] with a residual skip connection:

$$\hat{h} = h + \alpha \odot h_{\text{out}}. \quad (13)$$

7.3. Iterative Mask Refinement Details

As illustrated in Sec. 4.2, the Granularity-Aware Mask Refiner (GAMR) is designed to iteratively refine the instance mask throughout the ODE solving process, guided by the precision factor p and a predefined degradation schedule. We investigate three distinct schedules: (i) *Linear degradation*. The precision factor p decreases linearly from its initial value down to 0 throughout the inference steps. (ii) *Constant degradation*. The precision factor p remains fixed at its initial value throughout the inference process. (iii) *Instant degradation*. The precision factor p is set to its initial value only for the first step (*i.e.*, $k = 0$). For all subsequent steps (*i.e.*, $k \geq 1$), the factor is instantly set to $p = 1$, allowing only subtle mask refinement.

As shown in Fig. 8, the linear and constant schedules are generally suboptimal, because the mask is largely optimized during the initial step ($k = 0$). If p decays too slowly, the AMR is guided by an overly-coarse mask granularity. Consequently, the model struggles in predicting the target boundary, causing unnecessary shrinkage in the refined mask and leading to unnatural visual content (*e.g.*, the woman’s face is rendered in profile to fit the overly small region). Therefore, the best schedule strategy is the simple yet effective instant degradation, where the first step quickly converges to the target region, while subsequent steps precisely refine the mask.

To validate this conclusion, we randomly sample 40 video clips from the test split of AVISSET and calculate the Intersection over Union (IoU) [20] of the predicted masks. As presented in Tab. 5, the instant degradation strategy achieves the best score, demonstrating that its predicted masks best match the ground truth.

[†] Equal contribution.

[‡] Corresponding author.



Figure 6. Additional application scenarios of AVI-Edit.

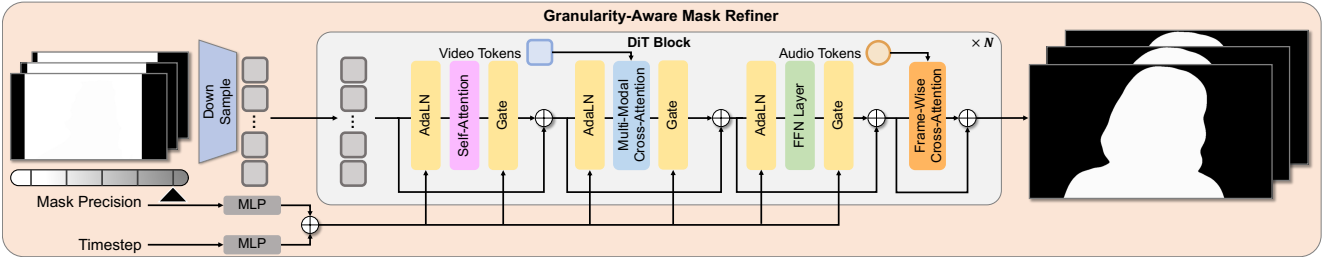


Figure 7. The detailed architecture of the granularity-aware mask refiner.

7.4. Self-Feedback Audio Agent Details

As introduced in Sec. 4.3, we design a dedicated audio agent to curate the accompanying audio. Since human speech separation is a mature technique that typically requires no text prompts, our prompt refinement loop primarily targets the more challenging non-speech separation and audio generation tasks. For clarity, we present a representative execution example in Fig. 9, illustrating the separate-generate-remix-rework pipeline, including plan initialization, separation, generation, remixing, and quality verification.

Judging MLLM. We adopt a judging MLLM to perform quality verification and evaluate the mixed audio across five dimensions: (i) *Separation accuracy* assesses whether the corresponding audio components are correctly retained from the original audio; (ii) *Generation accuracy* assesses whether the newly generated audio component matches the text description; (iii) *Acoustic harmony* assesses whether the remixed audio is acoustically harmonious; (iv) *Instruction adherence* assesses whether the remixed audio adheres

to the text instruction and editing goal; and (v) *Audio fidelity* assesses whether the remixed audio sounds realistic. We prompt the MLLM to score each dimension on a 3-point scale (*i.e.*, 0: Unacceptable, 1: Borderline acceptable, 2: Perfectly acceptable). Only remixed audios with a total quality score exceeding a predefined threshold $\tau = 7$ are accepted. Otherwise, the audio is rejected, and the MLLM generates an improvement instruction to rework with an iterative refinement loop.

Remix. We implement remix as an additive operator without optimization that composites the separation and generation audio components, as a *prerequisite* for final results.

Efficiency. We quantitatively assess the operational efficiency of our self-feedback audio agent on a subset of 200 randomly sampled video clips from the test set. On average, the agent requires only 1.67 rework iterations per clip to achieve satisfactory results, with a total processing time averaging 69.9 seconds. Specifically, the initial editing plan generation requires 27.3 seconds, leaving 42.6 seconds dedicated to the iterative self-feedback loop.




Figure 8. Visual comparison of the mask refinement process under different degradation schedules. **Top:** Original videos with coarse masks. **Rows 2-4:** Editing results with evolving refined masks.

Table 6. The inference time and peak VRAM of comparison with state-of-the-art methods.

Method	AvED	Ovi	VACE-Foley	AVI-Edit
Time (s) ↓	1140.5	232.6	1003.6	311.3
VRAM (GB) ↓	37.6	59.6	33.2	36.2

7.5. Inference Costs

We evaluate the inference time and peak VRAM of our framework and the comparison methods, averaged over 100 clips using a single NVIDIA A100 GPU. As shown in Tab. 6, AVI-Edit operates over $3\times$ faster than both AvED [41] and VACE-Foley [29, 54], while requiring significantly less peak VRAM than Ovi [43].

7.6. Evaluation Metrics Details

As described in Sec. 5.1, we adopt seven metrics to quantitatively evaluate performance and present the results in Tab. 1 of the main paper. For clarification, we detail these metrics as follows:

- **FVD** assesses the video quality, calculating the Fréchet video distance [59] between real and generated video features, extracted by the Inception Network.
- **IS** assesses the fidelity and diversity of generated frames, calculating the Inception score [52] using the softmax outputs of the Inception classifier.
- **FC** assesses the temporal consistency across video frames, calculating the cosine similarity between the embeddings of consecutive video frames, extracted by the CLIP visual encoder [48].
- **TC** assesses the text-video consistency, calculating the cosine similarity between the generated video and text embeddings extracted by the VideoCLIP-XL [62].
- **AC** assesses the audio-video consistency, calculating the cosine similarity between the generated audio and video embeddings extracted by the ImageBind [25].
- **Sync-C/D** are metrics that assess the audio-lip synchronization. They first calculate the Euclidean distance between audio embeddings and mouth-region video embed-

dings, extracted by the SyncNet [18]. Next, Sync-C is defined as the difference between the minimum and the median of these Euclidean distances, and Sync-D is defined as the minimum of all computed Euclidean distances.

7.7. Organization of Supplementary Video

We provide a supplementary video to dynamically showcase our audio-sync instance editing results. The video is structured as follows: (i) **Representative application scenarios:** We demonstrate four representative editing scenarios to highlight the diverse and compelling applications of our framework (Fig. 1). (ii) **Additional application scenarios:** We showcase four extended application scenarios to further demonstrate the versatility of our framework (Fig. 6). (iii) **Optional control contexts:** We visualize the integration of optional control contexts (*i.e.*, scribble, pose, and reference image) to facilitate fine-grained editing (Fig. 5). (iv) **Comparison and ablation study:** Finally, we provide audio-sync instance editing comparisons against state-of-the-art methods (Fig. 3), followed by visual ablations to verify the effectiveness of each component (Fig. 4).

Model Instruction



Vision Language Model

You are an expert visual captioner. You will be given a video, an instance mask, a summarized audio caption, and a user-provided text description. Your task is to edit the audio caption by removing the masked object's sound or to generate a new audio description that satisfies the user's intent.

(a) Audio Separation

First, analyze the provided mask, video, the original audio caption and the user-provided description to determine the removed instance and whether the removed instance is a human who is speaking or not.

Next, imagine the audio after removing that instance and generate a caption within 10 words, which should reflect the remaining audio after the instance is removed.

Output only two lines:

Line 1 — the remaining audio caption for separation.

Line 2 — “Human-speech” or “Non-speech”, representing the sound category of the removed audio components.

(b) Audio Generation

First, analyze the provided mask, video, and the user-provided description to determine the new target instance and classify its intended sound category.

Next, based on the video scene and the user's description, generate a scene-appropriate audio caption suitable for audio generation.

Output only two lines:

Line 1 — the new audio caption for generation.

Line 2 — “Speech”, “Music”, or “Effect”, representing the sound category of the generated audio components.



Judging Multi-Modal Language Model

You are an expert multimodal evaluator for an automatic audio editing system. Your goal is to evaluate the quality of the remixed audio and its adherence to a user's edit instruction. You should score the audio across five criteria and provide improved generative prompts if needed. You will be given the following 7 inputs: 1. Input Video (v): The video to be edited; 2. Accompanying Audio (a): The accompanying audio with the video; 3. Input Mask (m): A video identifying the target instance to be edited; 4. Remixed Audio (a'): The final, edited audio; 5. Text Description (y): The high-level text instruction describing the desired edits; 6. Separation Prompt: The text prompt used by the system to identify and separate non-target audio components. 7. Generation Prompt: The text prompt used by the system to generate the new target audio components.

You are an extremely harsh and critical evaluator whose primary job is to find flaws. You must rigorously compare the Original Audio (a) and Remixed Audio (a'). Then, evaluate a' based on the inputs (v, m, y) by assigning an integer score (0=Unacceptable, 1=Borderline Acceptable, 2=Perfectly Acceptable) to each of the five criteria. Calculate the Total Score (max 10); if the Total Score is less than 8, you should provide improved versions of the Separation Prompt and Generation Prompt; otherwise, output the original prompts as the edit is successful.

The five evaluation criteria are: (i) Separation accuracy assesses whether the corresponding audio components are correctly retained from the original audio. (ii) Generation accuracy assesses whether the newly generated audio component matches the text description. (iii) Acoustic harmony assesses whether the remix audio is acoustically harmonious. (iv) Instruction adherence assesses whether the remixed audio adheres to the text instruction and editing goal. (v) Audio fidelity assesses whether the remixed audio sounds realistic.

Overall Process

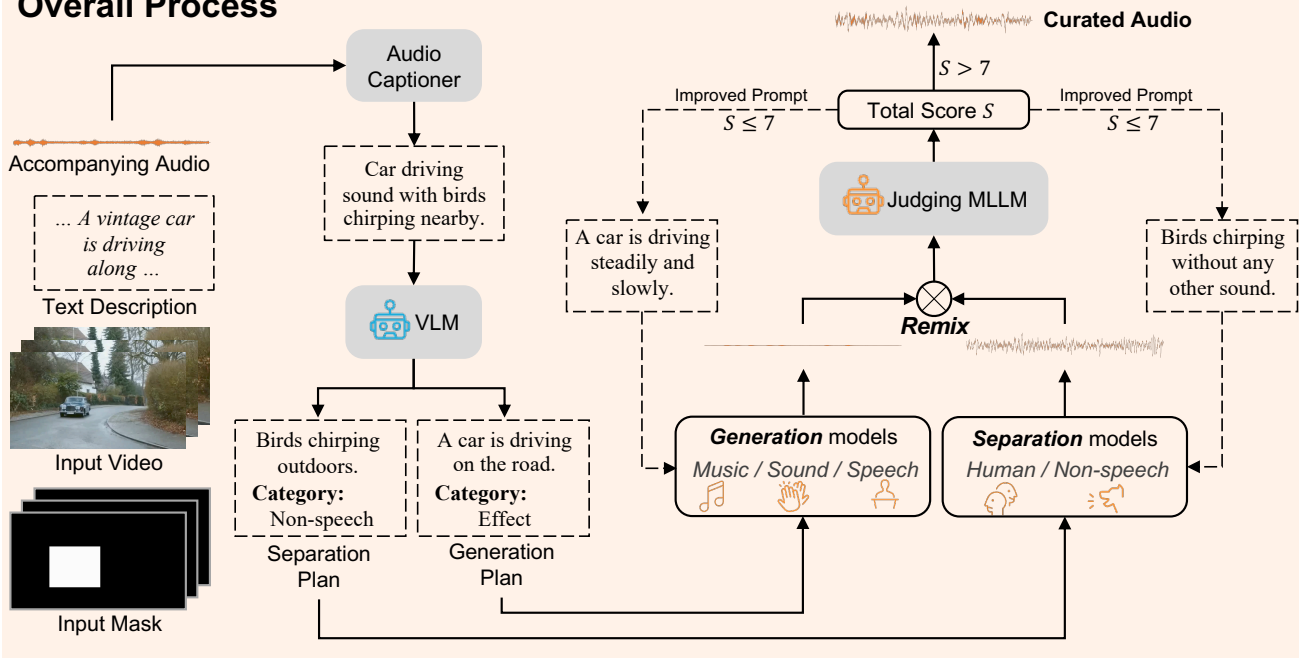


Figure 9. A representative example of the self-feedback audio agent.

References

- [1] ElevenLabs. <https://elevenlabs.io>. 5
- [2] ElevenLabs Scribe. <https://elevenlabs.io/blog/meet-scribe>. 3
- [3] Sora 2. <https://openai.com/index/sora-2/>. 2
- [4] Veo 3. <https://aistudio.google.com/models/veo-3>. 2
- [5] Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibao Song, Kai Dang, Peng Wang, Shijie Wang, Jun Tang, et al. Qwen2.5-vl technical report. *arXiv:2502.13923*, 2025. 3, 5
- [6] Max Bain, Arsha Nagrani, Andrew Brown, and Andrew Zisserman. Condensed movies: Story based retrieval with contextual embeddings. In *ACCV*, 2020. 3
- [7] Bibars Al Haj Bara, Nadezhda N Pokrovskaya, Marianna Yu Ababkova, Irina A Brusakova, and Anastasia A Korban. Artificial intelligence for advertising and media: machine learning and neural networks. In *EIconRus*, 2022. 2
- [8] Stanislav Beliaev and Boris Ginsburg. TalkNet 2: Non-autoregressive depth-wise separable convolutional model for speech synthesis with explicit pitch and duration prediction. *arXiv:2104.08189*, 2021. 3
- [9] Yuxuan Bian, Zhaoyang Zhang, Xuan Ju, Mingdeng Cao, Liangbin Xie, Ying Shan, and Qiang Xu. VideoPainter: Any-length video inpainting and editing with plug-and-play context control. In *SIGGRAPH*, 2025. 2
- [10] Andreas Blattmann, Tim Dockhorn, Sumith Kulal, Daniel Mendelevitch, Maciej Kilian, Dominik Lorenz, Yam Levi, Zion English, Vikram Voleti, Adam Letts, et al. Stable video diffusion: Scaling latent video diffusion models to large datasets. *arXiv:2311.15127*, 2023. 3
- [11] Jake Bruce, Michael D Dennis, Ashley Edwards, Jack Parker-Holder, Yuge Shi, Edward Hughes, Matthew Lai, Aditi Mavalankar, Richie Steigerwald, Chris Apps, et al. Genie: Generative interactive environments. In *Forty-first International Conference on Machine Learning*, 2024. 2
- [12] Brandon Castellano. PySceneDetect, 2025. <https://github.com/Breakthrough/PySceneDetect>. 3
- [13] Honglie Chen, Weidi Xie, Andrea Vedaldi, and Andrew Zisserman. VGGSound: A large-scale audio-visual dataset. In *ICASSP*, 2020. 3
- [14] Ming Chen, Liyuan Cui, Wenyuan Zhang, Haoxian Zhang, Yan Zhou, Xiaohan Li, Songlin Tang, Jiwen Liu, Borui Liao, Hejia Chen, et al. MIDAS: Multimodal interactive digital-human synthesis via real-time autoregressive video generation. *arXiv:2508.19320*, 2025. 3
- [15] Yiyang Chen, Xuanhua He, Xiujun Ma, and Yue Ma. ContextFlow: Training-free video object editing via adaptive context enrichment. *arXiv:2509.17818*, 2025. 2
- [16] Yuqing Chen, Junjie Wang, Lin Liu, Ruihang Chu, Xiaopeng Zhang, Qi Tian, and Yujiu Yang. O-DisCo-Edit: Object distortion control for unified realistic video editing. *arXiv:2509.01596*, 2025. 2
- [17] Xize Cheng, Siqi Zheng, Zehan Wang, Minghui Fang, Ziang Zhang, Rongjie Huang, Ziyang Ma, Shengpeng Ji, Jialong Zuo, Tao Jin, et al. OmniSep: Unified omni-modality sound separation with query-mixup. *arXiv:2410.21269*, 2024. 5
- [18] Joon Son Chung and Andrew Zisserman. Out of time: automated lip sync in the wild. In *ACCV Workshops*, 2017. 6, 11
- [19] Patrick Esser, Robin Rombach, and Björn Ommer. Taming transformers for high-resolution image synthesis. In *CVPR*, 2021. 3
- [20] Mark Everingham, Luc Van Gool, Christopher KI Williams, John Winn, and Andrew Zisserman. The pascal visual object classes (voc) challenge. *IJCV*, 2010. 9
- [21] Youquan Fu, Ruiyang Si, Hongfa Wang, Dongzhan Zhou, Jiacheng Sun, Ping Luo, Di Hu, Hongyuan Zhang, and Xuelong Li. Object-AVEdit: An object-level audio-visual editing model. *arXiv:2510.00050*, 2025. 2, 3
- [22] Songwei Ge, Thomas Hayes, Harry Yang, Xi Yin, Guan Pang, David Jacobs, Jia-Bin Huang, and Devi Parikh. Long video generation with time-agnostic vqgan and time-sensitive transformer. In *ECCV*, 2022. 3
- [23] Jort F Gemmeke, Daniel PW Ellis, Dylan Freedman, Aren Jansen, Wade Lawrence, R Channing Moore, Manoj Plakal, and Marvin Ritter. AudioSet: An ontology and human-labeled dataset for audio events. In *ICASSP*, 2017. 3
- [24] Ridouane Ghermi, Xi Wang, Vicky Kalogeiton, and Ivan Laptev. Short film dataset (SFD): A benchmark for story-level video understanding. *CoRR*, 2024. 3
- [25] 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. 6, 11
- [26] Yuwei Guo, Ceyuan Yang, Anyi Rao, Zhengyang Liang, Yaohui Wang, Yu Qiao, Maneesh Agrawala, Dahua Lin, and Bo Dai. AnimateDiff: Animate your personalized text-to-image diffusion models without specific tuning. *arXiv:2307.04725*, 2023. 2
- [27] Romain Hennequin, Anis Khelif, Felix Voituret, and Manuel Moussallam. Spleeter: a fast and efficient music source separation tool with pre-trained models. *Journal of Open Source Software*, 2020. 5
- [28] Dasaem Jeong, Seunghoon Doh, and Taegyun Kwon. Träumerai: Dreaming music with stylegan. *arXiv:2102.04680*, 2021. 3
- [29] Zeyinzi Jiang, Zhen Han, Chaojie Mao, Jingfeng Zhang, Yulin Pan, and Yu Liu. VACE: All-in-one video creation and editing. *arXiv:2503.07598*, 2025. 2, 6, 7, 11
- [30] Ozgur Kara, Bariscan Kurtkaya, Hidir Yesiltepe, James M Rehg, and Pinar Yanardag. RAVE: Randomized noise shuffling for fast and consistent video editing with diffusion models. In *CVPR*, 2024. 2
- [31] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *CVPR*, 2019. 3
- [32] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv:1412.6980*, 2014. 6
- [33] Weijie Kong, Qi Tian, Zijian Zhang, Rox Min, Zuozhuo Dai, Jin Zhou, Jiangfeng Xiong, Xin Li, Bo Wu, Jianwei Zhang, et al. HunyuanVideo: A systematic framework for large video generative models. *arXiv:2412.03603*, 2024. 2

- [34] Black Forest Labs, Stephen Batifol, Andreas Blattmann, Frederic Boesel, Saksham Consul, Cyril Diagne, Tim Dockhorn, Jack English, Zion English, Patrick Esser, et al. FLUX. 1 Kontext: Flow matching for in-context image generation and editing in latent space. *arXiv:2506.15742*, 2025. 2
- [35] Cheng-Che Lee, Wan-Yi Lin, Yen-Ting Shih, Pei-Yi Kuo, and Li Su. Crossing you in style: Cross-modal style transfer from music to visual arts. In *ACM MM*, 2020. 2
- [36] 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. 3
- [37] Seung Hyun Lee, Sieun Kim, Innfarn Yoo, Feng Yang, Donghyeon Cho, Youngseo Kim, Huiwen Chang, Jinkyu Kim, and Sangpil Kim. Soundini: Sound-guided diffusion for natural video editing. *arXiv:2304.06818*, 2023. 3
- [38] Taegyeong Lee, Jeonghun Kang, Hyeonyu Kim, and Taehwan Kim. Generating realistic images from in-the-wild sounds. In *ICCV*, 2023. 2
- [39] Xirui Li, Chao Ma, Xiaokang Yang, and Ming-Hsuan Yang. VidToMe: Video token merging for zero-shot video editing. In *CVPR*, 2024. 2
- [40] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *ICCV*, 2017. 5
- [41] Yan-Bo Lin, Kevin Lin, Zhengyuan Yang, Linjie Li, Jianfeng Wang, Chung-Ching Lin, Xiaoferi Wang, Gedas Bertasius, and Lijuan Wang. Zero-shot audio-visual editing via cross-modal delta denoising. *arXiv:2503.20782*, 2025. 2, 3, 6, 7, 11
- [42] Yaron Lipman, Ricky TQ Chen, Heli Ben-Hamu, Maximilian Nickel, and Matt Le. Flow matching for generative modeling. *arXiv:2210.02747*, 2022. 3
- [43] Chetwin Low, Weimin Wang, and Calder Katyal. Ovi: Twin backbone cross-modal fusion for audio-video generation. *arXiv:2510.01284*, 2025. 3, 6, 7, 11
- [44] Hila Manor and Tomer Michaeli. Zero-shot unsupervised and text-based audio editing using ddpn inversion. In *ICML*, 2024. 8
- [45] Tae-Hyun Oh, Tali Dekel, Changil Kim, Inbar Mosseri, William T Freeman, Michael Rubinstein, and Wojciech Matusik. Speech2Face: Learning the face behind a voice. In *CVPR*, 2019. 2
- [46] William Peebles and Saining Xie. Scalable diffusion models with transformers. In *ICCV*, 2023. 9
- [47] 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
- [48] 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. 6, 11
- [49] Tianhe Ren, Shilong Liu, Ailing Zeng, Jing Lin, Kunchang Li, He Cao, Jiayu Chen, Xinyu Huang, Yukang Chen, Feng Yan, et al. Grounded SAM: Assembling open-world models for diverse visual tasks. *arXiv:2401.14159*, 2024. 3
- [50] James Robert. Pydub, 2025. <https://github.com/jiaaro/pydub>. 5
- [51] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *CVPR*, 2022. 2
- [52] Tim Salimans, Ian Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, and Xi Chen. Improved techniques for training gans. In *NIPS*, 2016. 6, 11
- [53] Christoph Schuhmann. Improved aesthetic predictor, 2025. <https://github.com/christophschuhmann/improved-aesthetic-predictor>. 3
- [54] Sizhe Shan, Qiulin Li, Yutao Cui, Miles Yang, Yuehai Wang, Qun Yang, Jin Zhou, and Zhao Zhong. HunyuanVideo-Foley: Multimodal diffusion with representation alignment for high-fidelity foley audio generation. *arXiv:2508.16930*, 2025. 6, 7, 11
- [55] Kaixin Shen, Ruijie Quan, Linchao Zhu, Jun Xiao, and Yi Yang. AudioScenic: Audio-driven video scene editing. *arXiv:2404.16581*, 2024. 3
- [56] Zachary Teed and Jia Deng. Raft: Recurrent all-pairs field transforms for optical flow. In *ECCV*, 2020. 3
- [57] Andros Tjandra, Yi-Chiao Wu, Baishan Guo, John Hoffman, Brian Ellis, Apoorv Vyas, Bowen Shi, Sanyuan Chen, Matt Le, Nick Zacharov, et al. Meta audiobox aesthetics: Unified automatic quality assessment for speech, music, and sound. *arXiv:2502.05139*, 2025. 3
- [58] Stefan Uhlich, Giorgio Fabbro, Masato Hirano, Shusuke Takahashi, Gordon Wichern, Jonathan Le Roux, Dipam Chakraborty, Sharada Mohanty, Kai Li, Yi Luo, Jianwei Yu, Rongzhi Gu, Roman A. Solovyev, Alexander L. Stempkovskiy, Tatiana Habruseva, Mikhail Sukhovei, and Yuki Mitsufuji. The sound demixing challenge 2023 - cinematic demixing track. *Trans. Int. Soc. Music. Inf. Retr.*, 2024. 5
- [59] Thomas Unterthiner, Sjoerd Van Steenkiste, Karol Kurach, Raphael Marinier, Marcin Michalski, and Sylvain Gelly. Towards accurate generative models of video: A new metric & challenges. *arXiv:1812.01717*, 2018. 6, 11
- [60] Team Wan, Ang Wang, Baole Ai, Bin Wen, Chaojie Mao, Chen-Wei Xie, Di Chen, Feiwei Yu, Haiming Zhao, Jianxiao Yang, et al. Wan: Open and advanced large-scale video generative models. *arXiv:2503.20314*, 2025. 2, 4, 6, 8
- [61] Duomin Wang, Wei Zuo, Aojie Li, Ling-Hao Chen, Xinyao Liao, Deyu Zhou, Zixin Yin, Xili Dai, Daxin Jiang, and Gang Yu. UniVerse-1: Unified audio-video generation via stitching of experts. *arXiv:2509.06155*, 2025. 3
- [62] Jiapeng Wang, Chengyu Wang, Kunzhe Huang, Jun Huang, and Lianwen Jin. VideoCLIP-XL: Advancing long description understanding for video clip models. *arXiv:2410.00741*, 2024. 6, 11
- [63] Zhenzhi Wang, Jiaqi Yang, Jianwen Jiang, Chao Liang, Gaojie Lin, Zerong Zheng, Ceyuan Yang, and Dahua Lin. InterActHuman: Multi-concept human animation with layout-aligned audio conditions. *arXiv:2506.09984*, 2025. 3
- [64] Shuchen Weng, Haojie Zheng, Zheng Chang, Si Li, Boxin Shi, and Xinlong Wang. Audio-sync video generation with multi-stream temporal control. In *NIPS*, 2025. 2

- [65] Shuchen Weng, Haojie Zheng, Peixuan Zhang, Yuchen Hong, Han Jiang, Si Li, and Boxin Shi. VIREs: Video instance repainting via sketch and text guided generation. In *CVPR*, 2025. 2
- [66] Jay Zhangjie Wu, Yixiao Ge, Xintao Wang, Stan Weixian Lei, Yuchao Gu, Yufei Shi, Wynne Hsu, Ying Shan, Xiaohu Qie, and Mike Zheng Shou. Tune-A-Video: One-shot tuning of image diffusion models for text-to-video generation. In *ICCV*, 2023. 2
- [67] Weijia Wu, Mingyu Liu, Zeyu Zhu, Xi Xia, Haoen Feng, Wen Wang, Kevin Qinghong Lin, Chunhua Shen, and Mike Zheng Shou. MovieBench: A hierarchical movie level dataset for long video generation. *arXiv:2411.15262*, 2024. 3
- [68] Jin Xu, Zhifang Guo, Hangrui Hu, Yunfei Chu, Xiong Wang, Jinzheng He, Yuxuan Wang, Xian Shi, Ting He, Xinfu Zhu, et al. Qwen3-omni technical report. *arXiv:2509.17765*, 2025. 3, 5
- [69] An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, et al. Qwen3 technical report. *arXiv:2505.09388*, 2025. 3
- [70] Shuzhou Yang, Xiaoyu Li, Xiaodong Cun, Guangzhi Wang, Lingen Li, Ying Shan, and Jian Zhang. GenComposer: generative video compositing with diffusion transformer. *arXiv:2509.02460*, 2025. 2
- [71] Zhuoyi Yang, Jiayan Teng, Wendi Zheng, Ming Ding, Shiyu Huang, Jiazheng Xu, Yuanming Yang, Wenyi Hong, Xiaohan Zhang, Guanyu Feng, et al. CogVideox: Text-to-video diffusion models with an expert transformer. In *ICLR*, 2025. 2
- [72] Guy Yariv, Itai Gat, Lior Wolf, Yossi Adi, and Idan Schwartz. AudioToken: Adaptation of text-conditioned diffusion models for audio-to-image generation. *arXiv:2305.13050*, 2023. 2
- [73] Peixuan Zhang, Zijian Jia, Kaiqi Liu, Shuchen Weng, Si Li, and Boxin Shi. Stage: Storyboard-anchored generation for cinematic multi-shot narrative. *arXiv preprint arXiv:2512.12372*, 2025. 2
- [74] Ruihan Zhang, Borou Yu, Jiajian Min, Yetong Xin, Zheng Wei, Juncheng Nemo Shi, Mingzhen Huang, Xianghao Kong, Nix Liu Xin, Shanshan Jiang, Praagya Bahuguna, Mark Chan, Khushi Hora, Lijian Yang, Yongqi Liang, Runhe Bian, Yunlei Liu, Isabela Campillo Valencia, Patricia Morales Tredinick, Ilia Kozlov, Sijia Jiang, Peiwen Huang, Na Chen, Xuanxuan Liu, and Anyi Rao. Generative AI for film creation: A survey of recent advances. In *CVPR Workshops*, 2025. 2