

Robust Sound-Guided Image Manipulation

Seung Hyun Lee^a, Gyeongrok Oh^a, Wonmin Byeon^c, Sang Ho Yoon^d,
Jinkyu Kim^{*b}, Sangpil Kim^{*a}

^a*Department of Artificial Intelligence, Korea University, Seoul, 02841, South Korea*

^b*Department of Computer Science and Engineering, Korea
University, Seoul, 02841, South Korea*

^c*NVIDIA Research, NVIDIA Corporation, CA, 95051, USA*

^d*Graduate School of Culture Technology, KAIST, Daejeon, 34141, South Korea*

Abstract

Recent successes suggest that an image can be manipulated by a text prompt, e.g., a landscape scene on a sunny day is manipulated into the same scene on a rainy day driven by a text input “raining”. These approaches often utilize a StyleCLIP-based image generator, which leverages multi-modal (text and image) embedding space. However, we observe that such text inputs are often bottlenecked in providing and synthesizing rich semantic cues, e.g., differentiating heavy rain from rain with thunderstorms. To address this issue, we advocate leveraging an additional modality, sound, which has notable advantages in image manipulation as it can convey more diverse semantic cues (vivid emotions or dynamic expressions of the natural world) than texts. In this paper, we propose a novel approach that first extends the image-text joint embedding space with sound and applies a direct latent optimization method to manipulate a given image based on audio input, e.g., the sound of rain. Our extensive experiments show that our sound-guided image manipulation approach produces semantically and visually more plausible manipulation results than the state-of-the-art text and sound-guided image manipulation methods, which are further confirmed by our human evaluations. Our downstream task evaluations also show that our learned image-text-sound joint embedding space effectively encodes sound inputs.

Keywords: Image manipulation, self-supervised learning, multi-modal representation learning, sound.

^{*}Corresponding authors: Sangpil Kim (spk7@korea.ac.kr) and Jinkyu Kim (jinkyukim@korea.ac.kr)




Figure 1: Our model manipulates the given source image (top row) guided by sound inputs, e.g., coughing. We compare our image manipulation results (bottom row) with existing approaches: text-driven modified image (Patashnik et al., 2021) and our prior work (Lee et al. (Lee et al., 2022)).

1. Introduction

Image manipulation is a task that naturally edits images according to the user’s intention. Recent studies provide guidance on manipulation using multi-modal information such as sketches (Park et al., 2019b), text (Xia et al., 2021; Patashnik et al., 2021; Gatys et al., 2016), and sound (Lee et al., 2022). To apply the user’s intention into the image, a mixture of sketches and text is used to perform image manipulation and synthesis (Park et al., 2019b; Xia et al., 2021). The user’s intention can be applied by drawing a picture (Park et al., 2019b) or writing text with semantic meanings (Xia et al., 2021; Gatys et al., 2016). We mainly focus on image modification based on user-provided sound. Because discrete input like text cannot fully express the user’s intention when it is used alone, sound provides an alternative option, as it provides other properties such as continuity and dynamics over time. Such properties of sound allow image manipulation to be more dynamic and create different levels of semantics or sentiment in a scene. For example, every video clip of “thunder” generates a different loudness and different characteristics of “*sound of thunder*”. In addition, when a piece of music is

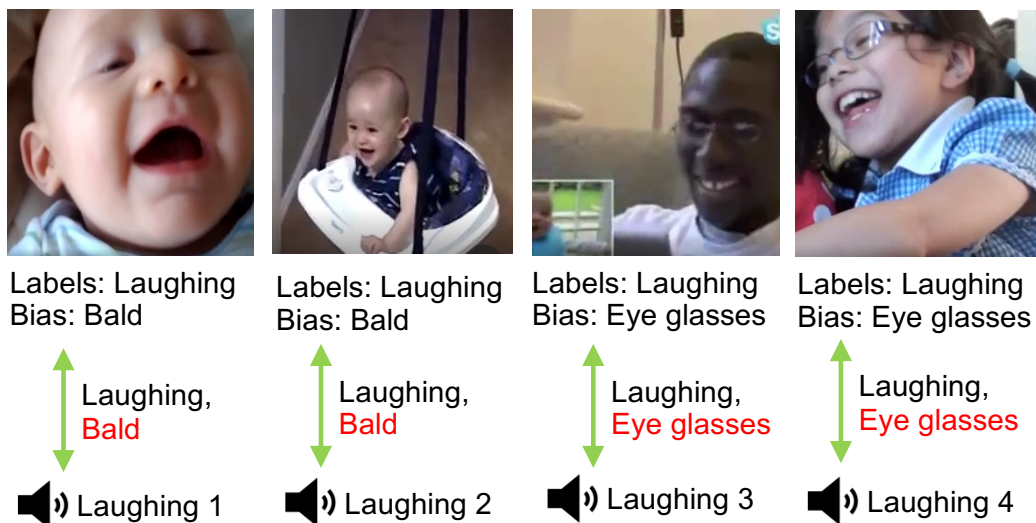


Figure 2: Examples of audio-visual representational biases observed in the VGG-Sound dataset (Chen et al., 2020), e.g., a scene of “laughing” often concurrently observed with baldness or eyeglasses.

expressed with a combination of several instruments, the atmosphere of the music can be expressed only with sound.

Recent text-conditional image manipulation methodologies show promising results (El-Nouby et al., 2019; Jiang et al., 2021; Li et al., 2020; Nam et al., 2018; Xia et al., 2021). Those works use a powerful joint multi-modal embedding space which is trained on a large scale image-text pairs. These text-based prior works provide a powerful tool for image manipulation. In particular, StyleCLIP (Patashnik et al., 2021), one of the best text-guided image manipulation method, considered leveraging the representational power of Contrastive Language-Image Pre-training (CLIP) (Radford et al.) models to produce text-relevant manipulations with given text input. StyleCLIP also maintains high-quality image generation ability using StyleGAN (Jeong et al., 2021) while allowing the style insertion of text semantics into the image. Both \mathcal{W} space and extended latent space $\mathcal{W}+$, the intermediate latent space of StyleGAN, are suitable for manipulation because they disentangle the feature in the representation spaces.

Previous sound-to-visual synthesis works (Chen et al., 2017; Hao et al., 2018; Oh et al., 2019; Qiu and Kataoka, 2018; Wan et al., 2019; Zhu et al., 2021) have attempted to convert sound signals into realistic visual content.

However, these works have a limitation in creating or changing intuitive image styles with various audio domains, as they do not leverage pre-trained joint multi-modal representations like CLIP (Radford et al.). To solve this problem, many recent studies try to match the sound embedding space with the CLIP embedding space (Wu et al., 2021; Lee et al., 2022; Guzhov et al., 2021). In particular, Lee *et al.* (Lee et al., 2022) proposed a method to provide sound-based guidance by matching the sound embedding space with the CLIP embedding space. After matching the multimodal embedding space, the StyleGAN source latent vector moves under CLIP-based guidance by minimizing the cosine distance of the manipulated image embedding and sound embedding. This approach allows the StyleGAN source latent vector in $\mathcal{W}+$ space to obtain a robust direction for semantic cues while being intensity-aware for various sound inputs.

However, the existing sound-guided image manipulation technique (Lee et al., 2022) has a weakness when guiding unintentional information that is not included in the sound inputs. As the sound latent vector does not project correctly into the image embedding space, the image attributes unrelated to the audio are changed. Fig. 3 shows how a sound sample gives audio-irrelevant direction such as eyeglasses or hair color. This is because the previous method (Lee et al., 2022) allows one video sample to have only one strong positive edge between its own audio and frame even if many same class video samples share common properties. Specifically, we find some biases in audio-visual pairs for sound representation learning which result in fatal issue about image manipulation (see Fig. 2). As a result, sampling only one strong positive pairs from one video affects the image manipulation results as shown in Fig. 3.

To tackle these problems, we propose a novel audio-visual weakly paired contrastive learning, which prevents sound representation from learning biases between audio-visual from other videos in minibatches. Our approach enlarges the number of positive weak edges between sound and image which are sampled from other videos. By leveraging weakly connected audio-visual pairs which shares same text description, we take advantage of the fact that videos having the same labels share common visual feature information.

We propose to minimize Kullback–Leibler divergence (Hinton et al., 2015) between visual-text cosine similarity probability distribution and audio-visual cosine similarity probability distribution. Our weakly paired contrastive learning proceeds as follows. First, we sample additional visual data with text queries during the sound representation pre-training step. After that, our



Figure 3: Failure examples of our prior sound-based image manipulation approach (Lee et al., 2022) with sound inputs of giggling and nose blowing. Though it produces visually plausible manipulation results, fine-grained visual details (e.g., hair color) are not properly retained in the manipulated images. For example, see unintentional hair color changes (1st column), hairstyle changes (2nd column), the appearance of glasses (3rd column), and disappearance of beard (4th column).

audio-visual similarity score aims to imitate the visual-text similarity score. Instead of estimating the identity matrix directly, the proposed method equalizes the audio-visual softmax-probability distribution and the visual-text softmax-probability distribution in minibatches. We find that our approach generates more robust images to fine-grained visual details, such as drastic color and hairstyle changes, than the previous method (Lee et al., 2022).

While expanding the scope and scale of the audio-visual relationship, we adopt (Lee et al., 2022)’s sound representation learning process to express the relationship between audio-text and audio-audio as shown in Fig 4 (a). The audio encoder is trained to produce a latent representation aligned with

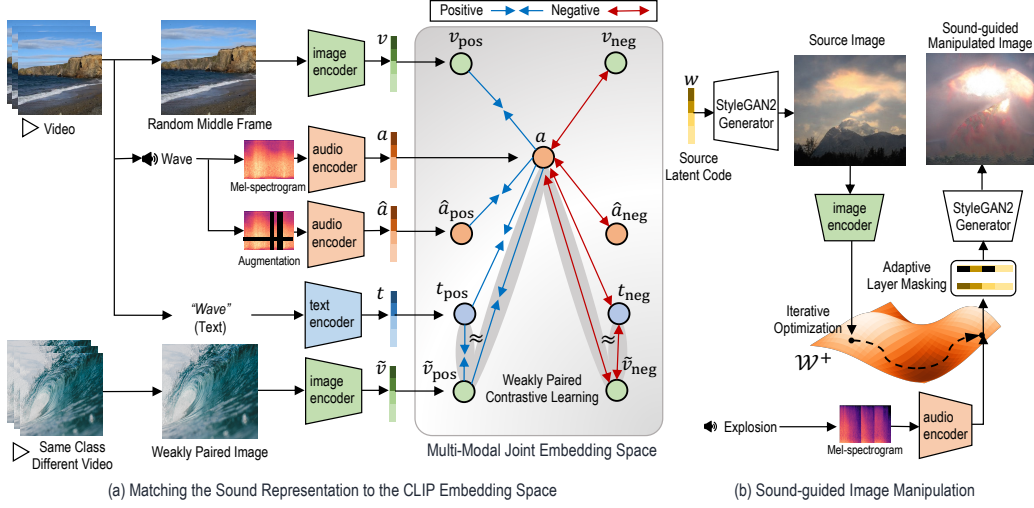


Figure 4: Overview of our proposed sound-guided image manipulation process. Our model mainly consists of two steps: (a) matching the sound representation to the CLIP embedding space and (b) the StyleGAN2-based sound-guided image manipulation step. In (a), we train an audio encoder to produce the matched latent representations from a given video. The latent representations for a positive triplet pair (e.g., audio input: “Wave”, text: “wave”, and corresponding image) are mapped close together, while that of negative pair samples are mapped further away in the (CLIP-based) embedding space (see right). To further enhance the representational power of the joint embedding space, we use a weakly-aligned image-audio pair, i.e., given a text input “Wave”, we randomly sample an image of the same class but from other video clips. We observe that our proposed Weakly Paired Contrastive Learning effectively retains fine-grained visual details, addressing issues raised in Fig. 3. In (b), we use a direct code optimization approach where a source latent code is modified in response to user-provided audio, producing a sound-guided image manipulation result (right).

textual and visual semantics by leveraging the representation power of pre-trained CLIP models. We perform self-supervised learning between audio samples to represent the inherent properties of audio. Self-supervised approach provides different views even in the same class, enabling richer sound intensity-aware image manipulation.

After we pre-train the multi-modal joint embedding space, we use the direct latent code optimization to produce a semantically meaningful image in response to a user-provided sound as shown in Fig 4 (b). StyleGAN’s source latent vector moves in StyleGAN latent space through an iterative process that minimizes the distance between input audio and the generated

image in CLIP space to generate an image that matches the meaning of the audio.

Our experimental results illustrate that the proposed method shows a robust image manipulation performance for the target attribute (see Fig. 1). Furthermore, the sound-based approach supports more diverse and detailed image manipulation results related to audio intensity information than state-of-the-art text-based image manipulation methods. In addition, we demonstrate that our sound representation outperforms existing audio-visual joint representation (Wu et al., 2021; Guzhov et al., 2021; Lee et al., 2022) for image manipulation. We provide a wide range of image manipulation scenarios for real-world applications, such as music style transfer and sound-guided artistic painting manipulation. This manuscript extends the preliminary conference version of this work (Lee et al., 2022). In this present study, our main contributions are listed as follows:

- We propose a novel method for robust sound-guided image manipulation. We obtain a more robust sound representation with audio-visual weakly paired contrastive learning. Furthermore, our approach reduces the fine-grained bias of sound-guided image manipulation results.
- We demonstrate that our method is effective for sound-guided image manipulation solely based on given audio features, including temporal context, tone, and volume. In particular, the visual change of style is closely proportional to the audio intensity.
- We conduct much more investigations and incorporate discussions, such as ablation studies of another sound representation for image manipulation and music style transfer, and we include discussions. We consider the method’s present limitations and potential future directions.

2. Related Work

2.1. Text-Guided Image Manipulation

Text-guided image manipulation approaches have been most commonly explored in the computer vision community. Several studies utilized a GAN-based encoder-decoder architecture to preserve the visual features in the manipulation process driven by descriptive text inputs (Dong et al., 2017; Li et al., 2020; Nam et al., 2018). Recently, StyleCLIP (Patashnik et al., 2021) utilized the latent space of the pre-trained StyleGAN and a CLIP (Radford

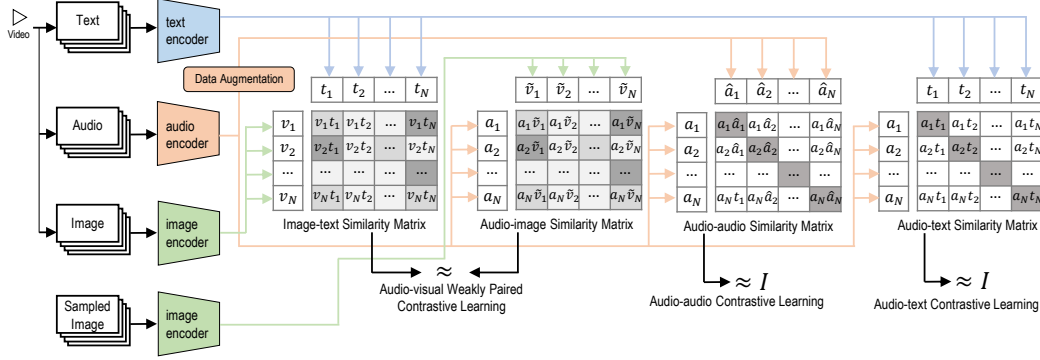


Figure 5: Sound representation learning for image manipulation. We adopt different strategies for representation due to the characteristics of each modality. First, we employ audio-visual weakly paired contrastive learning to obtain similarity scores between images and audio. Instead of performing direct contrastive learning between audio and image embedding, we learn that the audio-visual similarity matrix is equal to the text-visual similarity matrix. Second, we use intensity-aware self-supervised learning to separate different scenes in the embedding space, even if the audio class is the same. Finally, in order to have a joint embedding space between audio and text, we perform contrastive learning between the audio and text embeddings.

et al.)-based text-image joint embedding space to synthesize visually more plausible images (Patashnik et al., 2021). Further, TediGAN (Xia et al., 2021) also explored using a GAN inversion technique based on multi-modal mapping. These approaches generate manipulated images by minimizing the CLIP distance between an input text and an image generated by a StyleGAN-based generator. Similarly, in this work, we also follow this strategy to optimize distances in the CLIP embedding space, but we extend the CLIP space into an image-text-sound joint embedding space, improving its representational power for the image manipulation task. We empirically observe that a sound can convey a more complex semantic context in a scene (e.g., temporal context, music genre, tone, and volume), producing a more visually plausible image manipulation.

2.2. Sound-Guided Image Manipulation

A few works have explored the use of sound inputs in the image manipulation task. However, to our best knowledge, most of these works focused on music instead of sound semantics. Lee *et al.* (Lee et al., 2020) explored a music-to-visual style transfer model with a cross-modal learning strategy.

Jeong *et al.* (Jeong et al., 2021) introduced a neural music visualizer that maps the latent representation of given music to visual embeddings of StyleGAN. Similarly, *TräumerAI* (Jeong et al., 2021) used the latent transfer mapping of music-to-StyleGAN’s embedding to visually express a music input. However, they do not focus on sound semantics (but only on the reaction) in the navigational direction in the StyleGAN latent space, limiting the representational advantages of using sound. For example, *Crossing you in Style* (Lee et al., 2020) used the era of music to define the semantic relationship between sound and images, and thus it can only transfer to a limited number of image styles.

In our prior work (Lee et al., 2022), we studied a novel image manipulation approach with sound input. We first demonstrated that a CLIP-based image-text-sound joint embedding space could be learned via a contrastive learning technique. We also reported that such a joint embedding space can be used to manipulate images driven by various sound inputs (environmental sound, music, etc.). However, this model often fails to retain semantic visual cues that the user-provided sound should not manipulate. This work addresses this issue using audio-visual weakly paired contrastive learning and intensity-aware self-supervised learning techniques.

2.3. Interpreting Latent Space in StyleGAN

The intermediate latent space in pre-trained StyleGAN (Karras et al., 2019) solves the disentanglement issue and allows the generated images to be manipulated meaningfully according to changes in the latent space. Extended latent space $\mathcal{W}+$ allows image manipulation with interpretable controls by a pre-trained GAN generator (Abdal et al., 2019; Karras et al., 2019, 2020). For latent space analysis in audio sequences, *Audio-reactive StyleGAN* (Brouwer, 2020) generates an image at every time step by calculating the magnitude of the audio signal and moving it in the latent space of StyleGAN.

However, the method cannot control the meaning of sound in the latent space. The magnitude of the sound only determines the movement of the latent vector in the StyleGAN latent space. We can manipulate images with the properties of sound using the pre-trained sound representation. Instead of randomly traversing in StyleGAN latent space, our model can give reasonable sound semantics guidance.

2.4. Audio-Visual Representation Learning

Audio-visual representation learning aims to obtain the generalized relationship between different modalities to utilize in several audio-visual downstream tasks, such as audio classification. Most audio-visual representation learning studies (Aytar et al., 2017; Nagrani et al., 2018; Surís et al., 2018) are designed to map distinct modalities into the same embedding space. The correlation between modalities is learned by audio-visual contrastive learning (Chen et al., 2021; Mazumder et al., 2021; Sun et al., 2020). Contrastive learning allows the positive audio-visual pairs to have more similar representations than the representations of other negative pairs.

However, audio-visual representation learning is still challenging for image manipulation because there is as much positive pairs as CLIP (Radford et al.) to learn the correlation between different modalities. CLIP learned the relationship between image and text embedding through the multi-modal self-supervised learning of 400 million image-text pairs and show a zero-shot inference performance comparable to supervised learning in most image-text benchmark datasets.

To overcome this problem, recent studies (Wu et al., 2021; Guzhov et al., 2021; Lee et al., 2022) align the sound embedding space with the CLIP embedding space. In order to predict the same CLIP image representations from the audio stream, Wav2CLIP (Wu et al., 2021) freezes the CLIP image encoder and trains the audio encoder on the audio streams from videos. AudioCLIP (Guzhov et al., 2021) extends modality and scale by using a newly collected trimodal audio-visual-text dataset. Lee *et al.* (Lee et al., 2022) trains an audio encoder from scratch with a audio-text dataset and audio-visual streams from video to obtain a CLIP-based sound latent representation.

In this work, we additionally sample images that are suitable for text queries, and we propose a novel method of weakly paired contrastive learning to form a link between two modalities in sound representation (Lee et al., 2022). Given triplets, we sample other images from text descriptions and learn audio-visual weak supervision. In this way, the bias of the audio-visual representation is reduced, which produces a much more realistic sound-guided image manipulation result.

3. Method

In this work, we present a novel method for sound-guided image manipulation. Our model consists of the following two main steps: (i) the CLIP-based multi-modal latent representation learning (Section 3.1) and (ii) the sound-guided image manipulation (Section 3.2). As we will explain in Section 3.1, we utilize the CLIP-based image-text joint embedding space and extend it to deal with an additional modality, i.e., sound. However, learning such a joint representation from scratch is generally challenging due to the lack of multi-modal datasets. Thus, we instead leverage the pre-trained CLIP model, which optimized a visual-textual joint representation by using contrastive learning. Then, we train our audio encoder to produce representations that align with other modalities. We achieve this by applying contrastive learning losses for audio-text and audio-audio pairs. Additionally, we improve the quality of joint representations by applying a novel weakly paired contrastive loss for audio-image pairs.

In Section 3.2, we explain the sound-guided image manipulation step. Given the multi-modal joint embedding space, we use a direct latent code optimization method to manipulate images guided by an audio input. For example, a facial expression of a portrait photo is manipulated according to sound inputs, e.g., giggling or sobbing, as shown in Fig. 1. Note that our approach may be similar to existing text-guided image manipulation models, including StyleCLIP (Patashnik et al., 2021), but we focus on exploring a novel modality, i.e., sound.

3.1. Multi-Modal Latent Representation Learning

3.1.1. Extending CLIP Embedding Space with Sound

As shown in Fig. 5, we train a set of encoders with three different modalities {audio, text, and image} to produce the matched representations in the embedding space from a video dataset. Specifically, given audio, text, and image inputs from the same video, i.e. x_a , x_t , and x_v , we use three different encoders to obtain a set of d dimensional latent representations, i.e. \mathbf{a} , \mathbf{t} and $\mathbf{v} \in \mathbb{R}^d$, respectively.

To learn these joint latent representations, we use a typical contrastive learning approach that follows the work by Radford *et al.* (Radford et al.): the latent representations for a positive triplet pair are mapped close together in the embedding space while representations of negative pair samples are mapped farther away from each other. Specifically, we use InfoNCE

loss (Alayrac et al., 2020) to map positive audio-text pairs and audio-visual pairs close together in the CLIP-based joint embedding space, and we map negative pairs farther away from each other. Formally, given a minibatch of N audio-text representation pairs $\{\mathbf{a}_i, \mathbf{t}_j\}$ for $i \in \{1, 2, \dots, N\}$, we first compute the (i, j) component of the audio-text similarity matrix $M^{a \rightarrow t} \in \mathbb{R}^{N \times N}$ between the \mathbf{a}_i and \mathbf{t}_j latent representation as:

$$M_{(i,j)}^{a \rightarrow t} = \frac{\exp(\mathbf{a}_i \cdot \mathbf{t}_j / \tau)}{\sum_{k=1}^N \exp(\mathbf{a}_i \cdot \mathbf{t}_k / \tau)}, \quad (1)$$

where $\mathbf{a}_i \cdot \mathbf{t}_j$ represents vector dot product and τ is a temperature hyperparameter. We emphasize that \mathbf{a}_i and \mathbf{t}_j are normalized vectors, so that the magnitude of each vector is 1. This loss function is the log loss of an N -way classifier that seeks to predict the i -th diagonal component $\{\mathbf{a}_i, \mathbf{t}_i\}$ as the true representation pair in the similarity matrix $M^{a \rightarrow t}$.

$$l_i^{(a \rightarrow t)} = -\log(M_{(i,i)}^{a \rightarrow t}). \quad (2)$$

As the loss function is asymmetric, we also define the following text-to-audio similarity matrix $M^{t \rightarrow a}$:

$$M_{(i,j)}^{t \rightarrow a} = \frac{\exp(\mathbf{t}_i \cdot \mathbf{a}_j / \tau)}{\sum_{k=1}^N \exp(\mathbf{t}_i \cdot \mathbf{a}_k / \tau)}. \quad (3)$$

We maximize the diagonal component of the matrix $M^{t \rightarrow a}$ by minimizing the InfoNCE loss as:

$$l_i^{(t \rightarrow a)} = -\log(M_{(i,i)}^{t \rightarrow a}). \quad (4)$$

Concretely, we minimize the following loss function \mathcal{L}_{nce} as a sum of the two losses $l_i^{(a \rightarrow t)}$ and $l_i^{(t \rightarrow a)}$ for all positive audio-text representation pairs in each minibatch of size N :

$$\mathcal{L}_{\text{nce}}^{(a \leftrightarrow t)} = \frac{1}{N} \sum_{i=1}^N (l_i^{(a \rightarrow t)} + l_i^{(t \rightarrow a)}). \quad (5)$$

At the same time, we also calculate the diagonal component of the audio-visual similarity matrix $M^{a \rightarrow v}$ in the same way as in Eq. 4. We determine the following $\mathcal{L}_{\text{nce}}^{(a \leftrightarrow v)}$:

$$\mathcal{L}_{\text{nce}}^{(a \leftrightarrow v)} = \frac{1}{N} \sum_{i=1}^N (l_i^{(a \rightarrow v)} + l_i^{(v \rightarrow a)}), \quad (6)$$

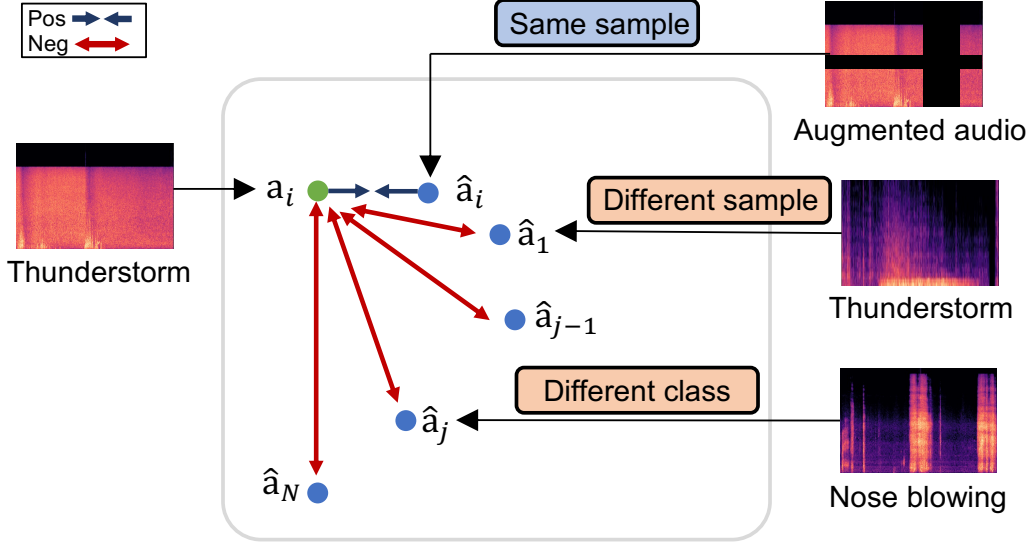


Figure 6: Intensity-aware self-supervised representation learning for audio inputs. Augmented audio of the same scene is considered a positive pair, and other samples of the different scenes are considered a negative pair (even if they are the same class).

where $l_i^{(v \rightarrow a)}$ and $l_i^{(a \rightarrow v)}$ denote the negative log-likelihood of $M_{(i,i)}^{v \rightarrow a}$ and $M_{(i,i)}^{a \rightarrow v}$ in each minibatch of size N .

Data Augmentations. We further apply a data augmentation strategy to improve the quality of representations and overcome the lack of large-scale audio-text multi-modal datasets. For the text inputs, we augment the text data by (i) replacing words with synonyms, (ii) applying a random permutation of words, and (iii) inserting random words. Note that, for (i), we find synonyms of the given word from WordNet (Fellbaum, 2010) and insert the synonym at random into the given text input. For example, we augment the original texts “rowboat, canoe, kayak rowing” to produce the new text “row canoe, kayak quarrel rowboat.”

3.1.2. Self-Supervised Learning towards Intensity-Aware Sound Representation

Unlike text inputs, audio inputs can provide more diverse and fine-grained semantic cues. For example, a “thunderstorm” can be expressed in various ways: e.g., weak or severe thunderstorms, thunderstorms with heavy rain, winds, hail, or no precipitation. It may be challenging for texts to verbalize

and differentiate these subtle differences, but it does not apply to sound. Thus, we apply self-supervised learning techniques between audio samples (e.g., weak vs. severe thunderstorms) to obtain an intensity-aware sound representation.

Existing self-supervised learning approaches depend on a contrastive loss that makes representations of the same-class different views close in the embedding space. In contrast, different-class views are pushed farther away. To this end, we regard the audio input of different augmented views as positive and different audio samples as negative. Note that we obtain a latent representation $\hat{\mathbf{a}} \in \mathbb{R}^d$ from an augmented audio input \hat{x}_a , which has been shown to be useful for improving the quality of the latent representation, as this is a common practice in self-supervised representation learning.

Formally, we define a similarity matrix $M^{a \rightarrow \hat{a}} \in \mathbb{R}^{N \times N}$ from embeddings obtained from audio and augmented audio. The (i, j) component of the audio-audio similarity matrix $M^{a \rightarrow \hat{a}}$ is as follows:

$$M_{(i,j)}^{a \rightarrow \hat{a}} = \frac{\exp(\mathbf{a}_i \cdot \hat{\mathbf{a}}_j / \tau)}{\sum_{k=1}^N \exp(\mathbf{a}_i \cdot \hat{\mathbf{a}}_k / \tau)}. \quad (7)$$

The intensity-aware representation is obtained by minimizing the InfoNCE loss of $M_{(i,i)}^{a \rightarrow \hat{a}}$. We apply this technique to improve the quality of audio representations by minimizing $\mathcal{L}_{\text{self}}^{(a \leftrightarrow \hat{a})}$ as follows:

$$\mathcal{L}_{\text{self}}^{(a \leftrightarrow \hat{a})} = \frac{1}{N} \sum_{i=1}^N (l_i^{(a \rightarrow \hat{a})} + l_i^{(\hat{a} \rightarrow a)}), \quad (8)$$

where $l_i^{(a \rightarrow \hat{a})}$ and $l_i^{(\hat{a} \rightarrow a)}$ are defined in a similar way, as in Eq. 2 and 4. This loss function is useful for learning subtle differences over sound inputs, as it must maximize the mutual information between two different views of the same inputs but to minimize the mutual information between two views of the different inputs. For example, as shown in Fig. 6, an audio sample a_i forms a negative pair with \hat{a}_j for $i \neq j$, which induces a diffusive effect in the embedding space. For audio inputs, we apply the SpecAugment (Park et al., 2019a), which visually augments Mel-spectrogram acoustic features by warping the features and masking blocks of frequency channels.

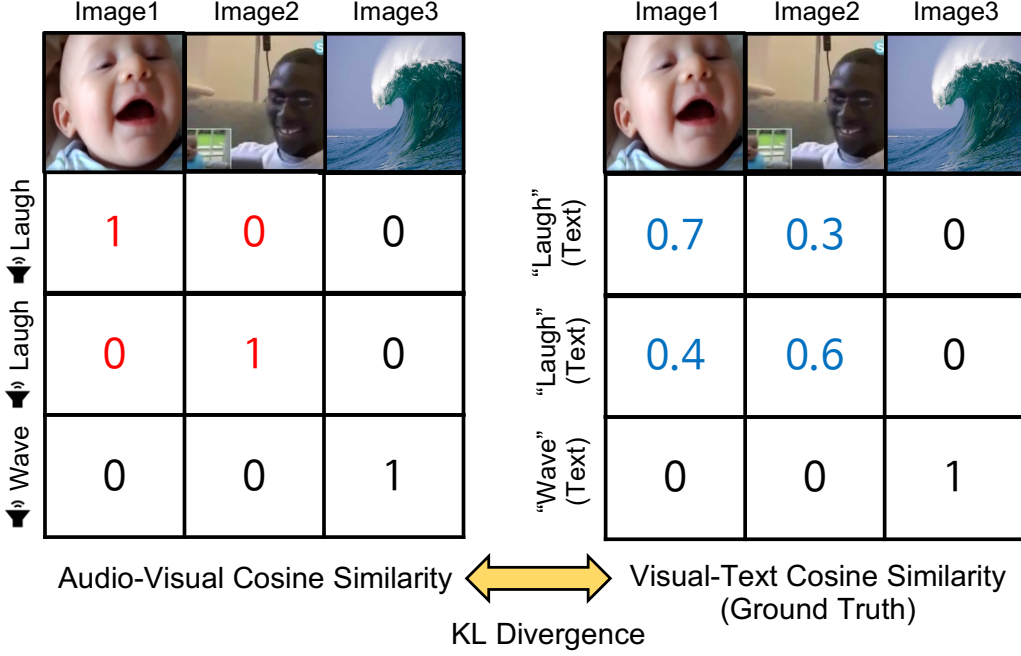


Figure 7: A graphical example of weakly paired contrastive learning for audio-visual joint embedding space. We minimize KL divergence between (left) audio-visual cosine similarity distribution and (right) visual-text cosine similarity distribution instead of simply enforcing it into the identity matrix, reducing the negative effect of audio-visual representational biases (see examples in Fig. 2).

3.1.3. Weakly Paired Contrastive Learning for Audio-Visual Joint Embedding Space

In addition, we propose weakly paired contrastive learning for audio-visual data to ensure that the representations of the input audio and different images with the same text labels are similar. The conventional contrastive loss only uses positive pairs which have audio and image from the same video. Therefore, audio and image from different videos are considered to be negative pairs although they may share visual features. Our new loss encourages these pairs to be closed to each other if they have the same text labels.

We first sample the new images \tilde{x}_v , which have the same text labels with given audio x_a . Then, we use the image encoder to obtain a new sampled image embedding $\tilde{v} \in \mathbb{R}^d$. We apply the knowledge distillation technique (Hin-

ton et al., 2015) to transfer the generalization ability of the CLIP to sound representation. The learning scheme allows the similarity matrix between audio-visual and visual-text to be equal in mini-batch (see Fig. 7).

The link knowledge between the audio and the new sampled image is transferred from CLIP to the diagonal component of a new audio-visual similarity matrix. After sampling a new triplet by sampling sound, text and new images of size N in a mini-batch, we define the audio-visual similarity matrix from the \mathbf{a}_i and $\tilde{\mathbf{v}}_j$ embeddings as:

$$M_{(i,j)}^{a \rightarrow \tilde{v}} = \frac{\exp(\mathbf{a}_i \cdot \tilde{\mathbf{v}}_j / \tau)}{\sum_{k=1}^N \exp(\mathbf{a}_i \cdot \tilde{\mathbf{v}}_k / \tau)}, \quad (9)$$

where τ is a temperature hyperparameter to refine the softmax-probability distributions, as was done in (Hinton et al., 2015). Specifically, we compute the pseudo cosine similarity matrix $M^{t \rightarrow \tilde{v}}$ between the latent representations of text and those of images. The audio-visual link knowledge is transferred from the fixed CLIP to $M_{(i,i)}^{a \rightarrow \tilde{v}}$. Then, we match the two types of logits in the sound-image and text-image similarity matrix instead of using hard labels such as 1 for positive pairs. We minimize the following Kullback-Leibler divergence (Hinton et al., 2015) for the i -th diagonal component:

$$l_i^{KD} = \text{KL}(M_{(i,i)}^{a \rightarrow \tilde{v}}, M_{(i,i)}^{t \rightarrow \tilde{v}}), \quad (10)$$

where KL denotes the Kullback-Leiber divergence. The KL can be rewritten as:

$$\text{KL}(M_{(i,i)}^{a \rightarrow \tilde{v}}, M_{(i,i)}^{t \rightarrow \tilde{v}}) = -M_{(i,i)}^{t \rightarrow \tilde{v}} \log M_{(i,i)}^{a \rightarrow \tilde{v}} \quad (11)$$

The KL divergence loss \mathcal{L}_{KL} for N batches is as follows:

$$\mathcal{L}_{\text{KL}}^{(a \leftrightarrow \tilde{v})} = \frac{1}{N} \sum_{i=1}^N l_i^{KL}. \quad (12)$$

To summarize, we minimize the following loss function $\mathcal{L}_{\text{total}}$:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{nce}}^{(a \leftrightarrow t)} + \mathcal{L}_{\text{nce}}^{(a \leftrightarrow v)} + \mathcal{L}_{\text{self}}^{(a \leftrightarrow \hat{a})} + \mathcal{L}_{\text{KL}}^{(a \leftrightarrow \tilde{v})}. \quad (13)$$

3.2. Sound-Guided Image Manipulation

After learning the multi-modal joint embedding space by minimizing Eq. 13, we use a direct latent code optimization method, similar to Style-CLIP (Patashnik et al., 2021), to manipulate the given image. Our method

optimizes for each image instance manipulated at the inference stage. Our model minimizes the distance between a given source latent code and an audio-driven latent code in the learned joint embedding space to produce sound-guided manipulated images. Moreover, we propose an *Adaptive Layer Masking* technique, which adaptively manipulates the latent code.

3.2.1. Direct Latent Code Optimization

Given user-provided audio input x_a , we employ the direct latent code optimization for sound-guided image manipulation by solving the following optimization problem:

$$\arg \min_{w_a \in \mathcal{W}} \mathcal{L}_{\text{cos}}(G(w_a), x_a) + \lambda_{\text{reg}} \mathcal{L}_{\text{reg}} + \lambda_{\text{ID}} \mathcal{L}_{\text{ID}}(w_a), \quad (14)$$

where a given source latent vector $w_s \in \mathcal{W}^{L \times D}$ (the intermediate latent space in StyleGAN), audio-driven latent code $w_a \in \mathcal{W}^{L \times D}$ is optimized. L denotes the number of style layers at different levels, and D denotes the dimension. \mathcal{L}_{reg} is regularization loss. \mathcal{L}_{ID} and G are the identity loss and StyleGAN-based generator, respectively. By minimizing the hinge loss \mathcal{L}_{cos} , we find the sound-guided latent vector w_a which generates an sound-guided image in the StyleGAN latent space. λ_{reg} and λ_{ID} are hyperparameters that control the strengths of the similarity loss term and the identity loss function \mathcal{L}_{ID} . High values of λ_{reg} and λ_{ID} lead to maintenance of the content of the source image, whereas low values do not. The source latent code w_s means the randomly generated latent code from G or the latent code obtained from the existing input image through GAN inversion (Richardson et al., 2021; Tov et al., 2021).

With such an optimization scheme, the hinge loss \mathcal{L}_{cos} is calculated as:

$$\mathcal{L}_{\text{cos}}(G(w_a), x_a) = \max(d_{\text{cos}}(f_v(G(w_s)), f_a(x_a)) - d_{\text{cos}}(f_v(G(w_a)), f_a(x_a)) + 1, 0). \quad (15)$$

In the CLIP embedding space, the distance from the source latent vector w_s is increased and closer to the audio-driven latent vector w_a . We minimize cosine distance between the latent representation of the manipulated image $G(w_a)$ and audio latent representation $f_a(x_a)$. f_a and f_v denote the pre-trained audio and visual encoders.

3.2.2. Adaptive Layer Masking

Generally, L_2 regularization effectively keeps the image generated from the moved latent code from being different from the original (Patashnik et al.,

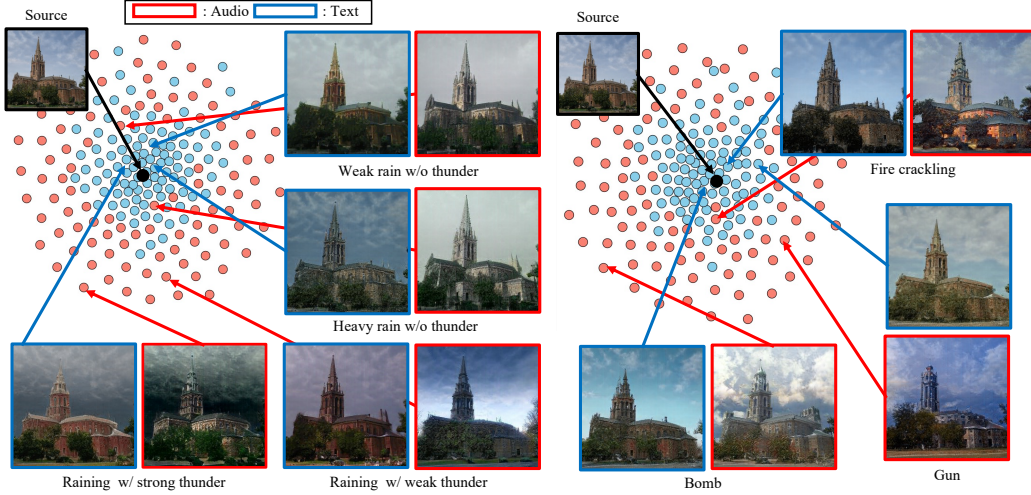


Figure 8: Comparison of manipulation results between our approach (red) and the existing text-driven manipulation approach, StyleCLIP (Patashnik et al., 2021) (blue). Unlike the text-based approach, we can produce more varied manipulation results, depending on the intensity of the sound. These are the results of image manipulation with weak rain, heavy rain, raining with weak thunder, raining with strong thunder, fire crackling, gun sounds, and bomb sounds from the left column to the right. We also report audio and text-guided manipulation directions with t-SNE visualization (van der Maaten and Hinton, 2008) (red: audio, blue: text).

2021). In the StyleGAN latent space, we maintain the image similarity by minimizing the l_2 - distance between w_a and w_s as follows:

$$\mathcal{L}_{\text{reg}} = \|w_a - w_s\|_2. \quad (16)$$

However, StyleGAN’s latent code has different properties at each layer, so different weights should be applied to each layer if the user-provided attribute changes. We control style changes with adaptive layer masking. Layerwise masking keeps compact content information within style latent code as follows:

$$\mathcal{L}_{\text{reg}} = \frac{1}{L} \sum_{i=1}^L g_i \cdot \|(w_i^a - w_i^s)\|_2. \quad (17)$$

In StyleGAN2 (Karras et al., 2020), the latent code is represented as $w \in \mathcal{W}^{L \times D}$, where L is the number of layers and D is the latent code’s dimension. We declare a parameter vector $g \in \mathbb{R}^L$. g is a trainable vector that masks adaptively the specific style layer. g ranks which of the L layers is

regulated by the softmax function. In the latent optimization step, g and w are multiplied according to the number of layers. g is iteratively updated, which adaptively manipulates the latent code.

3.2.3. Identity Loss

The similarity to the input image is also controlled by the identity loss function \mathcal{L}_{ID} , which is defined as:

$$\mathcal{L}_{\text{ID}}(w_a) = 1 - \langle R(G(w_s)), R(G(w_a)) \rangle. \quad (18)$$

where R is the pre-trained ArcFace (Deng et al., 2019) model for face recognition. Thus, this loss function minimizes the cosine distance $\langle R(G(w_s)), R(G(w_a)) \rangle$ between its arguments in the latent space of the ArcFace network. This allows for manipulating facial expressions without changing the personal identity. Note that we disable the identity loss by setting $\lambda_{\text{ID}} = 0$ for all other non-human image manipulations.

4. Experiments

4.1. Implementation and Training Details

To learn the multi-modal joint embedding space, we use the pre-trained CLIP (Radford et al.)-based image and text encoders, and we train a Swin transformer (Liu et al., 2021)-based audio encoder, which outputs a latent vector of the length (i.e., 512) that is the same as the output of the image and text encoders. Note that we convert audio inputs into Mel-spectrogram acoustic features by following a common practice in various audio processing tasks. For the image manipulation step, we leverage the pre-trained StyleGAN2 (Karras et al., 2020) generator. Note that we adjust the size of the latent code depending on the resolution of the output image, i.e., 18×512 latent code for images of size 1024×1024 and 14×512 for 256×256 . To implement such a generator, we utilize the official PyTorch implementation from StyleGAN2-ADA.

We train our model for 50 epochs using the Adam optimizer with the cosine cyclic learning rate scheduler (Smith, 2017). We set the learning rate to 10^{-3} with the momentum of 0.9, and we set the weight decay to 10^{-4} . The batch size is set to 384. For audio augmentation, we use SpecAugment (Park

<https://github.com/NVlabs/stylegan2-ada-pytorch>




Figure 9: Given the (a) input image, we compare the face image manipulation results between (b-c) text-driven image manipulation approaches (i.e. TediGAN (Xia et al., 2021) and StyleCLIP (Patashnik et al., 2021)), (d) Lee *et al.* (previous sound-guided image manipulation approach) and (e) ours. The attributes for driving such manipulations include baby crying, coughing, giggling, screaming, sobbing and nose blowing.

et al., 2019a) with a frequency mask ratio of 0.15 and a time masking ratio of 0.3. For direct latent code optimization, λ_{reg} and λ_{ID} in Eq. (14) are set to 0.008 and 0.004 for the FFHQ dataset; and 0.002 and 0 for the LSUN dataset.

4.2. Datasets

To learn a multi-modal embedding space, we use the audio-text pair from a publicly available video dataset called AudioSet (Gemmeke et al., 2017), which contains 2,084,320 sound clips (in 10-second segments) extracted from

YouTube videos and provides annotations of 632 audio classes. Because some of the associated URLs on YouTube were missing, we use 17,153 audio clips (out of 20,371 in the balanced subsets) and 1,617,939 clips (out of 2,041,789 in the unbalanced subsets) for training.

We also utilize VGG-Sound (Chen et al., 2020), which similarly provides over 200k human-labeled video clips (in 10-second segments) and the 310 corresponding audio classes. We finally use 182,342 videos for training due to missing clips. Note that a pair of images from the middle frame of each video clip and its corresponding class label is used as an image-text input pair for training our model.

Further, we train our StyleGAN2-based generator with the following high-resolution image datasets: Flickr-Faces-HQ (FFHQ) (Karras et al., 2019), Large-Scale Scene Understanding (LSUN) Challenge (Yu et al., 2015) dataset, WikiArt (Saleh and Elgammal, 2016), or Landscapes High-Quality (LHQ) (Skorokhodov et al., 2021). Flickr-Faces-HQ (FFHQ) (Karras et al., 2019) contains the 70,000 high-quality human face images with a resolution of 1024×1024 , and the Large-scale Scene Understanding (LSUN) Challenge (Yu et al., 2015) contains church images with a resolution of 256×256 and car images with a resolution of 384×512 . WikiArt (Saleh and Elgammal, 2016) contains images of the paintings with a resolution of 1024×1024 drawn by 195 different artists. Lastly, the Landscapes High-Quality (LHQ) (Skorokhodov et al., 2021) contains nature landscape images of the resolution 256×256 .

4.3. Qualitative Analysis

4.3.1. Comparison Against Text-Guided Manipulation

We first compare our sound-guided image manipulation results with existing text-guided manipulation approaches, including TediGAN (Xia et al., 2021) and StyleCLIP (Patashnik et al., 2021). We also compare our approach with the sound-guided image manipulation approach from our prior work (Lee et al., 2022). As can be observed in Fig. 9, image manipulation driven by sounds produces better manipulation quality (compare (b) and (c) vs. (d) and (e)). Furthermore, unlike text-guided methods, the audio-guided approach achieves natural image style transfer and is capable of representing multiple labels. For example, TediGAN emphasizes crying, whereas StyleCLIP focuses on the baby when the text “baby crying” is given. On the contrary, our proposed method can simultaneously handle “baby” and “crying”.

Moreover, our proposed method produces more visually plausible manipulation results than does our prior work (Lee et al., 2022), which often produced color shifts (e.g., skin tones).

4.3.2. *Audio Intensity-Aware Image Manipulation*

We demonstrate that each audio sample has its own context, which makes the guidance in StyleGAN latent space richer than the text-based image manipulation method (see Fig. 8). Specifically, our sound representations provide guidance that correlates with the audio intensity in the StyleGAN latent space. If the magnitude of *Thunder* is altered or if a specific attribute like *Rain* is added to the audio, the style variation for the image is better correlated to the audio intensity than text-guided image manipulation. Likewise, *Fire crackling*, *Gunfire*, and *Bomb* sounds contain similar audio characteristics where something explodes, but our image manipulation results show the changes from the audio intensity. In particular, we would like to emphasize that no such explicit representation, such as *Weak thunder* in the text descriptions, is used for pre-training the multi-modal embedding space.

We visualize the direction vector with t-SNE (van der Maaten and Hinton, 2008) (see Fig. 8). By subtracting the vectors of the latent code guided by each modality and the source latent code, we show the distribution of the direction of manipulation. We select the attributes in VGG-Sound (Chen et al., 2020) and randomly manipulate the audio and text prompts. Although we randomly sample the audio and text having the similar semantics, the sound-guided latent code shows a more significant transition than the text-guided latent code.

4.3.3. *Comparison with Existing State-of-the-Art Sound-Based Style Transfer Models*

We compare our sound-guided image manipulation model with existing sound-based style-transfer models, including TräumerAI (Jeong et al., 2021) and Crossing you in Style (Lee et al., 2020). Fig. 10 compares image manipulation results in response to given audio inputs, e.g., audio of fire crackling and raining. We observe that our model produces a better quality of manipulated images. Existing models fail to capture semantic cues from the given audio input (See 2nd and 3rd columns). In contrast, our model manipulates images to be consistent with the semantic cues from the input sound.




Figure 10: Comparison of sound-guided manipulation results. Given fire crackling (top) and raining (bottom) audio inputs, we manipulate the input image with TräumerAI (Jeong et al., 2021), Crossing you in style (Lee et al., 2020), and our method.

4.3.4. Comparison with Existing Multi-modal Embedders

Similar to our image-text-sound joint embedder, several works (Wu et al., 2021; Guzhov et al., 2021; Lee et al., 2022) have attempted to extend the CLIP model to jointly represent audio, text, and images. Here, we also compare those works with our learned embedding space to analyze their effectiveness for the image manipulation task. As shown in Fig. 11, we qualitatively compare image manipulation results based on different multi-modal joint embedding spaces learned by Wav2CLIP (Wu et al., 2021), AudioCLIP (Guzhov et al., 2021), Lee *et al.* (Lee et al., 2022), and our joint embedder. Given different sound inputs (e.g., bird song, bomb, weak thunder, etc.), our learned embedding space generally produces the best results compared with the other embedders (see 2nd, 3rd, and 4th rows vs. 5th row). This may confirm that our multi-modal joint embedding space is more effective for the image manipulation task than are existing approaches.

4.3.5. Effect of Audio-Visual Weakly Paired Contrastive Learning

Weakly paired audio-visual contrastive learning makes the manipulation results robust. As demonstrated in Fig. 12, we demonstrate that using weak supervision is helpful for reasonable sound-guided image manipulation. For




Figure 11: Image manipulation according to sound embedding space. We manipulate the input image with Wav2CLIP (Wu et al., 2021), AudioCLIP (Guzhov et al., 2021), Lee *et al.* (Lee et al., 2022) and our method. The sound attributes for image manipulation include *Bird song*, *Gun fire*, *Bomb*, *Fire crackling*, *Chainsaw*, *Weak Thunder*, *Strong Thunder*, and *Wave* sounds from left to right. Our model produces the most realistic sound-guided image manipulation results.

example, manipulated images maintain the audio-irrelevant attributes like glasses, hair styles and beards. Moreover, it is more robust for manipulating target attributes. It shows the manipulation results that satisfy the keywords "baby" and "laughing" at the same time.

4.3.6. Effect of Adaptive Layer Masking

In StyleGAN (Jeong et al., 2021), it is necessary to adaptively regularize the style layer, because each layer of latent code has different style attributes. Each layer of latent code multiplies the trainable parameter that controls the diversity during regularization. The ablation study shows a qualitative comparison of the mechanism for applying adaptive layer masking to the style layer, as illustrated in Fig. 13. The adaptive masking rectifies the direction by changing the latent code based on the semantic cue. When the gate function is applied, sound-guided image manipulation is semantically reasonable. For




Figure 12: Ablation study of weakly paired contrastive learning for audio-visual joint embedding space. The leftmost is the input source, the middle is the manipulated result without KL divergence loss and the rightmost is the manipulation result after applying the KL divergence loss. In the middle, fine features such as hair loss change, whereas our method with KL divergence loss preserves the characteristics of the input source.

example, a thunderstorm is a mixture of thunder and rain sound. Although thunder and lightning are not seen in the second row, lightning and rain appear in the last row.

Table 1: Comparison of the quality of audio representations between ours and alternatives. We report the classification accuracy (top-1 in %) of a linear classifier on the ESC-50 (Piczak, 2015) and the Urban sound 8k (Salamon et al., 2014) datasets as well as their zero-shot inference results. *Abbr. S: supervised setting.*

Model	S	Zero-shot	Dataset	
			ESC-50 (\uparrow)	Urban sound 8k (\uparrow)
ResNet50 (Hershey et al., 2017)	✓	-	66.8 %	71.3 %
Lee <i>et al.</i> (Lee et al., 2022) (LR)	-	-	72.2 %	66.8 %
Ours (LR)	-	-	89.5 %	72.4 %
Wav2clip (Wu et al., 2021)	-	✓	41.4 %	40.4 %
AudioCLIP (Guzhov et al., 2021)	-	✓	69.4 %	68.8 %
Lee <i>et al.</i> (Lee et al., 2022)	-	✓	57.8 %	45.7 %
Ours w/o $\mathcal{L}_{\text{KL}}^{(a \leftrightarrow \tilde{v})}$	-	✓	63.6 %	53.2 %
Ours	-	✓	74.4 %	58.6 %

4.4. Quantitative Analysis

4.4.1. Zero-Shot Transfer

We compare our model to the supervised method and the existing zero-shot audio classification method (see Table 1). First, we compare audio embeddings trained by the supervised methods such as logistic regression and ResNet50 (Hershey et al., 2017) supervised by random initialization of weights as a baseline model. Our audio encoder shows better classification performance than the baseline. Secondly, we compare the zero-shot audio classification accuracy with Wav2CLIP (Wu et al., 2021) and AudioCLIP (Guzhov et al., 2021). Our proposed loss function learns three modalities in the CLIP embedding space and obtains a richer audio representation through contrastive loss. However, Wav2CLIP only understands the relationship between the audio and the visual context. Furthermore, Wav2CLIP learns from an audio-visual having a dataset with a relatively small data scale to distillate the CLIP representation to the audio encoder (the VGG-Sound dataset: 200K). AudioCLIP uses a dataset scaled up from contrastive learning between modalities from the audio, image, and text given as a triplet to obtain a more powerful sound representation. The AudioCLIP’s audio encoder was pre-trained on a large-scale audio dataset, namely AudioSet (Gemmeke et al., 2017), which consists of over 2M sound clips.

As a result, Table 1 also illustrates that the power of sound representation improves when the scale of positive audio-visual data is increased with weakly

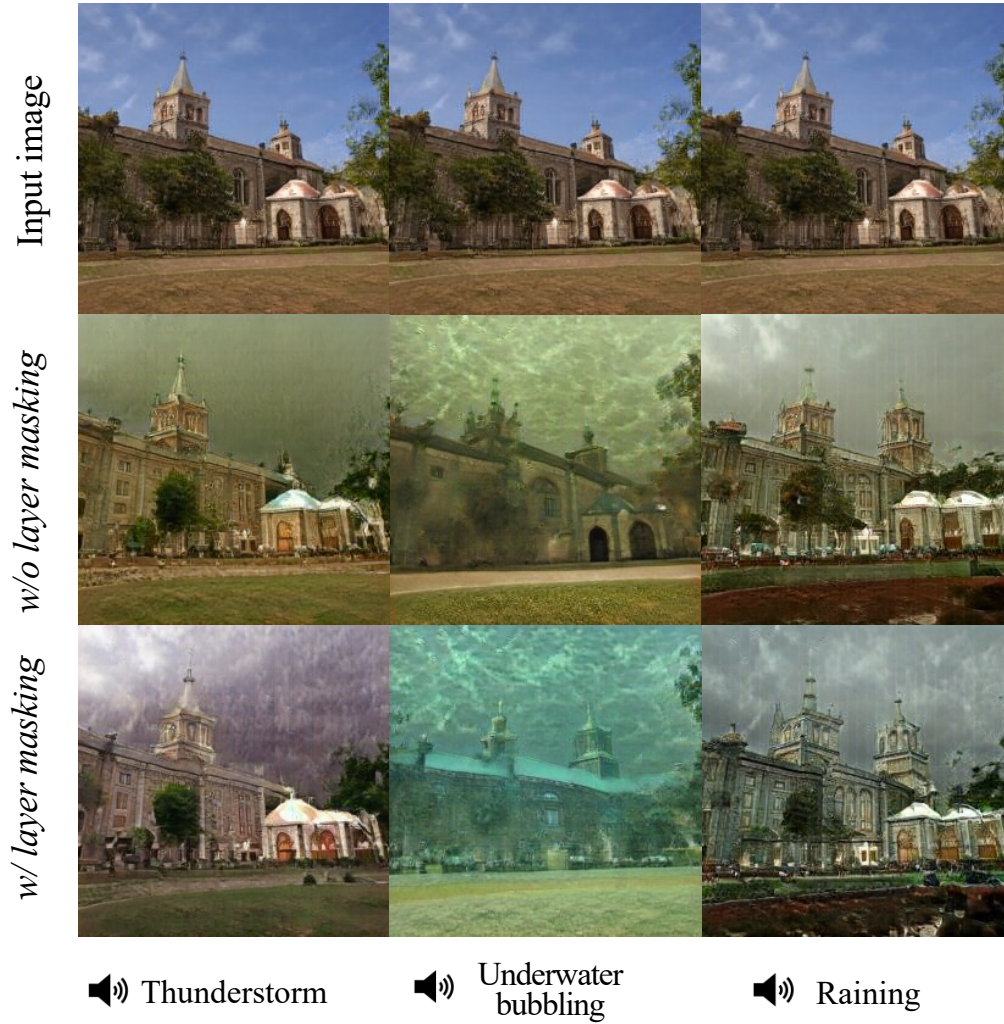


Figure 13: Ablation study of adaptive layer masking. The first row shows the input image, the second row shows the manipulation results without adaptive layer masking, and the third row shows the sound-guided image manipulation results after applying the adaptive layer masking.

paired contrastive learning than Lee *et al.* (Lee et al., 2022)’s work. Moreover, our learned embedding space is more suitable for the image manipulation task.

Table 2: Cosine similarity comparison between text-guided latent code w_t and sound-guided latent code w_a . Note that w_t is the latent code from StyleCLIP (Patashnik et al., 2021)’s text-driven latent optimization. w_s indicates the source latent code.

Latent code	Metric	Attribute											
		Giggling	Sobbing	Nose blowing	Fire crackling	Wind noise	Underwater bubbling	Explosion	Thunderstorm	Gun	Raining	Bird song	Average
(w_s, w_a)	Mean (\downarrow)	0.99315	0.99354	0.99201	0.97400	0.97333	0.97409	0.97230	0.97014	0.97645	0.97588	0.97613	0.97918
	Std (\uparrow)	0.00154	0.00168	0.00267	0.00868	0.00881	0.00929	0.01042	0.00984	0.00755	0.00898	0.00837	0.00708
(w_s, w_t)	Mean (\downarrow)	0.99866	0.99849	0.99779	0.99260	0.99095	0.99416	0.99437	0.99136	0.99208	0.99332	0.99238	0.99420
	Std (\uparrow)	0.00054	0.00034	0.00064	0.00301	0.00359	0.00246	0.00240	0.00369	0.00222	0.00259	0.00239	0.00217
(w_a, w_t)	Mean (\downarrow)	0.99554	0.99510	0.99448	0.97519	0.97280	0.97082	0.9687	0.96926	0.97716	0.97093	0.97488	0.97862
	Std (\uparrow)	0.00158	0.00170	0.00213	0.00937	0.01026	0.01035	0.01160	0.01045	0.00813	0.01023	0.00912	0.00772




Figure 14: Examples of sound-guided animal face image manipulation results.

4.4.2. Semantic Accuracy of Manipulation

We quantitatively analyze the effectiveness of our proposed audio-driven image manipulation approach. First, we measure performance on the semantic-level classification task. Given the audio embeddings from our pre-trained audio encoder, we train a linear classifier to recognize eight semantic labels, including giggling, sobbing, nose-blowing, fire crackling, wind noise, underwater bubbling, explosion, thunderstorm, gun, raining, and bird song. We use StyleGAN2 (Karras et al., 2020) weights pre-trained from the FFHQ (Karras et al., 2019) dataset for the guiding with giggling, sobbing, and nose-blowing attributes to compare the semantic-level classification accuracy between text and audio. Also, for the guiding with fire crackling, wind noise, underwater bubbling, explosion, and thunderstorm attributes, the weights of StyleGAN2 pre-trained on the LSUN (church) (Yu et al., 2015) dataset are used. As shown in Fig. 17 (a), we generally outperform the existing text-driven manipulation approach with better semantically-rich latent representation.

4.4.3. Distribution of Manipulation Direction

We can see how much the latent code has changed by the cosine similarity between the source latent code and the manipulated latent code. We compare the cosine similarity between text-guided and sound-guided latent representations. We evaluate the mean and variance of the cosine similar-

ity between w_s , a source latent code, w_a , an audio-driven latent code, and w_t , a text-driven latent code. The latent representations generally exhibit a high-level characteristic of the content (see Table 2.). In the latent space of StyleGAN2, the sound-guided latent code moves more from the source latent code than does the text-guided latent code, and the image generated from the sound-guided latent code is more diverse and dramatic than for the text-guided method.

4.5. Human Evaluations

We recruited fifty participants from Amazon Mechanical Turk (AMT) to evaluate the effectiveness of our proposed method from a human perspective. We generate a set of samples of manipulated images from 12 different sound sources: e.g., laughing, sobbing, nose-blowing, wind noise, underwater bubbling, explosion, and thunderstorms. Source images are sampled from FFHQ (Karras et al., 2019), LSUN (Yu et al., 2015), and LHQ (Skorokhodov et al., 2021) datasets. In our human study, participants need to complete the following two questionnaires: (i) Naturalness- *Which image manipulation result better expresses the target attribute?* and (ii) Perceptual Realism- *How realistic are the images below?* Note that we use the Likert scale ranging from 1 (low realistic) to 5 (high realistic) for perceptual realism, and we compare our manipulation results with StyleCLIP (Patashnik et al., 2021) and work by Lee *et al.* (Lee et al., 2022). As we summarized in Fig. 17 (b) and (c), our model clearly outperforms other state-of-the-art approaches (StyleCLIP and Lee *et al.*) in both human evaluations. A large portion of participants (58.4%) judge our generated images to be better manipulated than the images of other approaches, and our method obtains a score of 4.333 on average in terms of perceptual realism. This may confirm our model generates more perceptually realistic images than the other models.

4.6. Sound and Text-Guided Style Mixing

Creating new imaginary visual content by combining text and audio information is promising. Our method ensures that audio, text, and image share the same embedding space. We interpolated text and sound-guided latent code to demonstrate that multi-modal embedding lies in the same latent space (see Fig. 15). We obtain the interpolated latent style vector as follows:

$$w = (1 - \alpha) \cdot w_a + \alpha \cdot w_t. \quad (19)$$

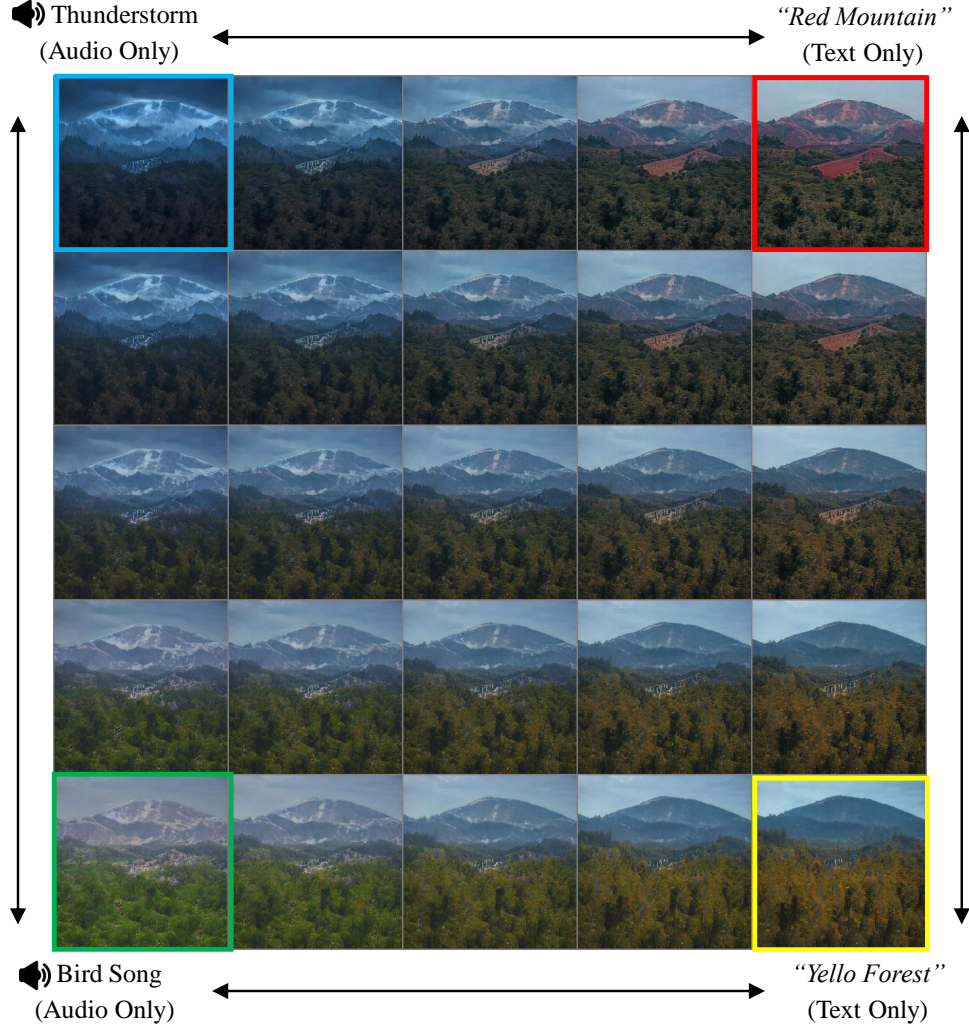


Figure 15: Manipulation results of interpolation between interpolation between text and sound-guided StyleGAN latent vectors. Even if the source latent vector is guided by a different modality, we can control the style changes linearly.

w denotes the interpolated latent vector which is the weighted sum of w_t , the text-guided latent vector, and w_a denotes the sound-guided latent vector. The results generated from the interpolated latent vector contains the continuous intermediate meaning of the two modalities.

We propose a multi-modal manipulation of audio and text based on the style mixing technique of StyleGAN. Different layers of w latent code in Style-

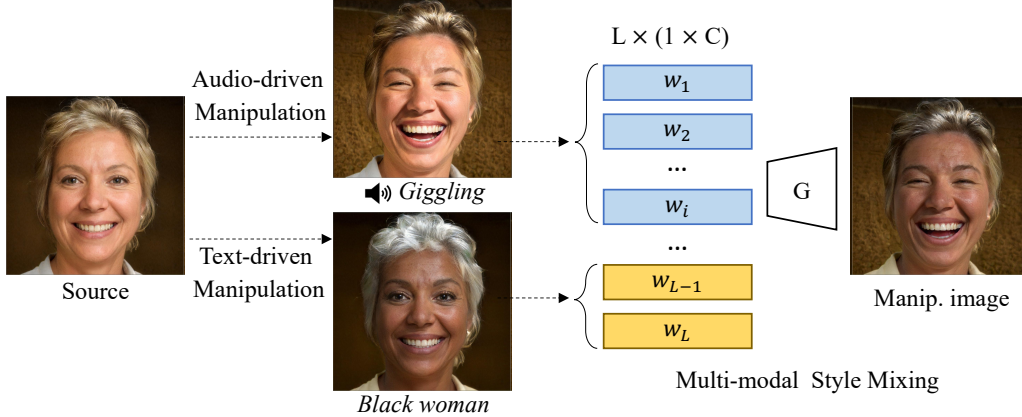


Figure 16: Overview of image style mixing jointly with the audio (*people giggling*) and text input (*black woman*).

GAN represent different properties. We perform multi-modal style mixing by selecting a specific layer of latent code and a mixing style that includes audio and text. Selecting a particular layer of each latent code guided by audio and text manipulates the image using the properties of audio and text (see Fig. 16).

We find that the sound source can effectively manipulate aspects of facial emotion such as “giggling” on the face and text information controls the background color of the target image. For the style-mixing details, we follow TediGAN’s StyleGAN layerwise analysis (Xia et al., 2021). In the 18×512 latent code, the style-mixing technique selects the 1st to 9th layers of the sound-guided latent code and the 10th to 18th layers of the text-guided latent code to mix the dynamic characteristics of sound and the human properties of text.

5. Applications

5.1. Sound-Guided Animal Face Image Manipulation

Sound-guided Animal Face Image Manipulation We further apply our method to manipulate the images of animal faces. We pre-train the StyleGAN3 (Karras et al., 2021) generator with the existing Animal FacesHQ (AFHQ) v2 (Choi et al., 2020) dataset, which contains over 15k images of multiple species (cat, dog, wildlife) and various breeds (\geq eight) per species. As shown in Fig. 14, our model successfully manipulates animal faces with




Figure 17: Quantitative evaluation and user study results. (a) Downstream task evaluation to compare the quality of representations between our approach and text-driven manipulation approaches on the FFHQ (Karras et al., 2019) dataset. A linear classifier is trained to predict 11 semantic labels, e.g. giggling. Participants answered a questionnaire on (b) Naturalness (“Which image manipulation result better expresses the target attribute?”) and (c) Perceptual Realism (“On the scale from 1 to 5, how realistic are the images below?”).

guidance by a sound input (e.g., laughing, dog barking, and little cat meowing).

5.2. Sound-Guided Artistic Painting Manipulation

We further apply our approach for sound-guided artistic painting manipulation. We train our StyleGAN2 (Karras et al., 2020)-based generator with a fine-art paintings dataset, called WikiArt (Saleh and Elgammal, 2016), which contains over 50k images of paintings from 195 different artists. We observe in Fig. 18 that our model successfully produces manipulated images for art paintings guided by audio inputs (e.g., laughing, crying, bomb, raining, or bird). For example, given a laughing sound as a semantic cue, a portrait painting of a woman is manipulated (see 1st column), but the existing approach by Lee *et al.* does not. Our human study (done in the same setting as our user study) also confirms that 61.2% of 50 participants judge that our results are better manipulated than those by Lee *et al.* (26.7%) and text-based approach (12.1%).

5.3. Music Style Transfer

The objective of music style transfer’s objective is that the source latent code is close to the keywords of the music, so the mood of the music appears in the image. For instance, *Funny* music manipulates the image




Figure 18: Examples of sound-guided artistic painting manipulation results.

with a fairy-tale style, whereas *Latin* music manipulates the image with red-colored theme that reflects the characteristic of *passion*. Lee *et al.* (Lee et al., 2022) demonstrate that CLIP-based sound representation can potentially reflect the music’s mood into the image style. In this study, we outperform previous works on visual style transfer for music. We collect well-known 10-second segments of well-known music and compare the image manipulation results (*The Lion King* OST, *The Toy Story* OST, *Despacito* title song, game sound effects, *Fortress* game song) for evaluation. Fig. 19 illustrates the results of image style transfer using various music genres. Visual information collected from text descriptions related to music genres in the AudioSet results in robust image manipulation for the music the model encounters for the first time.

6. Discussion

6.1. Limitations

Our method performs sound-guided image manipulation for various domains and synthesizes realistic manipulation results. However, our model



Figure 19: Examples of music style transfer using our method. Given 10 seconds of music, we manipulate the input image using Crossing you in style (Lee et al., 2020), Lee *et al.* (Lee et al., 2022) and our method.

requires a trained StyleGAN latent space, which increases the learning cost in real applications and limits the domain of visual content. In order to manipulate images in the open domains, it is necessary to obtain a prior that directly generates the image from the latent representation of the sound, rather than providing guidance to the StyleGAN latent space as a CLIP-based sound embedding space. This will be studied as future work.

6.2. Societal Impacts

It can have consequences for misuse in creating fake news or propaganda. Moreover, the method of sound-guided image manipulation is based on CLIP’s knowledge and may reflect social prejudice. In the CLIP paper,

the author says “CLIP is trained on text paired with images on the internet. These image-text pairs are unfiltered and uncured and result in CLIP models learning many social biases.”. Therefore, there is a possibility that certain social biases may appear when people faces are manipulated with sounds related to a thief, prisoner, criminal, or suspicious people, such as guns and fighting sounds. Also, when a human face is manipulated with a sound such as a vacuum cleaner or a women speaking, manipulation results with social prejudice like housekeepers may be generated.

7. Conclusion

We propose a method to manipulate images based on the semantic-level understanding from the given audio input. We take the user-provided audio input into the latent space of StyleGAN2 (Karras et al., 2020) and the CLIP (Radford et al.) embedding space. In particular, the strategy of adding a semantic edge between audio-visual from text obtains a more robust sound representation for fine-features. Then, the latent code is aligned with the audio to enable meaningful image manipulation while reflecting the context from the audio. Our model produces responsive manipulations based on various audio inputs. We observe that an audio input can successfully provide a semantic cue to manipulate images accordingly. Our method of traversing multi-modal embedding space can be used in many applications with multi-modal contexts.

Acknowledgments

This research was fully supported by Culture, Sports and Tourism R&D Program through the Korea Creative Content Agency grant funded by the Ministry of Culture, Sports and Tourism in 2022 (Project Name: 4D Content Generation and Copyright Protection with Artificial Intelligence, Project Number: R2022020068).

References

Abdal, R., Qin, Y., Wonka, P., 2019. Image2stylegan: How to embed images into the stylegan latent space?, in: Proceedings of the IEEE/CVF International Conference on Computer Vision, pp. 4432–4441.

- Alayrac, J.B., Recasens, A., Schneider, R., Arandjelovic, R., Ramapuram, J., De Fauw, J., Smaira, L., Dieleman, S., Zisserman, A., 2020. Self-supervised multimodal versatile networks. *NeurIPS* 2, 7.
- Aytar, Y., Vondrick, C., Torralba, A., 2017. See, hear, and read: Deep aligned representations. *CoRR* abs/1706.00932. URL: <http://arxiv.org/abs/1706.00932>, [arXiv:1706.00932](https://arxiv.org/abs/1706.00932).
- Brouwer, H., 2020. Audio-reactive latent interpolations with stylegan, in: *NeurIPS 2020 Workshop on Machine Learning for Creativity and Design*.
- Chen, H., Xie, W., Vedaldi, A., Zisserman, A., 2020. Vggsound: A large-scale audio-visual dataset, in: *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, IEEE. pp. 721–725.
- Chen, L., Srivastava, S., Duan, Z., Xu, C., 2017. Deep cross-modal audio-visual generation, in: *Proceedings of the on Thematic Workshops of ACM Multimedia 2017*, pp. 349–357.
- Chen, Y., Xian, Y., Koepke, A., Shan, Y., Akata, Z., 2021. Distilling audio-visual knowledge by compositional contrastive learning, in: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 7016–7025.
- Choi, Y., Uh, Y., Yoo, J., Ha, J.W., 2020. Stargan v2: Diverse image synthesis for multiple domains, in: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 8188–8197.
- Deng, J., Guo, J., Xue, N., Zafeiriou, S., 2019. 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.
- Dong, H., Yu, S., Wu, C., Guo, Y., 2017. Semantic image synthesis via adversarial learning, in: *Proceedings of the IEEE International Conference on Computer Vision*, pp. 5706–5714.
- El-Nouby, A., Sharma, S., Schulz, H., Hjelm, D., Asri, L.E., Kahou, S.E., Bengio, Y., Taylor, G.W., 2019. Tell, draw, and repeat: Generating and

- modifying images based on continual linguistic instruction, in: Proceedings of the IEEE/CVF International Conference on Computer Vision, pp. 10304–10312.
- Fellbaum, C., 2010. Wordnet, in: Theory and applications of ontology: computer applications. Springer, pp. 231–243.
- Gatys, L., Ecker, A., Bethge, M., 2016. A neural algorithm of artistic style. *Journal of Vision* 16, 326–326.
- Gemmeke, J.F., Ellis, D.P., Freedman, D., Jansen, A., Lawrence, W., Moore, R.C., Plakal, M., Ritter, M., 2017. Audio set: An ontology and human-labeled dataset for audio events, in: 2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), IEEE. pp. 776–780.
- Guzhov, A., Raue, F., Hees, J., Dengel, A., 2021. Audioclip: Extending clip to image, text and audio. [arXiv:2106.13043](#).
- Hao, W., Zhang, Z., Guan, H., 2018. Cmcgan: A uniform framework for cross-modal visual-audio mutual generation, in: Proceedings of the AAAI Conference on Artificial Intelligence.
- Hershey, S., Chaudhuri, S., Ellis, D.P., Gemmeke, J.F., Jansen, A., Moore, R.C., Plakal, M., Platt, D., Saurous, R.A., Seybold, B., et al., 2017. Cnn architectures for large-scale audio classification, in: 2017 IEEE international conference on acoustics, speech and signal processing (icassp), IEEE. pp. 131–135.
- Hinton, G., Vinyals, O., Dean, J., et al., 2015. Distilling the knowledge in a neural network. [arXiv preprint arXiv:1503.02531](#) 2.
- Jeong, D., Doh, S., Kwon, T., 2021. Träumerai: Dreaming music with style-gan. [arXiv preprint arXiv:2102.04680](#) .
- Jiang, W., Xu, N., Wang, J., Gao, C., Shi, J., Lin, Z., Liu, S., 2021. Language-guided global image editing via cross-modal cyclic mechanism, in: Proceedings of the IEEE/CVF International Conference on Computer Vision, pp. 2115–2124.
- Karras, T., Aittala, M., Laine, S., Härkönen, E., Hellsten, J., Lehtinen, J., Aila, T., 2021. Alias-free generative adversarial networks. *Advances in Neural Information Processing Systems* 34, 852–863.

- Karras, T., Laine, S., Aila, T., 2019. 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.
- Karras, T., Laine, S., Aittala, M., Hellsten, J., Lehtinen, J., Aila, T., 2020. Analyzing and improving the image quality of stylegan, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 8110–8119.
- Lee, C.C., Lin, W.Y., Shih, Y.T., Kuo, P.Y., Su, L., 2020. Crossing you in style: Cross-modal style transfer from music to visual arts, in: Proceedings of the 28th ACM International Conference on Multimedia, pp. 3219–3227.
- Lee, S.H., Roh, W., Byeon, W., Yoon, S.H., Kim, C., Kim, J., Kim, S., 2022. Sound-guided semantic image manipulation, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 3377–3386.
- Li, B., Qi, X., Lukasiewicz, T., Torr, P.H., 2020. Manigan: Text-guided image manipulation, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 7880–7889.
- Liu, Z., Lin, Y., Cao, Y., Hu, H., Wei, Y., Zhang, Z., Lin, S., Guo, B., 2021. Swin transformer: Hierarchical vision transformer using shifted windows, in: Proceedings of the IEEE/CVF International Conference on Computer Vision, pp. 10012–10022.
- van der Maaten, L., Hinton, G., 2008. Visualizing high-dimensional data using t-629 sne. *Journal of Machine Learning Research* 9, 630.
- Mazumder, P., Singh, P., Parida, K.K., Namboodiri, V.P., 2021. Avgzslnet: Audio-visual generalized zero-shot learning by reconstructing label features from multi-modal embeddings, in: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision, pp. 3090–3099.
- Nagrani, A., Albanie, S., Zisserman, A., 2018. Learnable pins: Cross-modal embeddings for person identity, in: Proceedings of the European Conference on Computer Vision (ECCV), pp. 71–88.

- Nam, S., Kim, Y., Kim, S.J., 2018. Text-adaptive generative adversarial networks: Manipulating images with natural language, in: Advances in Neural Information Processing Systems (NeurIPS).
- Oh, T.H., Dekel, T., Kim, C., Mosseri, I., Freeman, W.T., Rubinstein, M., Matusik, W., 2019. Speech2face: Learning the face behind a voice, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 7539–7548.
- Park, D.S., Chan, W., Zhang, Y., Chiu, C.C., Zoph, B., Cubuk, E.D., Le, Q.V., 2019a. SpecAugment: A Simple Data Augmentation Method for Automatic Speech Recognition, in: Proc. Interspeech 2019, p. 2613–2617. doi:10.21437/Interspeech.2019-2680.
- Park, T., Liu, M.Y., Wang, T.C., Zhu, J.Y., 2019b. Semantic image synthesis with spatially-adaptive normalization, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 2337–2346.
- Patashnik, O., Wu, Z., Shechtman, E., Cohen-Or, D., Lischinski, D., 2021. Styleclip: Text-driven manipulation of stylegan imagery, in: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV), pp. 2085–2094.
- Piczak, K.J., 2015. Esc: Dataset for environmental sound classification, in: Proceedings of the 23rd ACM international conference on Multimedia, pp. 1015–1018.
- Qiu, Y., Kataoka, H., 2018. Image generation associated with music data, in: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops, pp. 2510–2513.
- 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.
- Richardson, E., Alaluf, Y., Patashnik, O., Nitzan, Y., Azar, Y., Shapiro, S., Cohen-Or, D., 2021. 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.

- Salamon, J., Jacoby, C., Bello, J.P., 2014. A dataset and taxonomy for urban sound research, in: 22nd ACM International Conference on Multimedia (ACM-MM’14), Orlando, FL, USA. pp. 1041–1044.
- Saleh, B., Elgammal, A., 2016. Large-scale classification of fine-art paintings: Learning the right metric on the right feature. *International Journal for Digital Art History* .
- Skorokhodov, I., Sotnikov, G., Elhoseiny, M., 2021. Aligning latent and image spaces to connect the unconnectable, in: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV), pp. 14144–14153.
- Smith, L.N., 2017. Cyclical learning rates for training neural networks, in: 2017 IEEE winter conference on applications of computer vision (WACV), IEEE. pp. 464–472.
- Sun, Z., Sarma, P., Sethares, W., Liang, Y., 2020. Learning relationships between text, audio, and video via deep canonical correlation for multimodal language analysis, in: Proceedings of the AAAI Conference on Artificial Intelligence, pp. 8992–8999.
- Surís, D., Duarte, A., Salvador, A., Torres, J., Giró-i Nieto, X., 2018. Cross-modal embeddings for video and audio retrieval, in: Proceedings of the European Conference on Computer Vision (ECCV) Workshops, pp. 0–0.
- Tov, O., Alaluf, Y., Nitzan, Y., Patashnik, O., Cohen-Or, D., 2021. Designing an encoder for stylegan image manipulation. *ACM Trans. Graph.* 40. URL: <https://doi.org/10.1145/3450626.3459838>, doi:10.1145/3450626.3459838.
- Wan, C.H., Chuang, S.P., Lee, H.Y., 2019. Towards audio to scene image synthesis using generative adversarial network, in: ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), IEEE. pp. 496–500.
- Wu, H.H., Seetharaman, P., Kumar, K., Bello, J.P., 2021. Wav2clip: Learning robust audio representations from clip. *arXiv:2110.11499*.
- Xia, W., Yang, Y., Xue, J.H., Wu, B., 2021. 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.

Yu, F., Zhang, Y., Song, S., Seff, A., Xiao, J., 2015. Lsun: Construction of a large-scale image dataset using deep learning with humans in the loop .

Zhu, H., Luo, M.D., Wang, R., Zheng, A.H., He, R., 2021. Deep audio-visual learning: A survey. *International Journal of Automation and Computing* , 1–26.