











Sound-guided Semantic Video Generation

Seung Hyun Lee¹, Gyeongrok Oh¹, Wonmin Byeon², Chanyoung Kim¹,
Won Jeong Ryoo¹, Sang Ho Yoon³, Hyunjun Cho¹, Jihyun Bae¹,
Jinkyu Kim^{4*}, and Sangpil Kim^{1*}

¹Department of Artificial Intelligence, Korea University

²NVIDIA Research, NVIDIA Corporation

³Graduate School of Culture Technology, KAIST

⁴Department of Computer Science and Engineering, Korea University

Abstract. The recent success in StyleGAN demonstrates that pre-trained StyleGAN latent space is useful for realistic video generation. However, the generated motion in the video is usually not semantically meaningful due to the difficulty of determining the direction and magnitude in the StyleGAN latent space. In this paper, we propose a framework to generate realistic videos by leveraging multimodal (sound-image-text) embedding space. As sound provides the temporal contexts of the scene, our framework learns to generate a video that is semantically consistent with sound. First, our sound inversion module maps the audio directly into the StyleGAN latent space. We then incorporate the CLIP-based multimodal embedding space to further provide the audio-visual relationships. Finally, the proposed frame generator learns to find the trajectory in the latent space which is coherent with the corresponding sound and generates a video in a hierarchical manner. We provide the new high-resolution landscape video dataset (audio-visual pair) for the sound-guided video generation task. The experiments show that our model outperforms the state-of-the-art methods in terms of video quality. We further show several applications including image and video editing to verify the effectiveness of our method.

Keywords: Sound, Multi-modal Representation, Video Generation.

1 Introduction

Existing video generation methods rely on motion generation from a noise vector given an initial frame [34,35,43]. As they create trajectory from noise vectors without any guidance, the motion of the generated video is not semantically meaningful. On the other hand, sound provides the cue for motion and various context of the scene [23]. Specifically, sound may represent events of the scene such as ‘viola playing’ or ‘birds singing’. It can also provide a tone of the scene such as ‘Screaming’ or ‘Laughing’. This is an important cue for generating a video because the video’s temporal component and motion are closely associated with sound. Generating high-fidelity video from sound is crucial for a variety of applications, such as multimedia content creation and filmmaking.

However, most sound-based video generation works focus on generating a talking face matching from verbal sound using 2D facial landmarks [6,7,32] or 3D face representation [5,26,33,38,40,44,46].

Generating a realistic video from non-verbal sound is challenging. First, the spatial and temporal contexts of a generated video need to be semantically consistent with sound. Although some studies [3,20] have tried to synthesize video with sound as an independent variable, mapping from 1D semantic information from audio to visual signal is not straightforward. Furthermore, when a sound of waves enters as an input, these works can not generate a video with diverse waves. Second, the generated motion should be physically plausible and temporally coherent between frames. Recent works [1,12] use a latent mapper network with music as the input for navigating in StyleGAN [15] latent space. Since the direction in the latent space is randomly provided, however, the content semantics from the video are not realistic. Finally, there is a lack of high-fidelity video dataset for generating a realistic video from non-verbal sound.

To overcome these challenges, we propose a novel framework that can semantically invert sound to the StyleGAN latent space for video generation. The difficulty of inverting audio into the StyleGAN latent space is that the image regenerated after inverting the audio is visually irrelevant to the audio. To resolve this issue, we employ the joint embedding prior knowledge learned from large-scale multimodal data (image, text, sound) to sound inverting. A high fidelity video is generated by moving the latent vector in the $\mathcal{W}+$ space which is disentangled latent space of StyleGAN. Our sound-encoder learns to put the latent space $\mathcal{W}+$ from the sound input and find the trajectory from the initial latent code for video generation. As shown in Fig. 1, we generate a video whose meaning is consistent with the given sound-input.

Contrastive Language-Image Pre-Training (CLIP) [25] creates a very powerful joint multimodal embedding space with 400 million image-text pair data. We leverage a CLIP-based multimodal embedding space (sound-image, text-image) trained on a large scale at the same time. This multimodal embedding space provides guidance when walking in the $\mathcal{W}+$ space. StyleCLIP [24] and TediGAN [42] leverage the representational power of CLIP for text-driven image editing. Recently, the prior knowledge of CLIP has been transferred to the audio-visual relationship, making it possible to develop various applications using CLIP in sound modality [41,21]. In this paper, we exploit the prior knowledge of audio-visual to generate a complete video related to sound. Furthermore, we introduce a frame generator that finds the desired trajectory from the sound. The frame generator predicts the latent code of the next time by using audio

Acknowledgement. This work is partially supported by Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea government(MSIT) (No. 2019-0-00079, Artificial Intelligence Graduate School Program(Korea University)). J. Kim is partially supported by the National Research Foundation of Korea grant (NRF-2021R1C1C1009608), Basic Science Research Program (NRF-2021R1A6A1A13044830), and ICT Creative Consilience program (IITP-2022-2022-0-01819). Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the funding agency.

*Corresponding authors: Jinkyu Kim and Sangpil Kim. Code and more diverse examples are available at <https://kuai-lab.github.io/eccv2022sound/>.

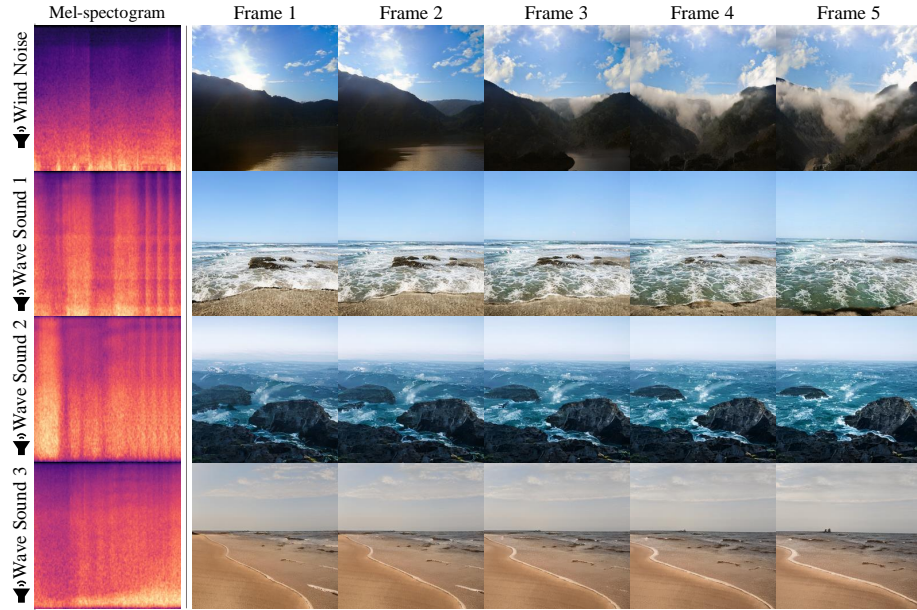


Fig. 1: An example of sound-guided semantic video generation. Our method generates semantic video reflecting temporal context using audio as an input feature.

as a constraint. StyleGAN’s style latent vector has a hierarchical representation, and our frame generator which is divided into three layers predicts coarse, mid, and fine style, respectively. Therefore, our proposed method provides nonlinear guidance for each time step beyond the linear interpolation of latent vectors (see Fig. 2). The latent code of our proposed method travels in the guidance of the audio in the latent space of StyleGAN with generating audio semantic meaningful and diverse output. In contrast, the general interpolation method travels in the latent space in one direction without any guide at regular intervals.

In this paper, our experimental results demonstrate that the proposed method supports a variety of sound sources and generates video that is a better reflection of given audio information. The existing audio-visual benchmark datasets [4, 16, 18] have limitations in high-fidelity video generation. We release a new high-resolution landscape video dataset, as the existing datasets are not designed for high-fidelity video generation [4] or do not support sound-video pairs [16, 18]. Our dataset is effective for experimenting with high-resolution audio-visual generation models.

Our main contributions are listed as follows:

- We propose a novel method for mapping audio semantics into the StyleGAN latent space by matching multi-modal semantics with CLIP space.
- We introduce a framework for semantic-level video generation solely based on the given audio. We demonstrate the effectiveness of the proposed method by

outperforming the quality of generated videos from state-of-the-art methods in the sound-guided video generation task.

- Our framework not only determines the direction and size of movement of the latent code in the StyleGAN latent space, but also can generate rich style video given various sound attributes.
- We provide the new high-resolution landscape video dataset with audio-visual pairs for this task.

2 Related Work

StyleGAN Latent Space Analysis.

StyleGAN [15] has a semantically rich latent space, so many studies [24,27,42] have analyzed the latent space of GAN. Most of the existing StyleGAN inversion modules minimize the difference between a latent vector and an embedding vector that encodes an image generated from the vector. Richardson *et al.* [27] have effectively inverted high-quality face images using prior knowledge related to the faces such as LPIPS [45], perceptual loss and pretrained ArcFace [8]. Patashnik *et al.* [24] defines a style latent mapper network that projects a given latent code to match the meaning of the text prompt. The latent mapper of the study learns given a pair of image and text prompts sampled from noise. However, there is a limitation in that randomly sampled image cannot have continuous motion like a real video.

Our goal is to indicate the direction in which the style latent code will move in an audio-visual multimodal embedding space. Our study embeds the input audio directly into the StyleGAN latent space. The sound-guided latent code enables video generation in the expanded audio-visual domain. CLIP [25] represents the powerful joint embedding space between image and text modalities with 400 million image-text pairs, and this joint embedding space provides direction for traversing the StyleGAN latent space. Recent studies [10,21,41] extend the modalities of CLIP to audio. Lee *et al.* [21] especially focused on audio-visual representation learning for image editing, and we also leverage that audio-visual multimodal space embedding for navigating the latent code. Our works differ from previous works in that we




Fig. 2: Comparison of traversing method for the video generation. The top grid shows the conventional interpolated way, and the bottom grid shows our audio-guided latent code movement method on StyleGAN’s latent space.

consider temporal style variation guided by the CLIP embedding space. Sound semantics generate images at each time step using the latent vector optimized by the CLIP guidance.

Sound-guided Video Generation. Many studies use the temporal dynamics in sound as a source for vivid video generation [3,12,20]. There are mainly two approaches to generate video from sound as a source. The first is a conditional variational autoencoder method [17] that predicts the distribution of future video frames in the latent space. VAE-based method [3,43,20] mainly solve video prediction tasks that predict the next frame for a given frame. Among them, Chatterjee, *et al.* [3] and Le, *et al.* [20] consider the context of a scene with a non-verbal sound. Ji *et al.* [13] and Lahiri *et al.* [19] generate a video about facial expression with a verbal sound.

In this study, we enable the video generation task with a non-verbal sound condition that is more difficult than the verbal sound condition. GAN-based video generation is to sample the video from a noise vector [28,35,37,39]. Tian *et al.* [34] and Fox *et al.* [9] synthesize continuous videos with StyleGAN [15], a pre-trained high-fidelity image generator. Unlike previous studies, we consider a novel sound-guided high resolution video generation method. Jeong *et al.* [12] explores StyleGAN’s latent space to generate a video. However, the domain of the sound is limited to music, and the guidance is not noticed by the user. We generate high-resolution video considering the semantics of sounds such as wind, raining, etc.

3 Sound-guided Video Generation

Our model takes sound information as an input to generate a sequence of video frames accordingly, as shown in Fig. 3. For example, given a Laughing sound input, our model generates a video with a facial expression of laughing. To achieve this goal, our model needs two main capabilities. (i) The ability to understand the sound input and to condition it in the trained video generator. (ii) The ability to generate a video sequence that is perceptually realistic and is temporally consistent.

We propose that such capabilities can be learned via our two novel modules: (1) *Sound Inversion Encoder*, which learns a mapping from the sound input to a latent code in the (pre-trained) StyleGAN [15] latent space (Section 3.1). (2) *StyleGAN-based Video Generator*, which is conditioned on the sound-guided latent code and generates video frames accordingly (Section 3.2). Furthermore, we leverage the representation power of the CLIP [25]-based multimodal (image, text, and audio) joint embedding space, which regularizes perceptual consistency between the sound input and the generated video.

3.1 Inverting Sound Into the StyleGAN Latent Space

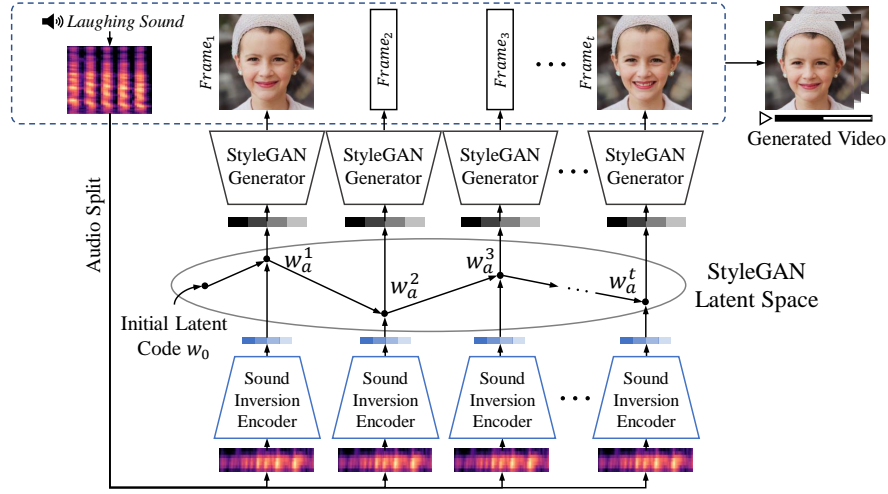


Fig. 3: An overview of our proposed sound-guided video generation model. Our model consists of two main modules: (i) Sound Inversion Encoder (Section 3.1), which takes a sequence of audio inputs as an input and outputs a latent code to generate video frames. (ii) StyleGAN-based Video Generator (Section 3.2), which generates temporally consistent and perceptually realistic video frames conditioned on the sound input.

As shown in Fig. 4, our Sound Inversion Encoder takes as an input Mel-spectrogram Acoustic features and outputs a latent feature $\mathbf{w}_a \in \mathcal{W}_+$ in the pre-trained StyleGAN feature space \mathcal{W}_+ . The sound-conditioned latent feature \mathbf{w}_a is then augmented by the element-wise summation with the randomly sampled latent code $\mathbf{w} \in \mathcal{W}_+$, yielding a sound-guided latent code $\hat{\mathbf{w}}_a \in \mathcal{W}_+$. Conditioned on $\hat{\mathbf{w}}_a$, we generate an image, which maintains the content of the original image (generated with the random latent code \mathbf{w}) but its style is transferred according to the semantic of sound. Formally,

$$\hat{\mathbf{w}}_a = E_a(\mathbf{x}_a) + \mathbf{w}. \quad (1)$$

where $E_a(\cdot)$ denotes our Sound Inversion Encoder given a sound input \mathbf{x}_a .

Matching Multimodal Semantics via CLIP Space. Lee *et al.* [21] introduced an extended CLIP-based multi-modal (image, text, and audio) feature

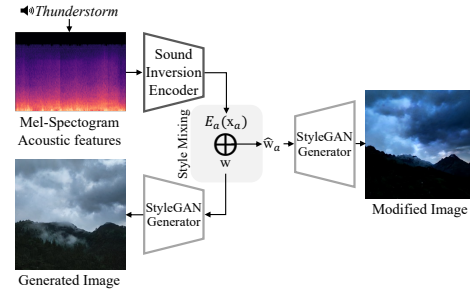


Fig. 4: Audio Inversion Encoder. From a given sound input, sound-guided latent code is obtained by the elementwise summation with the randomly sampled latent code.

space, which is trained to produce a joint embedding space where a positive triplet pair (e.g., audio input: “thunderstorms”, text: “thunderstorm”, and corresponding image) are mapped close together in the CLIP-based embedding space, while pushing that of negative pair samples further away. We utilize this pre-trained CLIP-based embedding space to generate images that are semantically well-aligned with the sound input. Specifically, we minimized the following cosine distance (in the CLIP embedding space) between the image generated from the latent code $\hat{\mathbf{w}}_a$ and an audio input \mathbf{x}_a .

$$\mathcal{L}_{\text{CLIP}}^{(a \leftrightarrow v)} = 1 - \frac{F_v(G(\hat{\mathbf{w}}_a)) \cdot F_a(\mathbf{x}_a)}{\|F_v(G(\hat{\mathbf{w}}_a))\|_2 \cdot \|F_a(\mathbf{x}_a)\|_2}. \quad (2)$$

where $G(\cdot)$ denotes the StyleGAN [15] generator, while $F_v(\cdot)$ and $F_a(\cdot)$ are CLIP’s image encoder and audio encoder, respectively. A similar loss function could be used for a pair of audio and text prompts. CLIP’s text embedding provide constraints on the semantics of audio chunks. Using audio labels (e.g. thunderstorm) as a text prompt, we minimize the cosine distance between representations for image and text in the CLIP embedding space.

Lastly, we use l_2 distance between the latent code $\hat{\mathbf{w}}_a$ and their averaged latent code, i.e. $\bar{\mathbf{w}}_a = \sum_t \hat{\mathbf{w}}_a^t$, to explicitly constrain the generated sequence of images to share similar contents and semantics over time. Ultimately, we minimize the following loss to train our Sound Inversion Encoders.

$$\mathcal{L}_{\text{enc}} = \mathcal{L}_{\text{CLIP}}^{(a \leftrightarrow v)} + \mathcal{L}_{\text{CLIP}}^{(a \leftrightarrow t)} + \lambda_b \|\hat{\mathbf{w}}_a - \bar{\mathbf{w}}_a\|_2^2. \quad (3)$$

where λ_b controls the strength of the regularization term.

3.2 Sound-guided Semantic Video Generation

Recurrent Module for Latent Code Sequence Generation. As shown in Fig. 5, our model recurrently generates latent code at each timestep to generate an image sequence with the StyleGAN [15] generator. Formally, we train a recurrent neural network $E_{\text{RNN}}(\cdot)$ that outputs $\hat{\mathbf{w}}_a^t$ conditioned on the previous audio segment \mathbf{x}_a^{t-1} and the hidden state \mathbf{h}_t at timestep t .

$$\hat{\mathbf{w}}_a^t, \mathbf{h}_t = E_{\text{RNN}}(\mathbf{h}_{t-1}, E_a(\mathbf{x}_a^{t-1}) + \hat{\mathbf{w}}_a^{t-1}), \quad (4)$$

where $E_{\text{RNN}}(\cdot)$ comprises a fully-connected layers to recurrently yield the hidden state \mathbf{h}_t and the current latent code $\hat{\mathbf{w}}_a^t$.

Multiple Recurrent Blocks for Rich Style Information. The input latent vectors are divided into three groups (0-2, 3-5, and 6-15 layers) and fed to different recurrent networks to generate different levels of detail. These three latent groups represent different styles (coarse, mid, and fine features) and control the attributes in StyleGAN [24]. Each recurrent block predicts the latent code of the next time step:

$$\hat{\mathbf{w}}_a^t = (E_{\text{RNN}}^{\text{coarse}}(\hat{\mathbf{w}}_a^{t-1}), E_{\text{RNN}}^{\text{mid}}(\hat{\mathbf{w}}_a^{t-1}), E_{\text{RNN}}^{\text{fine}}(\hat{\mathbf{w}}_a^{t-1})), \quad (5)$$

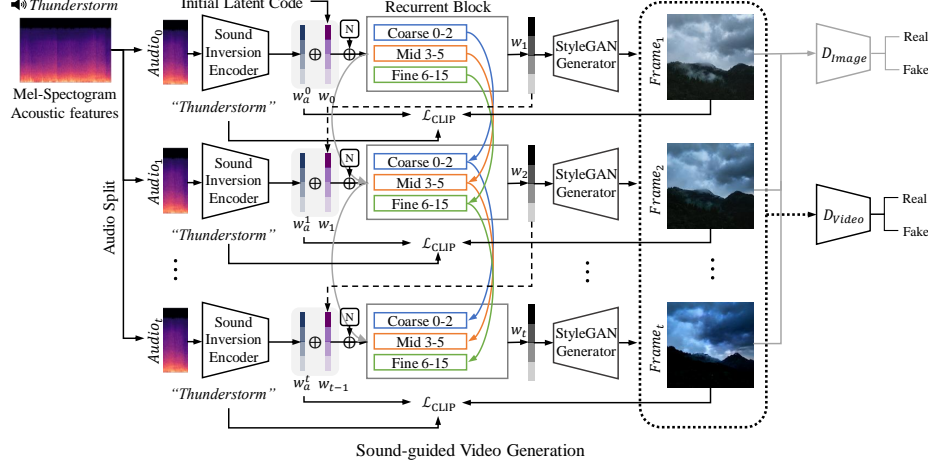


Fig. 5: An overview of our sound-guided video generation model, which consists of two main parts: (i) Sound Inversion Encoder, which iteratively generates sound-conditioned latent code \mathbf{w}_a^t from corresponding audio time segments. (ii) StyleGAN-based Video Generator, which recurrently generates a video frame that is trained to be consistent with neighboring frames. Moreover, we train image and video discriminators adversarially to generate perceptually realistic video frames.

where $E_{\text{RNN}}^{\text{coarse}}(\cdot)$, $E_{\text{RNN}}^{\text{mid}}(\cdot)$, and $E_{\text{RNN}}^{\text{fine}}(\cdot)$ denote coarse, mid, and fine recurrent encoder-decoder network, respectively. Generated video $\tilde{\mathbf{v}}$ is a sequence of images synthesized from each latent code as follows:

$$\tilde{\mathbf{v}} = [G(\hat{\mathbf{w}}_a^1), G(\hat{\mathbf{w}}_a^2), \dots, G(\hat{\mathbf{w}}_a^T)]. \quad (6)$$

where T is the total sequence length of the video.

Adversarial Image and Video Discriminators. As shown in Fig. 5, we train a video discriminator D_V adversarially by forwarding the generated video into the discriminator, which determines whether the input video is real or synthesized. Specifically, following MoCoGAN-HD [34], our video discriminator is based on the architecture of PatchGAN [11]. An input video (i.e. a real or fake example) is divided into small 3D patches, which are then classified as real or fake. The average response is used as the final output. We thus use the following adversarial loss \mathcal{L}_{D_V} :

$$\mathcal{L}_{D_V} = \mathbb{E}[\log D_V(\mathbf{v})] + \mathbb{E}[1 - \log D_V(\tilde{\mathbf{v}})], \quad (7)$$

where \mathbf{v} and $\tilde{\mathbf{v}}$ are the real and fake example, respectively. Additionally, we also apply an image discriminator D_I , which similarly trains the model adversarially to determine whether the input image is real or fake on the time axis. Concretely, we optimize the following loss function:

$$\mathcal{L}_D = \mathcal{L}_{D_V} + \mathcal{L}_{D_I} = \mathcal{L}_{D_V} + \mathbb{E}[\log D_I(\mathbf{v})] + \mathbb{E}[1 - \log D_I(\tilde{\mathbf{v}})]. \quad (8)$$

Table 1: Comparison table of the High Fidelity Audio-visual Landscape Video Dataset and other datasets. (\dagger): We report the smallest resolution for comparisons.

Dataset	# of Videos	Resolution †	# of Classes	Audio-Video Pairs	Total Length (hours)
HMDB51 [18]	6,766	340×256	51	\times	4.9
UCF-101 [31]	13,320	320×240	101	\checkmark	26.7
VGG-Sound [4]	200,000	640×360	310	\checkmark	560.0
Kinetics-400 [16]	306,245	640×360	400	\checkmark	850.7
Sub-URMP [22]	72	1920×1080	13	\checkmark	1.0
Landscape (ours)	9,280	1280×720	9	\checkmark	25.8

Table 2: Comparison to the state-of-the-art methods. We compare two version of our method (with and without sound inputs) to several state-of-the-art methods. We reproduce Sound2Sight [3], CCVS [20], and TrumerAI [12] as a baseline on two benchmark datasets: Sub-URMP [22] and our created Landscape.

Method	<i>use</i> sound inputs	<i>use</i> the first frame	Sub-URMP [22]		Landscape	
			IS (\uparrow)	FVD (\downarrow)	IS (\uparrow)	FVD (\downarrow)
Sound2Sight [3]	\checkmark	\checkmark	1.64 ± 0.11	282.48	1.55 ± 0.48	311.55
CCVS [20]	\checkmark	\checkmark	2.06 ± 0.07	274.01	1.78 ± 0.64	305.40
TrumerAI [12]	\checkmark	-	1.02 ± 0.14	350.80	1.07 ± 0.58	732.63
Ours (<i>w/o</i> sound inputs)	-	-	2.99 ± 0.48	272.27	1.68 ± 0.43	305.91
Ours	\checkmark	-	3.05 ± 0.42	271.03	1.82 ± 0.26	291.88

Loss Function. Our model can be trained end-to-end by optimizing the following objective:

$$\min_{\theta_G} \max_{\theta_{D_V}} \mathcal{L}_{D_V} + \min_{\theta_G} \max_{\theta_{D_I}} \mathcal{L}_{D_I} + \min_{\theta_{E_a}} \lambda_{\text{enc}} \mathcal{L}_{\text{enc}}, \quad (9)$$

where λ_{enc} controls the strength of the loss term.

4 Experiments

4.1 Datasets

There exist few publicly available Audio-Video paired datasets (e.g. UCF-101 [31], VGG-Sound [4], Kinetics-400 [16]), but they mostly support low-resolution video and contain a lot of noise in the audio. (see Table 1). To the best of our knowledge, there is the Sub-URMP (University of Rochester Musical Performance) dataset [22] that only provides pairs of high-fidelity audio-video. This dataset provides 72 video-audio pairs (≈ 1 hour in total) from recordings of 13 kinds of instruments played by different orchestra musicians. This dataset would be a good starting point to evaluate the performance of video generation models,

Table 3: Quantitative evaluation. (a) Semantic Consistency between a given audio and generated video. (b) Ablation study of CLIP Loss. We evaluate the Inversion Score on the LHQ dataset.

(a) Semantic Consistency.			(b) The LHQ Dataset Inversion Score.		
Model	Similarity (t \leftrightarrow v) \uparrow	Similarity (a \leftrightarrow v) \uparrow	Model	LPIPS \downarrow	MSE \downarrow
Sound2Sight [3]	0.2491	0.1771	$\mathcal{L}_2 + \mathcal{L}_{reg}$	0.648	0.052
CCVS [20]	0.2514	0.1764	$\mathcal{L}_2 + \mathcal{L}_{LPIPS} + \mathcal{L}_{reg}$	0.468	0.070
TräumerAI [12]	0.1932	0.1416	$\mathcal{L}_2 + \mathcal{L}_{LPIPS} + \mathcal{L}_{reg} + \mathcal{L}_{CLIP}$	0.432	0.048
Ours	0.2556	0.1897			

but it only focuses on orchestra playing scenes with a limited number of video clips, limiting training efficiency. Thus, we create a new dataset that provides a high-fidelity audio-video paired dataset on landscape scenes.

High Fidelity Audio-Video Landscape Dataset (Landscape). We collect 928 high-resolution (at least 1280×720) video clips, where each clip is divided into 10 (non-overlapped) different clips of 10 seconds each. Overall, the total number of video clips available is 9,280, and the total length is approximately 26 hours. Our Landscape dataset contains 9 different scenes, such as thunderstorms, waterfall burbling, volcano eruption, squishing water, wind noise, fire crackling, raining, underwater bubbling, and splashing water. We provide details of our created Landscape dataset in the supplemental material.

4.2 Comparison to Existing Sound-guided Video Generation

Quantitative Evaluation. We use the latest VAE-based models as baselines, including Sound2Sight [3], CCVS [20] and StyleGAN-based model TräumerAI [12]. For our model and TräumerAI, we first pre-train StyleGAN on the high fidelity benchmark datasets (the Sub-URMP [22] and the LHQ datasets [30]) then train to navigate the latent space with the fixed image generator. The videos are generated with randomly sampled initial frames. In contrast, for Sound2Sight and CCVS, the first frames are provided. Table 2 shows that our approach produces the best quality results, and sound information is effective for video generation.

Table 3a demonstrates that our method produces more visually correlated videos from sound than other methods. To measure if the generated videos are semantically related to sound, we compare the cosine similarity between text-audio and video embedding. We obtain 512-dimensional video embeddings corresponding to the number of each frame with the CLIP [25] image encoder, and average them. Additionally, text embeddings are obtained from CLIP’s text encoder, a 512-dimensional vector, and audio embeddings sharing CLIP space are obtained from Lee’s [21] multi-modal embedding space. So leveraging the multi-modal embedding space helps to achieve semantic consistency in sound-guided video generation.

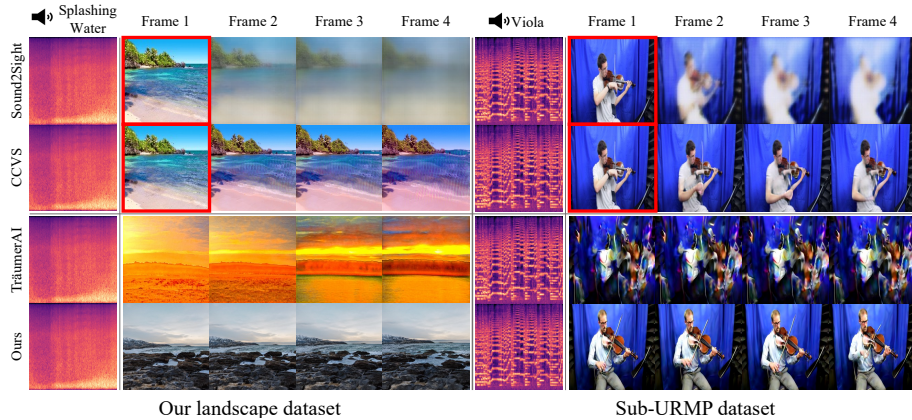


Fig. 6: Qualitative comparison between our sound-guided semantic video generation method and previous video generation results on our landscape dataset and Sub-URMP dataset [22]. Sound2Sight [3] and CCVS [20] (top) generate a video by conditioning on an image as the first frame highlighted in red, whereas TrumerAI [12], and our method (bottom) generate a video conditioned on audio input only.

Evaluation Metrics. Evaluation of the quality of generated videos is known to be challenging. We first use the two widely-used quantitative metrics: Inception Score (IS) [29] and Frechet Video Distance (FVD) [36]. The former is widely used to evaluate the outputs of GANs by measuring the KL-divergence between each image’s label distribution and the marginal label distribution. Further, FVD quantifies the video quality by measuring the distribution gap between the real vs. synthesized videos in the latent space. We use an Inception3D [2] network, which is pre-trained on Kinetics-400 [16] and is fine-tuned on each benchmark accordingly.

Qualitative Evaluation. In Fig. 6, we visually compare the quality of the generated video with other models including Sound2Sight [3], CCVS [20] and TrumerAI [12] for the two benchmark datasets (Sub-URMP [22] and our landscape dataset). We found that TrumerAI is not able to produce any realistic video at all. Also, the generated video by Sound2Sight and CCVS is mostly distorted. Our method, on the other hand, produces high fidelity videos.

4.3 Ablation Studies

CLIP Loss for StyleGAN Inversion. In this study, we show that CLIP [25] prior knowledge is helpful for StyleGAN inversion. Table 3b compares the mean squared error (MSE) and LPIPS [45] between the original and the reconstruction image from the inversion module using the LHQ dataset. By minimizing the cosine distance between CLIP embeddings, the inversion reconstruction performance is improved in landscape images.

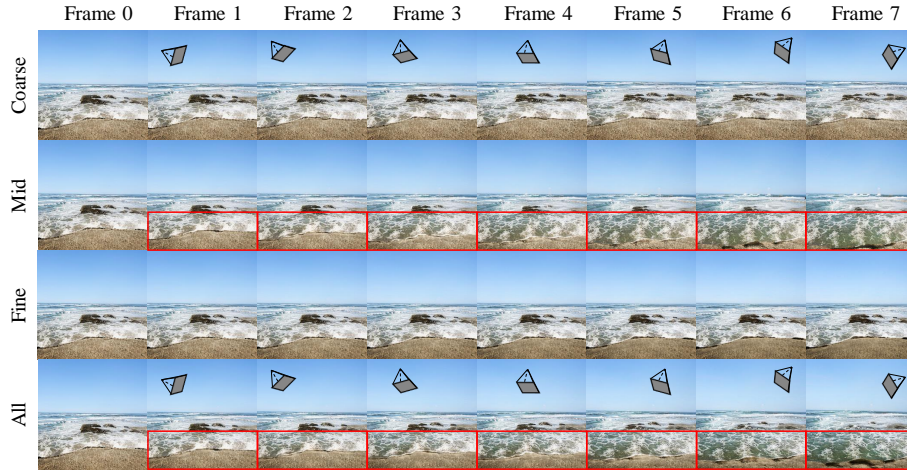


Fig. 7: Ablation of multiple recurrent blocks for rich style information. The first row shows the video generation result when only the coarse recurrent block is applied, the second is mid, the third is the fine recurrent block, and the last is the entire block. The red box is the area where the style has changed.

Effect of Multiple Recurrent Blocks for Rich Style Information. In sound-guided video generation, multiple recurrent blocks control the StyleGAN attribute. Fig. 7 compares video generation using only one coarse, mid, and fine recurrent block. Each block contains diverse style information. Specifically, the coarse and middle recurrent blocks control viewpoint and semantically meaningful motion changes (such as a wave strike) of the scene. Finally, the fine recurrent block handles fine texture changes.

4.4 User Study

In order to evaluate our proposed method, we request one hundred Amazon Mechanical Turk (AMT) participants. We show participants three types of generated videos that are generated by Sound2Sight [3], CCVS [20], TraumerAI [12], our model and ground truth. Participants answer the following questionnaire based on five-point Likert scale: (i) Realness - *Please evaluate the realness of the video* (“1 - very unrealistic” to “5 - very realistic”) and (ii) Naturalness - *Which video generation result better expresses the target attribute?* (iii) Semantic Consistency - *Please evaluate the semantic consistency between video and audio* (“1 - very inconsistent” to “5 - very consistent”).

The survey procedure is as follows. First, we have participants watch ground truth videos and measure realness and semantic consistency on a scale of 1 to 5. Then have participants measure realness against the results generated by other models. To evaluate naturalness, we ask participants to choose which of the videos generated by each model best expresses the target attribute. As




Fig. 8: User Study Results on our landscape dataset. (a) Realness. (b) Naturalness. (c) Semantic Consistency.

shown in Fig. 8, our method significantly outperforms other state-of-the-art approaches (*Realness*, *Naturalness*, *Semantic Consistency*).

4.5 Sound-guided Video Editing with Text Constraints

The combination of text and sound makes it possible to generate various videos with rich information. Fig. 9 demonstrates that the combination of text-query constraints and sound results in a much more versatile style of video. Because our model shares the CLIP [25] space trained with large-scale text and image pairs, video generation with text constraint is available. In addition, it is possible to apply styles of different levels with the style mixing technique. For video editing, we randomly sample multiple latent codes. The initial latent code for generating the first frame is the sample with the closest CLIP embedding cosine distance of the text prompt.

5 Discussion

Limitations. In our method, we sample the first frame from the StyleGAN [14,15] latent space, and the space covers only the train domain. So, color change is rarely observed when the pre-trained StyleGAN and the video datasets have a domain gap. Still, we observe that our model is applicable to many cases such as sound-guided face video generation (see sup.).

In addition, in order for our model to learn to generate a video, the weights of the pre-trained image generator are required, which increases the training time. And maintaining the movement size and identity of the latent code predicted by the recurrent block has a tradeoff relationship (see sup.).

Societal Impact. Following this characteristic of our method, the user can generate a real video with the user’s desire, and this usefulness allows the user to see the video that existed only in the user’s imagination. However, although



Fig. 9: Examples of sound-guided semantic video generation from sound with text prompts. Our model generates the videos by guiding sound from latent code containing text content. The time interval between each column is five frames.

these features can provide a suitable output video to the user if the input video source is a work of art or is not ethically correct, the result produced may imply the contrary intentions of the original video creator or may raise ethical concerns.

6 Conclusion

This paper shows a new way to make sound-guided realistic videos by taking advantage of the multi-modal embedding space, which uses sound, images, and text. To be more specific, we use CLIP [25] space to encode the sound information into the StyleGAN [14,15] latent space, which allows generating videos that contain the corresponding source sound semantics. We design our model with the recurrent neural network since video consists of multiple frames and requires time-domain consistency. Multiple recurrent blocks temporarily represent rich style information (view point, sound-dependent frame’s fine features). Additionally, we curate a new high-fidelity audio-video landscape dataset to validate our proposed method, which is superior to other methods. We demonstrate that our method qualitatively and quantitatively outperformed state-of-the-art methods for sound-guided video generation. The proposed model can be used in various applications such as video generation with text constraints.

References

1. Brouwer, H.: Audio-reactive latent interpolations with stylegan. In: NeurIPS 2020 Workshop on Machine Learning for Creativity and Design (2020)
2. Carreira, J., Zisserman, A.: Quo vadis, action recognition? a new model and the kinetics dataset. In: proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 6299–6308 (2017)
3. Chatterjee, M., Cherian, A.: Sound2sight: Generating visual dynamics from sound and context. In: European Conference on Computer Vision. pp. 701–719. Springer (2020)
4. Chen, H., Xie, W., Vedaldi, A., Zisserman, A.: Vggsound: A large-scale audio-visual dataset. In: ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP). pp. 721–725. IEEE (2020)
5. Chen, L., Cui, G., Liu, C., Li, Z., Kou, Z., Xu, Y., Xu, C.: Talking-head generation with rhythmic head motion. In: ECCV (2020)
6. Chen, L., Maddox, R.K., Duan, Z., Xu, C.: Hierarchical cross-modal talking face generation with dynamic pixel-wise loss. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 7832–7841 (2019)
7. Das, D., Biswas, S., Sinha, S., Bhowmick, B.: Speech-driven facial animation using cascaded gans for learning of motion and texture. In: ECCV (2020)
8. Deng, J., Guo, J., Xue, N., Zafeiriou, S.: Arcface: Additive angular margin loss for deep face recognition. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 4690–4699 (2019)
9. Fox, G., Tewari, A., Elgharib, M., Theobalt, C.: Stylevideogan: A temporal generative model using a pretrained stylegan. arXiv preprint arXiv:2107.07224 (2021)
10. Guzhov, A., Raue, F., Hees, J., Dengel, A.: Audioclip: Extending clip to image, text and audio (2021)
11. Isola, P., Zhu, J.Y., Zhou, T., Efros, A.A.: Image-to-image translation with conditional adversarial networks. CVPR (2017)
12. Jeong, D., Doh, S., Kwon, T.: Träumerai: Dreaming music with stylegan. arXiv preprint arXiv:2102.04680 (2021)
13. Ji, X., Zhou, H., Wang, K., Wu, W., Loy, C.C., Cao, X., Xu, F.: Audio-driven emotional video portraits. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 14080–14089 (2021)
14. Karras, T., Aittala, M., Laine, S., Härkönen, E., Hellsten, J., Lehtinen, J., Aila, T.: Alias-free generative adversarial networks. *Advances in Neural Information Processing Systems* **34** (2021)
15. Karras, T., Laine, S., Aila, T.: A style-based generator architecture for generative adversarial networks. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 4401–4410 (2019)
16. Kay, W., Carreira, J., Simonyan, K., Zhang, B., Hillier, C., Vijayanarasimhan, S., Viola, F., Green, T., Back, T., Natsev, P., et al.: The kinetics human action video dataset. arXiv preprint arXiv:1705.06950 (2017)
17. Kingma, D.P., Welling, M.: Auto-encoding variational bayes. arXiv preprint arXiv:1312.6114 (2013)
18. Kuehne, H., Jhuang, H., Garrote, E., Poggio, T., Serre, T.: HMDB: a large video database for human motion recognition. In: Proceedings of the International Conference on Computer Vision (ICCV) (2011)

19. Lahiri, A., Kwatra, V., Frueh, C., Lewis, J., Bregler, C.: Lipsync3d: Data-efficient learning of personalized 3d talking faces from video using pose and lighting normalization. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 2755–2764 (2021)
20. Le Moing, G., Ponce, J., Schmid, C.: Ccvs: Context-aware controllable video synthesis. *Advances in Neural Information Processing Systems* **34** (2021)
21. Lee, S.H., Roh, W., Byeon, W., Yoon, S.H., Kim, C.Y., Kim, J., Kim, S.: Sound-guided semantic image manipulation. *arXiv preprint arXiv:2112.00007* (2021)
22. Li, B., Liu, X., Dinesh, K., Duan, Z., Sharma, G.: Creating a multitrack classical music performance dataset for multimodal music analysis: Challenges, insights, and applications. *IEEE Transactions on Multimedia* **21**(2), 522–535 (2018)
23. Mesaros, A., Heittola, T., Virtanen, T., Plumbley, M.D.: Sound event detection: A tutorial. *IEEE Signal Processing Magazine* **38**(5), 67–83 (2021). <https://doi.org/10.1109/MSP.2021.3090678>
24. Patashnik, O., Wu, Z., Shechtman, E., Cohen-Or, D., Lischinski, D.: Styleclip: Text-driven manipulation of stylegan imagery. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*. pp. 2085–2094 (October 2021)
25. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al.: Learning transferable visual models from natural language supervision. *Image* **2**, T2
26. Richard, A., Lea, C., Ma, S., Gall, J., de la Torre, F., Sheikh, Y.: Audio- and gaze-driven facial animation of codec avatars. In: *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*. pp. 41–50 (January 2021)
27. Richardson, E., Alaluf, Y., Patashnik, O., Nitzan, Y., Azar, Y., Shapiro, S., Cohen-Or, D.: Encoding in style: a stylegan encoder for image-to-image translation. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 2287–2296 (2021)
28. Saito, M., Matsumoto, E., Saito, S.: Temporal generative adversarial nets with singular value clipping. In: *Proceedings of the IEEE international conference on computer vision*. pp. 2830–2839 (2017)
29. Salimans, T., Goodfellow, I., Zaremba, W., Cheung, V., Radford, A., Chen, X.: Improved techniques for training gans. *Advances in neural information processing systems* **29** (2016)
30. Skorokhodov, I., Sotnikov, G., Elhoseiny, M.: Aligning latent and image spaces to connect the unconnectable. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*. pp. 14144–14153 (October 2021)
31. Soomro, K., Zamir, A.R., Shah, M.: Ucf101: A dataset of 101 human actions classes from videos in the wild. *arXiv preprint arXiv:1212.0402* (2012)
32. Suwajanakorn, S., Seitz, S.M., Kemelmacher-Shlizerman, I.: Synthesizing obama. *ACM Transactions on Graphics (TOG)* **36**, 1 – 13 (2017)
33. Thies, J., Elgharib, M.A., Tewari, A., Theobalt, C., Nießner, M.: Neural voice puppetry: Audio-driven facial reenactment. In: *ECCV* (2020)
34. Tian, Y., Ren, J., Chai, M., Olszewski, K., Peng, X., Metaxas, D.N., Tulyakov, S.: A good image generator is what you need for high-resolution video synthesis. In: *International Conference on Learning Representations* (2021), <https://openreview.net/forum?id=6puCSjH3hwA>
35. Tulyakov, S., Liu, M.Y., Yang, X., Kautz, J.: Mocogan: Decomposing motion and content for video generation. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 1526–1535 (2018)

36. Unterthiner, T., van Steenkiste, S., Kurach, K., Marinier, R., Michalski, M., Gelly, S.: Towards accurate generative models of video: A new metric & challenges. arXiv preprint arXiv:1812.01717 (2018)
37. Vondrick, C., Pirsiaavash, H., Torralba, A.: Generating videos with scene dynamics. *Advances in neural information processing systems* **29**, 613–621 (2016)
38. Wang, T.C., Mallya, A., Liu, M.Y.: One-shot free-view neural talking-head synthesis for video conferencing. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* (2021)
39. Wang, Y., Bilinski, P., Bremond, F., Dantcheva, A.: G3an: disentangling appearance and motion for video generation. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 5264–5273 (2020)
40. Wu, H., Jia, J., Wang, H., Dou, Y., Duan, C., Deng, Q.: Imitating arbitrary talking style for realistic audio-driven talking face synthesis. In: *Proceedings of the 29th ACM International Conference on Multimedia*. pp. 1478–1486 (2021)
41. Wu, H.H., Seetharaman, P., Kumar, K., Bello, J.P.: Wav2clip: Learning robust audio representations from clip (2021)
42. Xia, W., Yang, Y., Xue, J.H., Wu, B.: Tedigan: Text-guided diverse face image generation and manipulation. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 2256–2265 (2021)
43. Yan, W., Zhang, Y., Abbeel, P., Srinivas, A.: Videogpt: Video generation using vq-vae and transformers. arXiv preprint arXiv:2104.10157 (2021)
44. Yi, R., Ye, Z., Zhang, J., Bao, H., Liu, Y.J.: Audio-driven talking face video generation with learning-based personalized head pose (2020)
45. Zhang, R., Isola, P., Efros, A.A., Shechtman, E., Wang, O.: The unreasonable effectiveness of deep features as a perceptual metric. In: *CVPR* (2018)
46. Zhou, Y., Han, X., Shechtman, E., Echevarria, J., Kalogerakis, E., Li, D.: Makeittalk: Speaker-aware talking-head animation. *ACM Transactions on Graphics* **39**(6) (2020)