

SALSA-V: SHORTCUT-AUGMENTED LONG-FORM SYNCHRONIZED AUDIO FROM VIDEOS

Anonymous authors

Paper under double-blind review

ABSTRACT

We propose SALSA-V, a multimodal video-to-audio generation model capable of synthesizing highly synchronized, high-fidelity long-form audio from silent video content. Our approach introduces a masked diffusion objective, enabling audio-conditioned generation and the seamless synthesis of audio sequences of unconstrained length. Additionally, by integrating a shortcut loss into our training process, we achieve rapid generation of high-quality audio samples in as few as eight sampling steps, paving the way for near-real-time applications without requiring dedicated fine-tuning or retraining. We demonstrate that SALSA-V significantly outperforms existing state-of-the-art methods in both audiovisual alignment and synchronization with video content in quantitative evaluation and a human listening study. Furthermore, our use of random masking during training enables our model to match spectral characteristics of reference audio samples, broadening its applicability to professional audio synthesis tasks such as Foley generation and sound design.

1 INTRODUCTION

Video-to-audio (V2A) generation, sometimes referred to as “computational Foley”, aims to produce realistic sounds for the visual events occurring in a silent video clip. Unlike background music or speech synthesis, Foley focuses on diegetic sounds, which are sounds implied by the current on-screen content (e.g., the sound of rain and thunder when a storm is shown, or a dog’s bark echoing in a room). Achieving realism requires semantic (the model must recognize what is happening so it can select the right acoustic event) as well as temporal alignment (it must identify when that event occurs). Especially temporal alignment is crucial, as humans are sensitive to as few as tens of milliseconds of asynchrony (Keetels & Vroomen, 2005).

Early generative machine learning models for video-to-audio were trained from scratch on modestly-sized audio-visual corpora and struggled to cover the acoustic diversity of in-the-wild video. Recent work has addressed this issue by borrowing scale from adjacent modalities. One line of research adapts large text-to-audio diffusion models with lightweight video-conditioned control modules (Zhang et al., 2024), while another trains end-to-end multimodal generators using cross-modal attention, latent diffusion, or rectified-flow objectives (Cheng et al., 2025; Wang et al., 2025).

While recent approaches have resulted in significant improvements being made in this domain, existing models still face issues with precise on-screen synchronization and controllability of the generated sounds. Additionally, most existing works focus solely on generating audio for short, isolated snippets, often limited to a maximum length of 10 seconds due to either fundamental architectural restrictions or memory constraints. Attempting to utilize these approaches for longer sequences often significantly deteriorates performance, or can lead to stitching artifacts if shorter sequences are processed independently and later re-joined. In addition, users are offered limited control over the generated sounds, with most models opting for text conditioning for this purpose. Apart from limiting sound characteristics to those describable through text, this type of conditioning by itself does not natively support any editability of the produced audio. Furthermore, generation speed is usually treated as an afterthought by most approaches, requiring (in the case of diffusion models) many sampling steps to achieve high-quality results. This restricts the usability of these models to slow-feedback workflows and offline-only uses.

054 Our work aims to address these aspects, achieving the ability to handle audio conditioning and
055 extended-length generations without large reductions in quality. Furthermore, our model enables
056 near-real-time high-quality audio generations by using a dedicated training paradigm that enables
057 few-step sampling natively (i.e., without special fine-tuning or re-training). We extensively evaluate
058 our model against the existing state-of-the-art, including conducting a human preference study.

059 Our contributions can be summarized as follows:
060

- 061 • We introduce SALSA-V, a shortcut-augmented latent flow matching model for video-
062 to-audio, enabling high-fidelity few-step sampling without distillation or additional fine-
063 tuning.
- 064 • Support for audio conditioning and outpainting through a masked training objective, en-
065 abling the generation of samples based on reference audio sections and long-form genera-
066 tions using iterative extension.
- 067 • A contrastively-trained audio-visual alignment model using a backbone with large-scale
068 pretraining, which yields high-resolution synchronization features, and results in state-of-
069 the-art temporal synchronization and strong human MOS.

071 **Samples are available at** <https://anonymous.4open.science/w/salsav>.
072

073 2 RELATED WORK

074 2.1 VIDEO-TO-AUDIO GENERATION

075
076
077
078 Early efforts approached video-to-audio generation as a self-supervised prediction problem. Visu-
079 ally Indicated Sounds (Owens et al., 2016) synthesized impact noises from silent video using an
080 RNN trained to regress sound features, demonstrating that visual cues can encode material and ac-
081 tion properties. With the advent of more capable generative models, the focus has recently shifted to
082 high-fidelity, temporally aligned synthesis. Diff-Foley (Luo et al., 2023) was the first work to focus
083 on video-to-audio generation with diffusion models, demonstrating the feasibility of this approach.
084 Frieren (Wang et al., 2025) replaced the standard diffusion training process with a rectified flow
085 matching approach (Liu et al., 2022). By using model distillation (reflowing the original model),
086 they can reduce the number of sampling steps necessary, but do so by re-training and distilling their
087 model in dedicated fine-tuning phases. Complementarily, MMAudio (Cheng et al., 2025) introduced
088 multimodal joint training on large text–audio corpora and a frame-level synchronization module,
089 achieving state-of-the-art results, both in terms of general audio quality, as well as synchronization
090 to onscreen events. An autoregressive approach was recently explored in V-AURA (Viertola et al.,
091 2024), which natively enables long-form generation capabilities. However, V-AURA’s approach is
092 somewhat limited in that it does not allow for text conditioning and is restricted to a comparatively
093 short context window of 2.56 seconds.

094 2.2 CROSS-MODAL AUDIOVISUAL LEARNING

095
096 Existing multimodal representation learning approaches focus mostly on learning similarities be-
097 tween visual features and audio in a “global manner”, meaning that the temporal information inher-
098 ent in both modalities is either ignored completely or treated coarsely (e.g. by subsampling videos
099 at a low frame rate). These approaches include ImageBind (Girdhar et al., 2023), which learns a
100 latent space covering various different modalities, and VATT (Akbari et al., 2021). More recently,
101 foundational models, such as InternVideo2 (Wang et al., 2024), have also been used for this purpose.

102 These approaches are highly capable of learning broad semantic similarities, but are not concerned
103 with high-resolution temporal alignment between individual video frames and audio. Another line
104 of work does place their focus on this topic, including Synchformer (Iashin et al., 2024), which
105 has a contrastive pretraining stage (*Segment-AVCLIP*) trained on short video-audio snippets. These
106 embeddings are then used to predict a global synchronization score between a given audio and video
107 segment. The embeddings learned by Synchformer’s first stage have been utilized by previous V2A
models to provide synchronization signals for downstream generative models (Cheng et al., 2025).

108
109
110
111
112
113
114
115
116
117
118
119
120
121
122
123
124
125
126
127
128
129
130
131
132
133
134
135
136
137
138
139
140
141
142
143
144
145
146
147
148
149
150
151
152
153
154
155
156
157
158
159
160
161

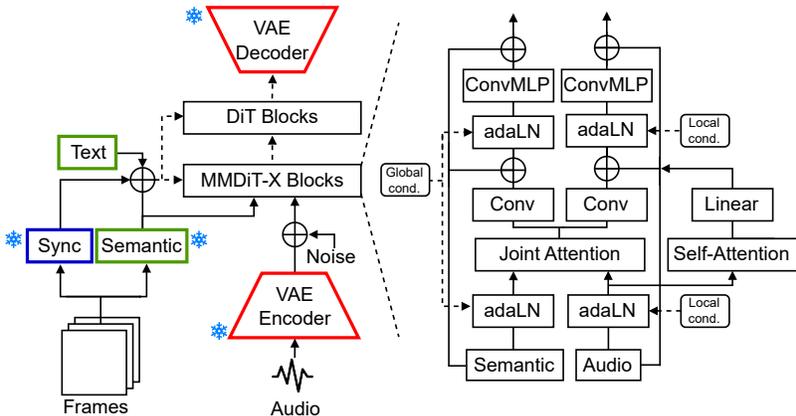


Figure 1: Architectural diagram of SALSA-V. We utilize a mixture of modified MMDiT-X blocks operating jointly over the combined sequence of audio and semantic features, as well as single-stream DiT blocks, which operate on the latent representation of a VAE (red). The overall global conditioning signal consists of the text embedding and pooled semantic features (green), which are extracted from a pretrained SigLIP 2 encoder, synchronization features (blue) obtained from a dedicated contrastive model, and the time-step embedding. The local conditioning tensor equals the global conditioning summed to the sequence-aligned synchronization features.

2.3 EFFICIENT GENERATIVE MODELS

Due to the high number of sampling steps required for high-quality generative results when using diffusion models (usually dozens in most domains), most generative audio models are not suited for real-time applications. Consequently, there has recently been some interest in improving this aspect by reducing the number of sampling steps necessary to achieve higher fidelity. In diffusion and flow-matching models, this is commonly done by employing additional fine-tuning stages, where the denoising trajectory of the model is straightened (Liu et al., 2022). This enables few-step generations with significantly lower errors by letting the model take larger straight-line steps (along the newly learned path) before having to recalculate the velocity field. While multiple of these approaches have been successfully employed in the image generation domain, the only video-to-audio generation model which employs such a strategy is Friren (Wang et al., 2025), which uses a rectified flow training objective (Liu et al., 2022).

2.4 LONG-FORM AUDIO GENERATION

While standard V2A models are not in principle restricted to inputs of a few seconds, in practice, memory (growing with the square of the sequence length in transformer models) and data constraints lead to this restriction being adopted (often resulting in maximum model inputs of 10 seconds). This leads to models that only perform well on this limited set of durations, with many experiencing a sharp drop-off in performance when extended even modestly beyond that point. LoVA (Cheng et al., 2024) first placed a focus on long-form generation. The authors attempt to address this issue by using a standard diffusion transformer (Peebles & Xie, 2023), but training it on dedicated long-form data (up to 30 seconds). Although this leads to performance improvements over training solely with short, fixed-size snippets, it does not address the underlying issue of unrestricted memory growth, and the generated audio snippets are still limited by the maximum length observed during training.

3 METHODOLOGY

A high-level schematic of our model and training process can be seen in Figure 1. We use a latent flow matching approach, utilizing an audio encoder based on Stable-Audio VAE (Evans et al., 2024), which was fine-tuned to increase its frame rate to 43 Hz in order to achieve higher audio quality (see Appendix A.5). We denote this input latent sequence as x_1 . Its shape is $(t_a, 64)$, where t_a is

the encoded sequence length. x_1 then undergoes the flow-matching noise-injection process on all sequence items, except for a randomly sampled contiguous subsequence (see Section 3.2). x_t is then processed through a diffusion transformer model (Peebles & Xie, 2023), which employs a similar architecture to recent multi-modal diffusion models (Cheng et al., 2025; Esser et al., 2024). In particular, we utilize a mixture of transformer blocks which operate over audio and coarse (semantic) video sequences. We modify the MMDiT-X (Stability AI, 2024) layer architecture for our use-case, which, in addition to a joint attention operation applied to the sequence-wise concatenation of the query, key, and value projections of its inputs, uses an optional additional self-attention layer applied solely to the audio stream. We only apply this additional self-attention operation over the first n blocks of the model. The output of these layers is then processed by additional standard single-sequence DiT blocks (Peebles & Xie, 2023). We utilize rotary positional embeddings (Su et al., 2023) in all transformer blocks.

Representations. Our model relies on three sources of conditioning information: low-resolution semantic visual features v_s , high-resolution visual synchronization features v_{syn} , and global text features t_c . v_s is utilized both as a global feature (by pooling along the sequence axis), and in its original shape as an input to the multi-modal transformer blocks. v_{syn} is used in sequence-level adaLN (Xu et al., 2019) conditioning layers after undergoing a learned upsampling operation to align its length with that of the audio sequence. For semantic video features, we use an off-the-shelf frozen SigLIP 2 (Tschannen et al., 2025) encoder, which is trained by contrastively aligning visual embeddings produced by a ViT (Dosovitskiy et al., 2021) with the corresponding text embeddings. We extract visual features frame-wise, using a lower frame rate of 8 FPS, with the intention of providing rich semantic information to the model, without yet focusing on precise synchronization. This process yields a sequence of shape $(t_{v_s}, 1152)$ where t_{v_s} is the number of original video frames at 8 FPS. Text features are calculated using the corresponding text encoder, which yields a single 1152-dimensional vector per sequence. We train a dedicated model to learn synchronization features, which produces embeddings of shape $(t_{v_{\text{syn}}}, 768)$ which are sampled at 24 FPS.

3.1 SYNCHRONIZATION FEATURES & CONTRASTIVE PRE-TRAINING

Recent V2A models have employed dedicated contrastively-learned synchronization features for conditioning (Cheng et al., 2025), mostly based on the Synchformer model (Iashin et al., 2024). While these models are generally able to capture audio-visual time alignment features well, they can suffer from particular failure modes where some actions within the video are not recognized as such, which leads to no corresponding sound being generated. We hypothesize that this is due to the limited amount of data seen during pre-training, which shifts data requirements to the diffusion training stage. For this reason, we train our own contrastive synchronization model based on a significantly larger pre-trained vision backbone (VideoPrism) (Zhao et al., 2025). VideoPrism itself is based on the ViViT (Arnab et al., 2021) architecture, which employs factorized attention for the temporal and spatial axes.

We use the audio-spectrogram transformer model (AST) (Gong et al., 2021) as the audio encoder, which operates on mel-spectrogram features extracted from audio snippets with a sampling rate of 16 kHz. Following previous literature (Iashin et al., 2024), we use 128-dimensional mel features extracted from 25 ms audio segments with a hop size of 10 ms. We initialize this model with pretrained AST weights (Gong et al., 2021). The entire audio backbone is unfrozen and fine-tuned during training, while we add additional trainable layers and a MAP head (Zhai et al., 2022) on top of the otherwise frozen VideoPrism encoder. During training, both encoders receive snippets of length 0.667 seconds of their respective modality (see Figure 2). The output embedding sequences are then averaged over their respective temporal axes to provide a single token for the input snippet for both modalities, \tilde{x}_v and \tilde{x}_a . We then align these two vectors using a SigLIP loss objective (Zhai et al., 2023):

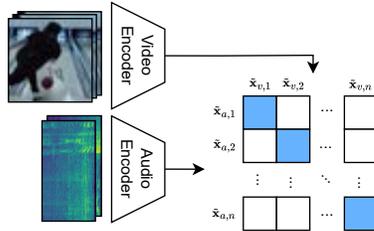


Figure 2: Contrastive alignment process. Groups of subsequent video frames are aligned with the corresponding audio snippets. Only patches corresponding to the same time range in the same batch are counted as positives.

216
217
218
219
220
221
222
223
224
225
226
227
228
229
230
231
232
233
234
235
236
237
238
239
240
241
242
243
244
245
246
247
248
249
250
251
252
253
254
255
256
257
258
259
260
261
262
263
264
265
266
267
268
269

$$\mathcal{L}_{\text{SigLIP}} = -\frac{1}{|\mathcal{B}|} \sum_{i=1}^{|\mathcal{B}|} \sum_{j=1}^{|\mathcal{B}|} \log \frac{1}{1 + e^{z_{ij}(-t\bar{\mathbf{x}}_{v,i} \cdot \bar{\mathbf{x}}_{a,j} + b)}} \quad (1)$$

where z_{ij} is the label for a given video/audio pair, equal to 1 if they are paired (that is, correspond to the same time range from the same video) and -1 otherwise, \mathcal{B} is the batch size, and $\bar{\mathbf{x}}_{v,i}$ and $\bar{\mathbf{x}}_{a,j}$ are the encoder outputs for video & audio at batch positions i and j . t and b are learned temperature and bias parameters, respectively.

An important aspect of this contrastive training process is the choice of batch size. In most contrastive pretraining tasks, higher batch sizes are usually preferred due to the model being exposed to more hard negatives against which the loss is evaluated. When training for precise synchronization, however, a larger batch size will result in an effective decrease of the proportion of hard negatives in the batch (snippets from the same video are generally more difficult to differentiate than snippets from other videos due to their high semantic similarity), transforming the task from that of learning synchronization and precise time alignment, to that of a more general semantic alignment process. We find that large batch sizes are worse for the learning features relevant to synchronization, which reflects similar findings made in Synchformer (Iashin et al., 2024), where it is noted that increasing the batch size beyond 28 (two videos yielding 14 snippets each) did not improve performance.

3.2 MASKED FLOW MATCHING

In order to enable audio-conditioned generation, we utilize a modified flow matching objective modified for in- and outpainting, using a similar approach as Chen et al. (2025). We do so with the intent of providing the model with the ability to seamlessly extend a provided ground-truth audio, enabling chunked long-form generation.

During training, we randomly mask a subset of the sequences in each batch with a given probability, with the mask spanning a contiguous range from 5 ms to 2.5 s. The masked subsequence contains the clean audio, and is thus not subject to the noising process resulting from flow matching training. Accordingly, the corresponding token positions are not included in the loss calculation, with the velocity prediction only being taken into account for noised tokens (ground-truth reference latents therefore do not contribute to the gradient). In order for the model to differentiate clean masked audio from the audio to be generated, we add two different learned tokens corresponding to masked and unmasked tokens to the relevant positions, following a technique employed in VampNet (Garcia et al., 2023). These masks are then also utilized during inference to mark the reference audio section (for details, see Appendix A.1). Beyond facilitating the conditioning on reference audio tracks, enabling users to generate audio that exhibits the spectral characteristics of a desired sample, this approach also enables in- and outpainting. The latter aspect is particularly relevant for long-form video-to-audio generation, which remains an issue in existing approaches. The vast majority of V2A models focus on the 8-12 second range, and experience a major reduction in performance when evaluated on longer sequences. Our model aims to circumvent this issue by continually extending samples generated on small snippets where synchronization remains high.

3.3 SHORTCUT LOSS

We adopt a shortcut formulation (Frans et al., 2025) for our generative model to enable fast sampling without needing to retrain specifically for this purpose. To our knowledge, this is the first adoption of this formulation in the video-to-audio domain, as well as the larger field of general generative audio.

The input to the standard generative model used in flow matching is augmented with the desired step size d . This allows the model to generalize its velocity prediction to larger step sizes, essentially anticipating future curvature of the denoising path. Shortcut models are trained by enforcing a self-consistency property, in that one shortcut step must equal two shortcut steps of half that size (Frans et al., 2025). Training with the consistency property is done by setting a fraction (usually 25%) of the batch to self-consistency targets instead of the standard flow-matching denoising targets. Specifically, the training objective can be written as follows:

$$\mathcal{L}^S(\theta) = E_{x_0, x_1, (t, d) \sim p(t, d)} [\underbrace{\|v_\theta(x_t, t, 0) - (x_1 - x_0)\|^2}_{\text{Flow-Matching}} + \underbrace{\|v_\theta(x_t, t, 2d) - v_{\text{target}}\|^2}_{\text{Self-Consistency}}], \quad (2)$$

where x_0 and x_1 represent samples from the noise and data distributions respectively, t is the time step, d is the sampled step size, and x_t is the noisy latent (corresponding to t). v_θ represents the velocity prediction made by the model, and v_{target} is the target velocity for the self-consistency loss.

4 EXPERIMENTS

We conduct experiments to evaluate our model’s capabilities along multiple aspects: **1)** General capabilities under the standard V2A setting, i.e., text-conditioned audio generation for input clips from a variety of domains. This includes audio quality and synchronization. **2)** Audio-conditioned generation, in order to evaluate our model’s ability to utilize a provided audio sample faithfully and realistically in a video clip. **3)** The capacity for few-step generation of high-quality results. **4)** Long-form capabilities, including the capacity to retain contextual aspects over longer horizons.

Datasets. In terms of audio-visual data, we use multiple datasets for training, including VGGSound (Chen et al., 2020) (~ 550 h), and, to expand coverage, Moments-in-Time (Monfort et al., 2019) and a collection of in-the-wild YouTube videos obtained from Panda70M (Chen et al., 2024). While VGGSound provides good audio-visual alignment, we filter the rest of our collected data using our audio-visual synchronization model, where only those samples whose overall temporal and global semantic alignment exceed a relative threshold are kept. In total, this process results in ~ 900 hours of audio-visual training data. For text supervision, we use AudioSetCaps captions for VGGSound (Bai et al., 2024), as well as WavCaps (Mei et al., 2024) and a high-fidelity studio Foley set for audio-text data. The total amount of audio-text data is thus ~ 8.5 k hours. Samples are cropped to a maximum length of 15 seconds. Evaluations are conducted on a test set composed of variable-length videos of up to 30 seconds, which contains samples from a holdout set of in-the-wild videos, the VGGSound test set, and UnAV-100 (Geng et al., 2023).

Evaluation Metrics. We assess the quality of our model using multiple metrics to measure audio similarity in learned representational spaces (distribution matching), overall quality, semantic & temporal alignment, and human preference. In particular, we compute the Fréchet Distance (Kilgour et al., 2019) using VGGish (Hershey et al., 2017) as the embedding model (FAD_{VGG}), and the Kullback-Leibler divergence using PaSST (Koutini et al., 2022) and PANNs (Kong et al., 2020). General audio quality is evaluated using the Inception Score (IS) (Salimans et al., 2016) based on PANNs, which is a reference-free measure to assess the clarity & diversity of generated samples. Semantic alignment between video and audio is measured using ImageBind (Girdhar et al., 2023) following the approach used in MMAudio and V-AURA (Cheng et al., 2025; Viertola et al., 2024). Features are extracted from the audio track and subsampled video frames, after which their embedding’s average cosine similarity is calculated. Temporal alignment is evaluated using SynchFormer (Iashin et al., 2024) (DeSync), following the approach of MMAudio. This metric produces a misalignment value (in seconds) with a resolution of 0.2 seconds.

We also conduct a human listening study in which listeners are tasked with rating 10 videos (chosen for thematic diversity and to cover a range of sounds) from our test set along the axes of overall quality, temporal synchronization, and audio fidelity using a 5-point MOS scale. Audio tracks are shuffled blindly between trials (see Appendix A.3). The listening study was conducted with 18 qualified participants and included samples generated by our model, MMAudio (Cheng et al., 2025), and FoleyCrafter (Zhang et al., 2024).

4.1 RESULTS

Table 1 summarizes our model’s performance on our test set against existing state-of-the-art models. **It should be noted that this evaluation is conducted on mixed-length videos, as opposed to clips of fixed durations.** SALSA-V surpasses previous approaches particularly in terms of synchronization, which is reflected in our human evaluation and in the DeSync metric. We also reach competitive scores in terms of distribution matching metrics, demonstrating the best FAD_{VGG} and KL_{PANNs} scores, while AudioX (Tian et al., 2025) performs best in terms of KL_{PANNs} . With regards

Table 1: Comparison with prior work. Left: objective metrics. Right: human evaluation (higher is better). ‘‘Sync.’’ refers to synchronization (i.e., how well video frames and audio align). SALSA-V outperforms previous approaches on both objective (DeSync) and subjective human evaluation, highlighting its ability to better generate aligned audio.

| Method | Params | Objective metrics | | | | | | Human evaluation | | |
|--------------|--------|----------------------|-----------------------|-----------------------|--------------|--------------|--------------|------------------|-------------|-------------|
| | | FAD _{VGG} ↓ | KL _{PANNs} ↓ | KL _{PaSST} ↓ | IS↑ | IB↑ | DeSync↓ | Overall | Quality | Sync. |
| Frieren | 159M | 1.42 | 2.61 | 2.55 | 13.20 | 23.64 | 0.981 | - | - | - |
| V-AURA | 695M | 2.93 | 2.33 | 1.98 | 9.87 | 27.29 | 0.696 | - | - | - |
| LoVA | 1.06B | 1.79 | 2.14 | 2.06 | 16.86 | 27.95 | 1.205 | - | - | - |
| AudioX | 1.17B | 1.11 | 1.70 | 1.63 | 18.05 | 26.57 | 0.862 | - | - | - |
| FoleyCrafter | 1.22B | 2.43 | 2.21 | 2.15 | 16.39 | 26.37 | 1.319 | 2.78 | 2.46 | 2.37 |
| MMAudio | 1.03B | 1.12 | 1.77 | 1.72 | 18.13 | 32.89 | 0.521 | 3.29 | 3.16 | 2.92 |
| SALSA-V | 643M | 1.07 | 1.81 | 1.63 | 17.85 | 33.76 | 0.497 | 3.43 | 2.96 | 3.52 |

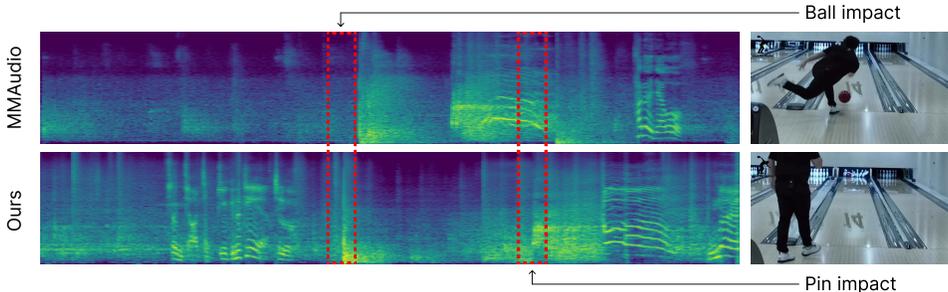


Figure 3: Qualitative example showing Mel-spectrogram of a generated audio clip with marked impacts, comparing SALSA-V with MMAudio. We observe that SALSA-V is able to better predict event alignment compared to previous methods.

to reference-free audio quality as measured by IS, and semantic alignment (IB) both our model and MMAudio achieve close scores. Human rating scores for overall quality are closer, with SALSA-V leading with a value of 3.43 versus MMAudio’s 3.29.

We visualize the audio generated by our model on a representative sample from our validation set in Figure 3, which displays the generated mel-spectrogram. The sample is chosen for the marked impact sounds created by the bowling ball hitting first the track and then the pins. We find that SALSA-V is able to accurately pinpoint said events with high precision while also capturing the general duration of a given impact sound (e.g., the pin hit results in a more sustained sound, while the ball hitting the track has a fast attack and shorter duration). Moreover, our model retains the ability to generate appropriate background sounds during periods of limited video activity, as illustrated by the sections of human speech generated at the start and end of the segment.

Table 2: Ablation study on model size and diffusion variant. We report metrics of differently-sized model variants trained with either a default flow-matching loss or a shortcut objective. Utilizing the shortcut training objective does not reduce performance compared to flow matching. All models are evaluated at 32 sampling steps.

| Variant | Params | FAD _{VGG} ↓ | KL _{PANNs} ↓ | KL _{PaSST} ↓ | IS↑ | IB↑ | DeSync↓ |
|---------------|--------|----------------------|-----------------------|-----------------------|--------------|--------------|--------------|
| Flow Matching | 347M | 1.09 | 1.87 | 1.71 | 18.01 | 31.21 | 0.504 |
| Flow Matching | 643M | 1.10 | 1.80 | 1.63 | 17.91 | 33.71 | 0.499 |
| Shortcut | 347M | 1.12 | 1.85 | 1.68 | 18.03 | 31.09 | 0.501 |
| Shortcut | 643M | 1.07 | 1.81 | 1.63 | 17.85 | 33.76 | 0.497 |

Table 3: Comparison of metrics of Frieren (reflow) and SALSA-V at 4 sampling steps. SALSA-V demonstrates better performance in the few-step setting across all metrics.

| Method | Steps | FAD _{VGG} ↓ | KL _{PANNs} ↓ | KL _{PaSST} ↓ | IS↑ | IB↑ | DeSync↓ |
|------------------|-------|----------------------|-----------------------|-----------------------|--------------|--------------|--------------|
| Frieren (reflow) | 4 | 2.17 | 2.54 | 2.52 | 9.46 | 22.10 | 0.917 |
| SALSA-V | 4 | 1.19 | 1.83 | 1.67 | 14.57 | 31.18 | 0.536 |

Few-Step Generation. In order to test the few-step generation capabilities of SALSA-V, we compare the samples generated in the small-sampling-step regime with those of MMAudio (Cheng et al., 2025). The results of this analysis are shown in Figure 5, which includes FAD as a measure of semantic similarity to reference audio, and DeSync to measure synchronization. While MMAudio’s performance decreases significantly across every metric as the number of sampling steps decreases, our model is able to stay competitive in the minimal sampling step regime, incurring almost no performance reduction at as little as 8 sampling steps in both metrics. It should be noted that our synchronization performance in particular is robust across all settings, as most impact sounds are already present after a few sampling steps, with the difference from results at a larger number of sampling steps being mainly composed of high-frequency textural details (see Figure 8). We also compare the performance of our shortcut model against an identical variant trained with a vanilla flow-matching loss. As can be seen in Table 2, we incur no performance reduction compared to the flow-matching variant at the full number of sampling steps (32 in this case), which indicates that the shortcut model training procedure does not compromise the model’s ability to produce high-fidelity samples when needed. This stands in contrast to other approaches commonly used to reduce the number of sampling steps of flow matching and diffusion models, which usually entail fundamentally altering the learned velocity field in a post-training stage (e.g. the *reflow* procedure for rectified flow models (Liu et al., 2022)), which can harm fidelity at a larger number of sampling steps compared to the original model version.

To evaluate our model’s ability in the few-step setting, we conduct an evaluation against Frieren (Wang et al., 2025) (using the *reflow* checkpoint, which is optimized for this task), with both models evaluated at 4 sampling steps. The results of this comparison can be seen in Table 3. SALSA-V displays better performance in all objective metrics in the few-step setting as well.

We assess the time taken by different models to generate a single 10-second sample in Table 4, evaluating dedicated few-step models in the full- and few-step sampling setting. SALSA-V exhibits competitive durations for high-quality samples, especially in the few-step regime. All models are executed on an NVIDIA A100 GPU at a batch size of 1.

Conditional & Long-Form Generation. Conditionally extended audio samples generally retain the spectral characteristics of their reference audio sequence, as qualitatively displayed in Figure 4, which uses a reference sample from our validation set. It can be seen that the extended sample faithfully transfers audio characteristics from the conditioning section, while still responding to unseen

Table 4: Generation efficiency of different models, measuring the time taken to generate a 10-second sample using a batch size of 1. SALSA-V is able to generate higher-quality results at 4 sampling steps, resulting in higher generation speeds than most competing models.

| Method | Params | Time (s) |
|--------------------|--------|----------|
| V-AURA | 695M | 52.77 |
| FoleyCrafter | 1.22B | 5.53 |
| Frieren (25 steps) | 159M | 0.52 |
| Frieren (4 steps) | 159M | 0.12 |
| MMAudio | 1.03B | 6.49 |
| AudioX | 1.17B | 14.76 |
| LoVA | 1.06B | 8.45 |
| SALSA-V (32 steps) | 643M | 4.84 |
| SALSA-V (4 steps) | 643M | 0.71 |

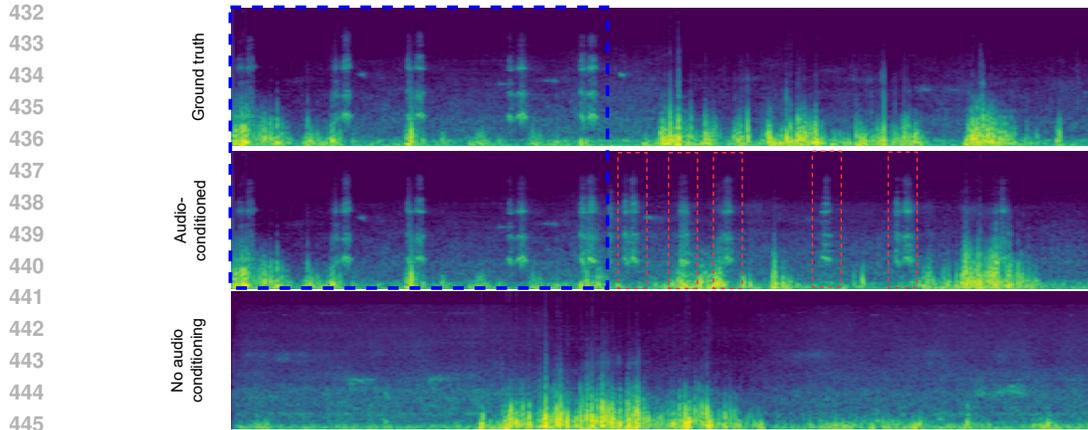


Figure 4: Qualitative example of outpaiting capabilities of SALSAs-V. The video used for this example contains a turkey gobbling (with its appropriate sound). The dashed blue region is given as context. The mel-spectrograms visualize a generation with and without audio conditioning. With audio conditioning, the model uses the characteristic sound present in the first 2 seconds and performs outpaiting beyond that point (outside of the blue box). Both samples use the ground-truth video frames as visual information. With conditioning, the characteristic sound is utilized (the bird’s call, highlighted by red boxes) and other spectral features, such as the frequency top-end, are also better-preserved.

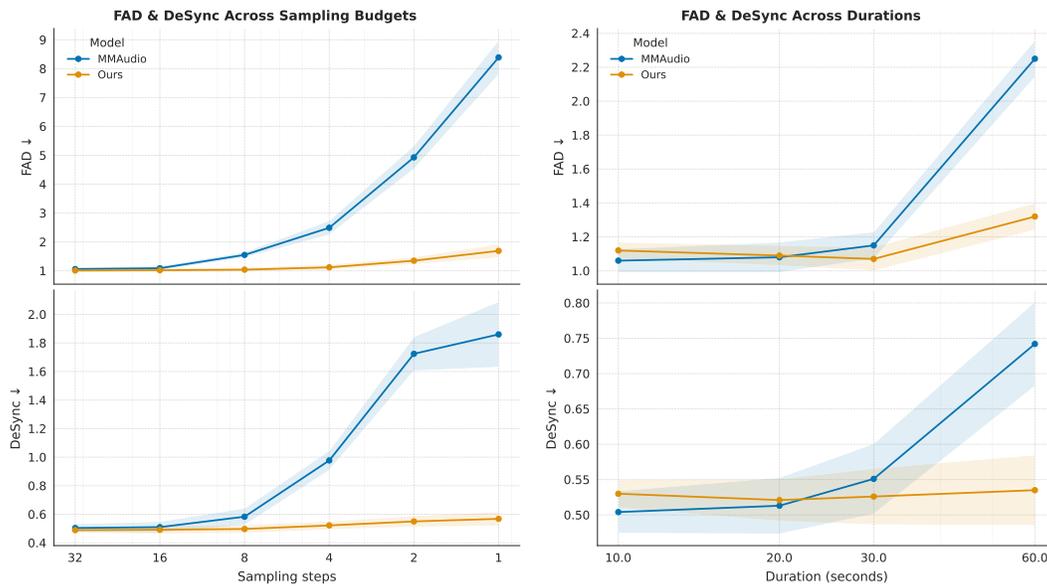


Figure 5: FAD and DeSync across different sampling steps (left) and generation lengths (right). SALSAs-V is able to maintain generation quality with fewer sampling steps, and outperforms previous methods in long-form audio generation.

video events, as is evident by the increased frequency of repetition towards the end of the sample. Our model is able to accurately reuse a sample in the section to be generated, preserving both momentary effects and longer-range global frequency characteristics. We quantify these capabilities in Appendix A.2.

We investigate SALSAs-V’s long-form generation capabilities in terms of distribution matching and synchronization performance by calculating FAD and DeSync metrics for increasing generation lengths by using a set of long-form examples from our test set, with both models being evaluated on samples of a common, fixed duration. The corresponding results are displayed in Figure 5. After

Table 5: Long-form comparison between SALSA-V and LoVA on 30-second samples. SALSA-V displays superior results in all objective metrics.

| Method | FAD _{VGG} ↓ | KL _{PANNs} ↓ | KL _{PaSST} ↓ | IS↑ | IB↑ | DeSync↓ |
|---------|----------------------|-----------------------|-----------------------|--------------|--------------|--------------|
| LoVA | 1.88 | 2.01 | 1.96 | 14.72 | 26.6 | 1.225 |
| SALSA-V | 1.09 | 1.73 | 1.71 | 16.34 | 31.88 | 0.526 |

generation lengths of approximately 20 seconds, MMAudio’s performance begins to suffer, showing a reduction in both FAD and DeSync. By comparison, our model does not face these constraints, as we are able to repeatedly prepend the model’s previous generation, in this case using a 0.5-second long reference section to generate 9.5 new seconds of audio. MMAudio instead processes the full sequence at once. We hypothesize that this improves MMAudio’s performance in this comparison up to the 20-second range, as it has full access to the sequence, while we force our model to rely on a 10 second context in this setting as well. Beyond this range, MMAudio receives sequences with positional embeddings unseen during its training process, which explains the sudden decrease in synchronization performance and overall quality.

We conduct an evaluation of SALSA-V’s long-form generation capabilities on 30-second samples against LoVA (Cheng et al., 2024), the results of which can be seen in Table 5. SALSA-V exhibits superior long-form generation capabilities in all objective metrics, maintaining particularly high scores in terms of synchronization and audio-visual alignment.

Limitations. Despite strong synchronization performance and competitive scores in terms of audiovisual alignment, our human evaluation study (Table 1) reveals a gap in subjective audio quality compared to MMAudio. Given the significantly smaller parameter count of our model (643M vs. 1.03B parameters for MMAudio), it is reasonable to hypothesize that scaling our model further could close or surpass this gap, further improving subjective quality. Future investigations could thus focus on larger-scale models, or techniques such as model distillation to enhance perceptual audio quality without proportionally increasing computational costs.

In addition, our current mechanism for long-form generation relies on the progressive outpainting of a prepended reference sample. This approach is mostly usable for single-shot videos without cuts between scenes displaying considerably different events, as the model will attempt to carry over the previous sound characteristics to this new section. However, this is mostly an issue when relying on automatic progressive extension without manual intervention; in practice, users can take steps such as re-seeding with a new generation or tuning the guidance value for audio conditioning to largely alleviate this issue.

5 CONCLUSION

We introduce SALSA-V, a video-to-audio model that demonstrates significant improvements over baselines in several key aspects. SALSA-V is particularly effective at synchronization with video content and overall quality at increased generation lengths while retaining sampling efficiency. Through a masked training objective, our model is able to jointly support audio conditioning and outpainting, which enables seamless long-form generations. The effectiveness of this approach is underscored by the quality and synchronization maintained over significantly extended durations. Furthermore, through the incorporation of a shortcut loss objective, SALSA-V supports few-step sampling without sacrificing sample fidelity at higher step counts or requiring dedicated fine-tuning procedures. Our experiments demonstrate that this training objective equips our model with the ability to generate high-quality audio samples even with a significantly reduced number of sampling steps. This aspect directly addresses one of the main limitations of conventional diffusion models in the video-to-audio domain, where a large number of sampling steps restrict their practical utility in real-time or resource-constrained applications.

REFERENCES

Hassan Akbari, Liangzhe Yuan, Rui Qian, Wei-Hong Chuang, Shih-Fu Chang, Yin Cui, and Boqing Gong. Vatt: Transformers for multimodal self-supervised learning from raw video, audio and

- 540 text, 2021. URL <https://arxiv.org/abs/2104.11178>.
541
- 542 Anurag Arnab, Mostafa Dehghani, Georg Heigold, Chen Sun, Mario Lučić, and Cordelia Schmid.
543 Vivit: A video vision transformer, 2021. URL <https://arxiv.org/abs/2103.15691>.
544
- 545 Jisheng Bai, Haohe Liu, Mou Wang, Dongyuan Shi, Wenwu Wang, Mark D. Plumbley, Woon-
546 Seng Gan, and Jianfeng Chen. Audiosetcaps: An enriched audio-caption dataset using automated
547 generation pipeline with large audio and language models, 2024. URL <https://arxiv.org/abs/2411.18953>.
548
- 549 James Bradbury, Roy Frostig, Peter Hawkins, Matthew James Johnson, Chris Leary, Dougal
550 Maclaurin, George Necula, Adam Paszke, Jake VanderPlas, Skye Wanderman-Milne, and Qiao
551 Zhang. JAX: composable transformations of Python+NumPy programs, 2018. URL <http://github.com/jax-ml/jax>.
552
- 553 Honglie Chen, Weidi Xie, Andrea Vedaldi, and Andrew Zisserman. Vggsound: A large-scale audio-
554 visual dataset, 2020. URL <https://arxiv.org/abs/2004.14368>.
555
- 556 Tsai-Shien Chen, Aliaksandr Siarohin, Willi Menapace, Ekaterina Deyneka, Hsiang wei Chao,
557 Byung Eun Jeon, Yuwei Fang, Hsin-Ying Lee, Jian Ren, Ming-Hsuan Yang, and Sergey Tulyakov.
558 Panda-70m: Captioning 70m videos with multiple cross-modality teachers, 2024. URL <https://arxiv.org/abs/2402.19479>.
559
- 560 Ziyang Chen, Prem Seetharaman, Bryan Russell, Oriol Nieto, David Bourgin, Andrew Owens, and
561 Justin Salamon. Video-guided foley sound generation with multimodal controls, 2025. URL
562 <https://arxiv.org/abs/2411.17698>.
563
- 564 Ho Kei Cheng, Masato Ishii, Akio Hayakawa, Takashi Shibuya, Alexander Schwing, and Yuki Mit-
565 sutfuji. Mmaudio: Taming multimodal joint training for high-quality video-to-audio synthesis,
566 2025. URL <https://arxiv.org/abs/2412.15322>.
567
- 568 Xin Cheng, Xihua Wang, Yihan Wu, Yuyue Wang, and Ruihua Song. Lova: Long-form video-to-
569 audio generation, 2024. URL <https://arxiv.org/abs/2409.15157>.
570
- 571 Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas
572 Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszko-
573 reit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at
574 scale, 2021. URL <https://arxiv.org/abs/2010.11929>.
- 575 Patrick Esser, Sumith Kulal, Andreas Blattmann, Rahim Entezari, Jonas Müller, Harry Saini, Yam
576 Levi, Dominik Lorenz, Axel Sauer, Frederic Boesel, et al. Scaling rectified flow transformers
577 for high-resolution image synthesis. In *Forty-first international conference on machine learning*,
578 2024.
- 579 Zach Evans, Julian D. Parker, CJ Carr, Zack Zukowski, Josiah Taylor, and Jordi Pons. Stable audio
580 open, 2024. URL <https://arxiv.org/abs/2407.14358>.
581
- 582 Kevin Frans, Danijar Hafner, Sergey Levine, and Pieter Abbeel. One step diffusion via shortcut
583 models, 2025. URL <https://arxiv.org/abs/2410.12557>.
584
- 585 Hugo Flores Garcia, Prem Seetharaman, Rithesh Kumar, and Bryan Pardo. Vampnet: Music gener-
586 ation via masked acoustic token modeling, 2023. URL <https://arxiv.org/abs/2307.04686>.
587
- 588 Jort F. Gemmeke, Daniel P. W. Ellis, Dylan Freedman, Aren Jansen, Wade Lawrence, R. Channing
589 Moore, Manoj Plakal, and Marvin Ritter. Audio set: An ontology and human-labeled dataset for
590 audio events. In *Proc. IEEE ICASSP 2017*, New Orleans, LA, 2017.
591
- 592 Tiantian Geng, Teng Wang, Jinming Duan, Runmin Cong, and Feng Zheng. Dense-localizing audio-
593 visual events in untrimmed videos: A large-scale benchmark and baseline, 2023. URL <https://arxiv.org/abs/2303.12930>.

- 594 Rohit Girdhar, Alaaeldin El-Nouby, Zhuang Liu, Mannat Singh, Kalyan Vasudev Alwala, Armand
595 Joulin, and Ishan Misra. Imagebind: One embedding space to bind them all, 2023. URL <https://arxiv.org/abs/2305.05665>.
596
597
- 598 Yuan Gong, Yu-An Chung, and James Glass. Ast: Audio spectrogram transformer, 2021. URL
599 <https://arxiv.org/abs/2104.01778>.
600
- 601 Jonathan Heek, Anselm Levskaya, Avital Oliver, Marvin Ritter, Bertrand Rondepierre, Andreas
602 Steiner, and Marc van Zee. Flax: A neural network library and ecosystem for JAX, 2024. URL
603 <http://github.com/google/flax>.
604
- 605 Shawn Hershey, Sourish Chaudhuri, Daniel P. W. Ellis, Jort F. Gemmeke, Aren Jansen, R. Channing
606 Moore, Manoj Plakal, Devin Platt, Rif A. Saurous, Bryan Seybold, Malcolm Slaney, Ron J. Weiss,
607 and Kevin Wilson. Cnn architectures for large-scale audio classification, 2017. URL <https://arxiv.org/abs/1609.09430>.
608
- 609 Vladimir Iashin, Weidi Xie, Esa Rahtu, and Andrew Zisserman. Synchformer: Efficient synchro-
610 nization from sparse cues, 2024. URL <https://arxiv.org/abs/2401.16423>.
611
- 612 Mirjam Keetels and Jean Vroomen. The role of spatial disparity and hemifields in audio-visual
613 temporal order judgments. *Experimental Brain Research*, 167(4):635–640, 2005. doi: 10.1007/
614 s00221-005-0067-1.
615
- 616 Kevin Kilgour, Mauricio Zuluaga, Dominik Roblek, and Matthew Sharifi. Fréchet audio distance:
617 A metric for evaluating music enhancement algorithms, 2019. URL <https://arxiv.org/abs/1812.08466>.
618
- 619 Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization, 2017. URL
620 <https://arxiv.org/abs/1412.6980>.
621
- 622 Qiuqiang Kong, Yin Cao, Turab Iqbal, Yuxuan Wang, Wenwu Wang, and Mark D. Plumbley. Panns:
623 Large-scale pretrained audio neural networks for audio pattern recognition, 2020. URL <https://arxiv.org/abs/1912.10211>.
624
- 625 Khaled Koutini, Jan Schlüter, Hamid Eghbal-zadeh, and Gerhard Widmer. Efficient training of au-
626 dio transformers with patchout, September 2022. URL [http://dx.doi.org/10.21437/
Interspeech.2022-227](http://dx.doi.org/10.21437/Interspeech.2022-227).
627
- 628 Xingchao Liu, Chengyue Gong, and Qiang Liu. Flow straight and fast: Learning to generate and
629 transfer data with rectified flow, 2022. URL <https://arxiv.org/abs/2209.03003>.
630
- 631 Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization, 2019. URL <https://arxiv.org/abs/1711.05101>.
632
- 633 Simian Luo, Chuanhao Yan, Chenxu Hu, and Hang Zhao. Diff-foley: Synchronized video-to-
634 audio synthesis with latent diffusion models, 2023. URL [https://arxiv.org/abs/2306.
17203](https://arxiv.org/abs/2306.17203).
635
- 636 Xinhao Mei, Chutong Meng, Haohe Liu, Qiuqiang Kong, Tom Ko, Chengqi Zhao, Mark D. Plumb-
637 ley, Yuexian Zou, and Wenwu Wang. Wavcaps: A chatgpt-assisted weakly-labelled audio caption-
638 ing dataset for audio-language multimodal research. *IEEE/ACM Transactions on Audio, Speech,
639 and Language Processing*, 32:3339–3354, 2024. ISSN 2329-9304. doi: 10.1109/taslp.2024.
640 3419446. URL <http://dx.doi.org/10.1109/TASLP.2024.3419446>.
641
- 642 Mathew Monfort, Alex Andonian, Bolei Zhou, Kandan Ramakrishnan, Sarah Adel Bargal, Tom
643 Yan, Lisa Brown, Quanfu Fan, Dan Gutfrund, Carl Vondrick, and Aude Oliva. Moments in time
644 dataset: one million videos for event understanding, 2019. URL [https://arxiv.org/abs/
1801.03150](https://arxiv.org/abs/1801.03150).
645
- 646 Andrew Owens, Phillip Isola, Josh McDermott, Antonio Torralba, Edward H. Adelson, and
647 William T. Freeman. Visually indicated sounds, 2016. URL [https://arxiv.org/abs/
1512.08512](https://arxiv.org/abs/1512.08512).

- 648 William Peebles and Saining Xie. Scalable diffusion models with transformers, 2023. URL <https://arxiv.org/abs/2212.09748>.
649
650
- 651 Julien Rabin, Gabriel Peyré, Julie Delon, and Marc Bernet. Wasserstein barycenter and its applica-
652 tion to texture mixing. In Alfred M. Bruckstein, Bart M. ter Haar Romeny, Alexander M. Bron-
653 stein, and Michael M. Bronstein (eds.), *Scale Space and Variational Methods in Computer Vision*,
654 pp. 435–446, Berlin, Heidelberg, 2012. Springer Berlin Heidelberg. ISBN 978-3-642-24785-9.
- 655 Tim Salimans, Ian Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, and Xi Chen. Im-
656 proved techniques for training gans, 2016. URL <https://arxiv.org/abs/1606.03498>.
657
- 658 Stability AI. sd3.5: Stable diffusion 3.5 inference-only reference implementation. <https://github.com/Stability-AI/sd3.5>, 2024. MIT License. Accessed 2025-09-10.
659
- 660 Jianlin Su, Yu Lu, Shengfeng Pan, Ahmed Murtadha, Bo Wen, and Yunfeng Liu. Roformer: En-
661 hanced transformer with rotary position embedding, 2023. URL <https://arxiv.org/abs/2104.09864>.
662
- 663 Zeyue Tian, Yizhu Jin, Zhaoyang Liu, Ruibin Yuan, Xu Tan, Qifeng Chen, Wei Xue, and Yike
664 Guo. Audiox: Diffusion transformer for anything-to-audio generation, 2025. URL <https://arxiv.org/abs/2503.10522>.
665
666
- 667 Michael Tschannen, Alexey Gritsenko, Xiao Wang, Muhammad Ferjad Naeem, Ibrahim Alabdul-
668 mohsin, Nikhil Parthasarathy, Talfan Evans, Lucas Beyer, Ye Xia, Basil Mustafa, Olivier Hénaff,
669 Jeremiah Harmsen, Andreas Steiner, and Xiaohua Zhai. Siglip 2: Multilingual vision-language
670 encoders with improved semantic understanding, localization, and dense features, 2025. URL
671 <https://arxiv.org/abs/2502.14786>.
- 672 Ilpo Viertola, Vladimir Iashin, and Esa Rahtu. Temporally aligned audio for video with autoregres-
673 sion, 2024. URL <https://arxiv.org/abs/2409.13689>.
674
- 675 Yi Wang, Kunchang Li, Xinhao Li, Jiashuo Yu, Yanan He, Chenting Wang, Guo Chen, Baoqi Pei,
676 Ziang Yan, Rongkun Zheng, Jilan Xu, Zun Wang, Yansong Shi, Tianxiang Jiang, Songze Li,
677 Hongjie Zhang, Yifei Huang, Yu Qiao, Yali Wang, and Limin Wang. Internvideo2: Scaling
678 foundation models for multimodal video understanding, 2024. URL <https://arxiv.org/abs/2403.15377>.
679
- 680 Yongqi Wang, Wenxiang Guo, Rongjie Huang, Jiawei Huang, Zehan Wang, Fuming You, Ruiqi
681 Li, and Zhou Zhao. Frieren: Efficient video-to-audio generation network with rectified flow
682 matching, 2025. URL <https://arxiv.org/abs/2406.00320>.
- 683 Jingjing Xu, Xu Sun, Zhiyuan Zhang, Guangxiang Zhao, and Junyang Lin. Understanding and
684 improving layer normalization, 2019. URL <https://arxiv.org/abs/1911.07013>.
685
- 686 Xiaohua Zhai, Alexander Kolesnikov, Neil Houlsby, and Lucas Beyer. Scaling vision transformers,
687 2022. URL <https://arxiv.org/abs/2106.04560>.
- 688 Xiaohua Zhai, Basil Mustafa, Alexander Kolesnikov, and Lucas Beyer. Sigmoid loss for language
689 image pre-training, 2023. URL <https://arxiv.org/abs/2303.15343>.
690
- 691 Yiming Zhang, Yicheng Gu, Yanhong Zeng, Zhening Xing, Yuancheng Wang, Zhizheng Wu, and
692 Kai Chen. Foleyrafter: Bring silent videos to life with lifelike and synchronized sounds, 2024.
693 URL <https://arxiv.org/abs/2407.01494>.
- 694 Long Zhao, Nitesh B. Gundavarapu, Liangzhe Yuan, Hao Zhou, Shen Yan, Jennifer J. Sun, Luke
695 Friedman, Rui Qian, Tobias Weyand, Yue Zhao, Rachel Hornung, Florian Schroff, Ming-Hsuan
696 Yang, David A. Ross, Huisheng Wang, Hartwig Adam, Mikhail Sirotenko, Ting Liu, and Boqing
697 Gong. Videoprism: A foundational visual encoder for video understanding, 2025. URL <https://arxiv.org/abs/2402.13217>.
698
699
700
701

702
703
704
705
706
707
708
709
710
711
712
713
714
715
716
717
718
719
720
721
722
723
724
725
726
727
728
729
730
731
732
733
734
735
736
737
738
739
740
741
742
743
744
745
746
747
748
749
750
751
752
753
754
755

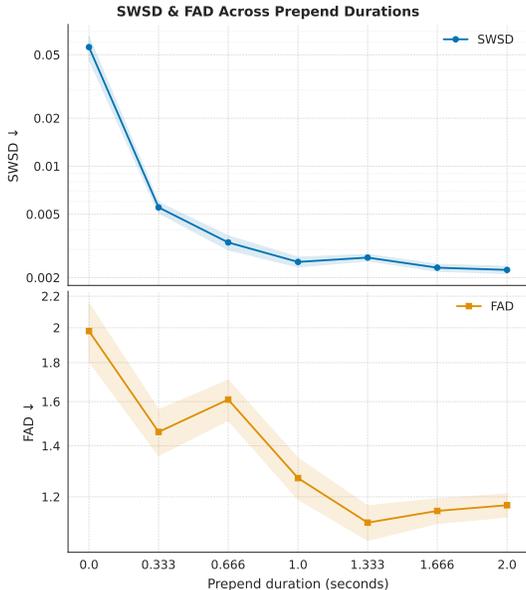


Figure 6: SWSD and FAD over increasing audio conditioning length. A prepending duration of 0 indicates a generation without audio-conditioning. Only the non-prepended section is considered for all calculations. Both metrics show increased similarity to the ground-truth audio for longer conditioning durations, with SWSD showing a more consistent reduction in spectral distance.

A APPENDIX

A.1 MASKED AUDIO INFERENCE

We utilize classifier-free-guidance in conjunction with audio conditioning during inference as such:

$$\hat{v}(x_t, c, t) = v_\theta(x_t, t) + w \cdot (v_\theta(\hat{x}_t, c, t) - v_\theta(x_t, t)) \tag{3}$$

where w is the guidance scale, c is a collection of conditioning signals (semantic visual, synchronization visual, and text in our case), v_θ is the model’s velocity prediction, and \hat{x}_t represents the sequence to be generated at step t with masking applied. More precisely, \hat{x}_t is composed of x_t and x_1 (the unconditional latent sequence, and the ground-truth sequence, respectively) in the following manner:

$$\hat{x}_t = (x_t^{(1)}, \dots, x_t^{(n)}) \parallel x_1 \parallel (x_t^{(n+|x_1|+1)}, \dots, x_t^{(T)})$$

where n depends on the desired starting position of the reference sequence, and $T = |x_t|$ is the length of the audio sequence to be generated.

A.2 SPECTRAL CHARACTERISTICS OF AUDIO CONDITIONING

In order to quantify our model’s ability to faithfully utilize audio conditioning information, we evaluate two different metrics: the per-sample FAD between the ground truth reference and the generated result, and a spectral difference metric intended to capture finer-grained frequency characteristics. This latter metric is intended to show whether specific characteristics of the frequency spectrum are preserved, while staying time-invariant (as we assess the quality of synchronization independently). We do this by treating the ground-truth and generated mel-spectrograms as probability distributions, calculating a Monte-Carlo estimate of the sliced Wasserstein distance (Rabin et al., 2012) between them. We refer to this metric as *Sliced-Wasserstein Spectral Distance* (SWSD).

To calculate SWSD, we first turn each signal (reference and generated) into a sequence of log-mel vectors (one per short-time frame). We then perform column-wise ℓ^1 normalization, such that each

vector behaves like a probability mass function over mel bands. Those vectors are then treated as empirical samples from two high-dimensional distributions $\mathcal{P}_{\text{ref}}, \mathcal{P}_{\text{gen}} \subset \mathbb{R}^{\text{n_mels}}$ between which we want to calculate a geometry-aware probability distance metric.

However, calculating the full Wasserstein distance between high-dimensional distributions is computationally infeasible, so we use a "sliced" variant, where the original high-dimensional samples are repeatedly projected on a number of randomly drawn unit vectors, between which we can efficiently calculate the 1-Wasserstein distance. This serves as an approximation of the full Wasserstein distance, with the estimate's accuracy increasing with the number of random unit vectors, which we set to 100 in our calculations. The mean of the 1-Wasserstein distances is used as our result. For details of our calculation, refer to Algorithm 1.

Both metrics are displayed in Figure 6, which shows the impact of audio-conditioned generation with different conditioning lengths on both FAD and SWSD. Values are calculated over a representative subset of audios with a maximum length of 10 seconds, which contain characteristic sounds corresponding to visual entities present in both the conditioning audio section, as well as in the section to be generated. Note that both metrics are only evaluated over these unseen sections (i.e., excluding the conditioning segment), in order to ensure a fair comparison. Our qualitative assessment is thus confirmed by both FAD and SWSD, which both show a significant decrease over the context length, before tapering off. While FAD is somewhat noisier, SWSD shows a steady decrease.

Algorithm 1 Sliced-Wasserstein Spectral Distance

Require: Audio signals $x_{\text{ref}}, x_{\text{gen}} \in \mathbb{R}^n$, sampling rate sr
Require: Hyper-parameters: n_mels , hop_ms , win_ms , number of projections K , numerical ε
Ensure: Distance $d \in \mathbb{R}_{\geq 0}$

```

1: function SLICEDWSMELDIST( $x_{\text{ref}}, x_{\text{gen}}$ )
2:    $F_{\text{ref}} \leftarrow \text{LOGMELFRAMES}(x_{\text{ref}})$  ▷  $T_{\text{ref}} \times \text{n\_mels}$ 
3:    $F_{\text{gen}} \leftarrow \text{LOGMELFRAMES}(x_{\text{gen}})$  ▷  $T_{\text{gen}} \times \text{n\_mels}$ 
4:   if  $T_{\text{ref}} \neq T_{\text{gen}}$  then
5:      $T \leftarrow \min(T_{\text{ref}}, T_{\text{gen}})$ 
6:     Randomly subsample each  $F$  to  $T$  rows
7:   else
8:      $T \leftarrow T_{\text{ref}}$  ▷ common frame count
9:   end if
10:  Sample  $K$  random unit vectors  $\{u_k\}_{k=1}^K \subset \mathbb{R}^{\text{n\_mels}}$ 
11:  for  $k = 1$  to  $K$  do
12:     $p_{\text{ref}} \leftarrow F_{\text{ref}} u_k$  ▷ project frames
13:     $p_{\text{gen}} \leftarrow F_{\text{gen}} u_k$ 
14:    Sort  $p_{\text{ref}}$  and  $p_{\text{gen}}$  ascending
15:     $d_k \leftarrow \frac{1}{T} \sum_{i=1}^T |p_{\text{ref}}[i] - p_{\text{gen}}[i]|$ 
16:  end for
17:  return  $d \leftarrow \frac{1}{K} \sum_{k=1}^K d_k$ 
18: end function
19: function LOGMELFRAMES( $x$ )
20:    $\text{hop} \leftarrow \lfloor sr \cdot \text{hop\_ms} / 1000 \rfloor$ 
21:    $\text{win} \leftarrow \lfloor sr \cdot \text{win\_ms} / 1000 \rfloor$ 
22:    $S \leftarrow \text{MELSPECTROGRAM}(x, sr, \text{n\_mels}, \text{win}, \text{hop})$ 
23:    $S \leftarrow \log(S + \varepsilon)$ 
24:   for each column  $j$  do  $S[:, j] \leftarrow S[:, j] / \sum_b S[b, j]$  ▷  $\ell^1$  normalisation
25:   return  $S^T$  ▷ shape  $(T, \text{n\_mels})$ 
26: end function

```

804 A.3 HUMAN EVALUATION

805
806 We develop an online interface to facilitate blind and randomized trials for our human listening
807 evaluation. Figure 7 shows an example trial page. Listeners are tasked with rating 3 different tracks
808 for a given video along 3 different aspects, each using a 5-point MOS scale:

- 809 • **Overall Fit:** How well the given audio track semantically corresponds to the video.

810
811
812
813
814
815
816
817
818
819
820
821
822
823
824
825
826
827
828
829
830
831
832
833
834
835
836
837
838
839
840
841
842
843
844
845
846
847
848
849
850
851
852
853
854
855
856
857
858
859
860
861
862
863

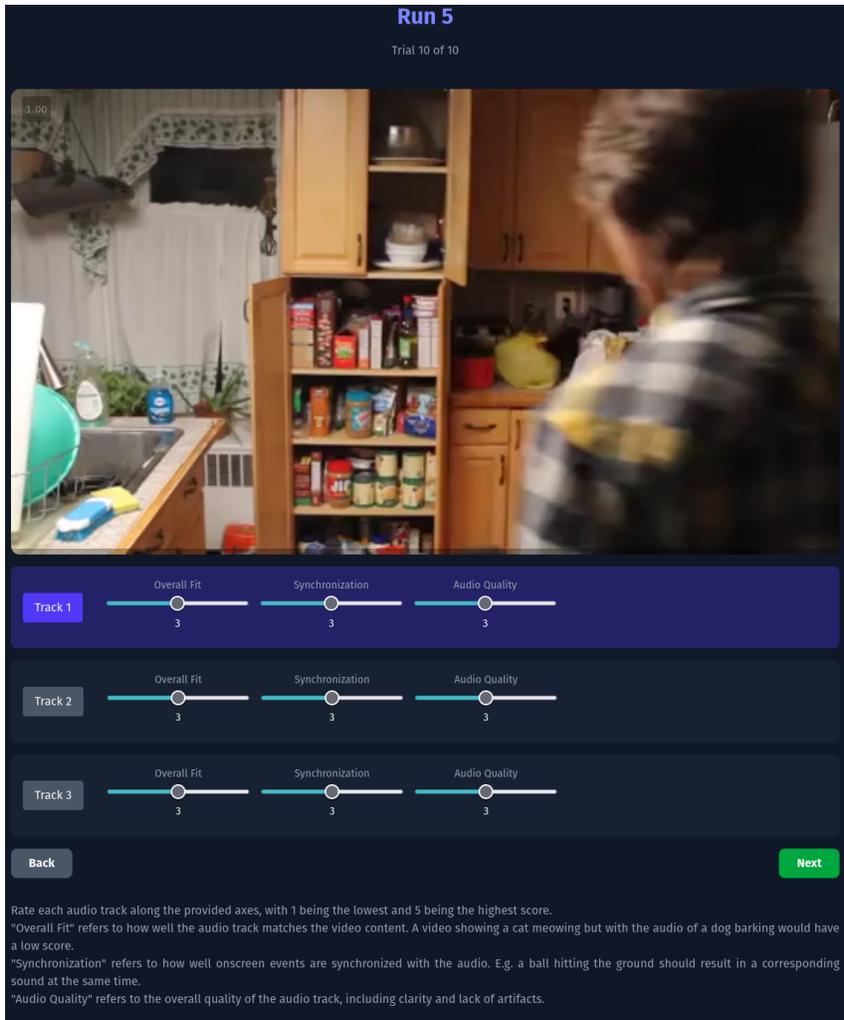


Figure 7: An example screenshot showing the interface of the listening test used for our human evaluation study. Audio track identifiers are randomized both between trials and runs.

- **Synchronization:** How precisely video events are synchronized to their corresponding sound.
- **Audio Quality:** General quality of the audio, encompassing aspects such as artifacts and clarity.

Each listener completes one test run, where a single run consists of 10 trials. In our case, this resulted in 180 individual scores. Particular care was taken to ensure that synchronization is maintained in the web interface when switching between audio tracks, as even small drifts over time can affect results.

A.4 IMPLEMENTATION DETAILS

Length Unification. To unify the sequence lengths of our synchronization conditioning signal to the audio input, we employ a combination of cubic interpolation and a learned upsampling layer, yielding features that are frame-aligned to the audio latent sequence.

More precisely, resampled synchronization video features are calculated as follows:

$$\tilde{v}_{\text{syn}} = \text{LayerNorm}(\text{stopgrad}(\text{interp}(v_{\text{syn}}))) + \tanh(g) \cdot \text{ConvTranspose}(v_{\text{syn}}) \quad (4)$$

864 where v_{syn} are the visual synchronization features at their original temporal resolution, g is a learned
 865 gating value, "interp" denotes cubic interpolation, and "stopgrad" is the stop-gradient operation. We
 866 find that stopping the gradient from flowing through the interpolation path helps stabilize training
 867 early on and prevents the model from collapsing the learned upsampling path.

868 **Training & Evaluation Details.** We use a batch size of 30 for training of the contrastive audio-
 869 visual alignment model, which is done on VGGSound (Chen et al., 2020). We keep the VideoPrism
 870 (Zhao et al., 2025) backbone frozen while adding 4 additional spatial, and 2 temporal encoder layers
 871 (using the same underlying architecture as ViViT (Arnab et al., 2021)) which are trainable. Further-
 872 more, we unfreeze the complete audio backbone and perform training at a learning rate of 3×10^{-5} .
 873 The new task layers are set to a learning rate of 1×10^{-4} . We use the AdamW (Loshchilov &
 874 Hutter, 2019) optimizer with a weight decay of 1×10^{-4} and a cosine learning rate scheduler with
 875 2000 warmup steps, starting at an initial learning rate of 0. All models are trained for 1M steps.
 876 We discard input samples with an audio/frame ratio outside of a 5% grace range centered around
 877 the value expected for our configuration ($\frac{16000 \text{ audio samples}}{24 \text{ frames}} \approx 666.67$). This is done to ensure precise
 878 synchronization between visual events and their corresponding sound.

879 The visual encoder outputs a tensor of shape (t_v, s, d) per input snippet, where t_v , s , and d are the
 880 temporal, spatial, and depth dimensions, respectively. The spatial dimension s is reduced using a
 881 MAP head, yielding a (t_v, d) tensor. This tensor is used as-is as the synchronization condition-
 882 ing feature for the generative model, while additionally undergoing a pooling operation during the
 883 contrastive model training stage (yielding a single d -sized vector per input snippet for the loss cal-
 884 culation).

885 The audio encoder outputs a tensor of shape (f, t_a, d) , where f , t_a , and d are the frequency, time,
 886 and depth dimensions. This tensor is processed using a similar aggregation operation as the visual
 887 output, but along the frequency dimension, which results in an output of shape (t_a, d) . After pooling
 888 along the time dimensions, the resulting vector is used in the contrastive loss calculation.

889 The main generative model is trained with a batch size of 300 and a consistency target ratio of 25%,
 890 using the Adam (Kingma & Ba, 2017) optimizer with a learning rate of 1×10^{-4} . Time steps for
 891 the noising process are sampled according to a log-normal distribution, and we use EMA with a
 892 target update rate of 0.999, training for 1M steps in total. All conditioning signals are dropped out
 893 with a probability of 0.1 during training, and 0.25 of the input sequences in a batch are modified
 894 by the audio masking process which affects a random contiguous subsection of the sequence with a
 895 randomly chosen size between 1 and 88 latent tokens. We use a CFG guidance value of 4.0 for our
 896 experiments while using author-recommended settings for all other models.

897 Our large model uses 6 multi-modal blocks where the first 3 include the optional self-attention
 898 operation over the audio sequence, followed by 6 additional single-modality DiT blocks. Our smaller
 899 model variant uses 4 multi-modal blocks, with the first 2 including separate self-attention, followed
 900 by 4 single-modality blocks.

901 All models are implemented in JAX (Bradbury et al., 2018) using the Flax (Heek et al., 2024)
 902 machine learning library and trained on TPU v4-8 nodes. Other relevant models used during training
 903 or preprocessing (AST (Gong et al., 2021), Stable-Audio VAE (Evans et al., 2024)) were converted
 904 to JAX.

905 906 907 908 A.5 VAE TRAINING

909 We fine-tune the Stable-Audio VAE (Evans et al., 2024), increasing its sampling rate from 21.5 Hz
 910 to 43 Hz by replacing the first encoder and last decoder layers. The VAE is then fine-tuned from this
 911 state for 800k steps, using a batch size of 128. This training process was conducted on audio samples
 912 from VGGSound (Chen et al., 2020), AudioSet (Gemmeke et al., 2017), and our high-fidelity studio
 913 Foley dataset.

914 This was done both to increase the quality of the audio codec, as it is limited to a large degree by
 915 the downsampling ratio used in the encoder, and to align the training data distribution more closely
 916 with our intended application, as Stable-Audio VAE is trained on a significant amount of speech
 917 data (while we focus on general environmental sounds).

918
919
920
921
922
923
924
925
926
927
928
929
930
931
932
933
934
935
936
937
938
939
940
941
942
943
944
945
946
947
948
949
950
951
952
953
954
955
956
957
958
959
960
961
962
963
964
965
966
967
968
969
970
971

A.6 ADDITIONAL RESULTS

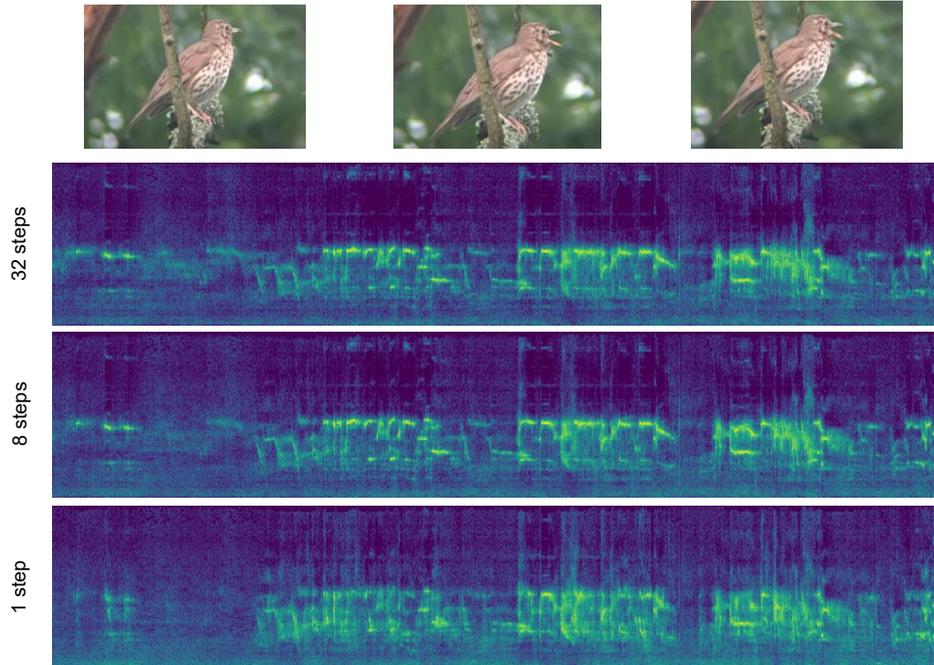


Figure 8: Spectrograms of samples for a 10-second video of a bird chirping generated at a different number of sampling steps (32, 8, and 1, i.e. single-shot generation) using our shortcut model. General audio structure is preserved among all settings, with the single-shot generation mostly losing fidelity in higher frequency ranges and short transients.

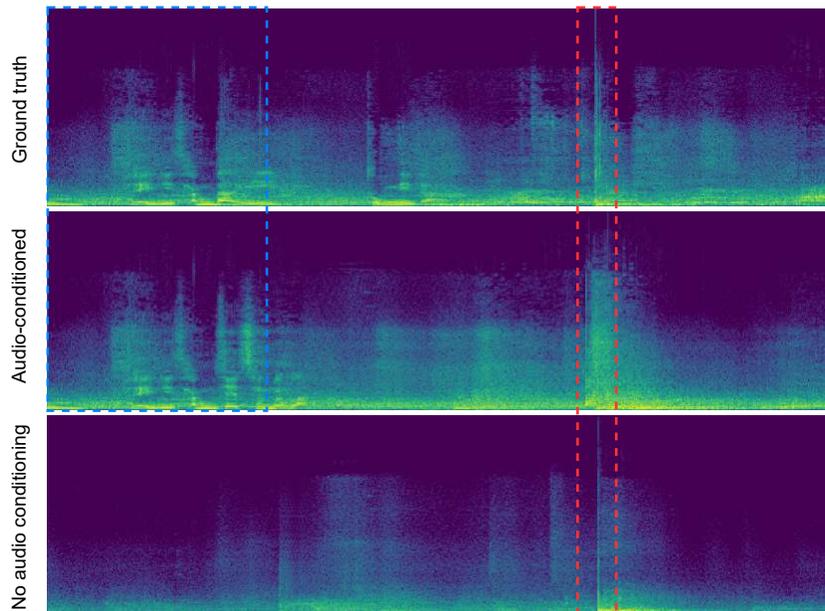


Figure 9: Spectrograms of samples for a 5-second video of a pronounced impact sound with visual correspondence (a baseball bat hitting a ball) with and without audio conditioning. The audio conditioning portion is in the blue box, while the impact is highlighted with red. As can be seen, audio-conditioned generations also remain responsive to visual cues (i.e. no over-reliance on ground-truth audio).