

# Generating Visually Aligned Sound from Videos

Peihao Chen, Yang Zhang, Mingkui Tan, Hongdong Xiao, Deng Huang, and Chuang Gan

**Abstract**—We focus on the task of generating sound from natural videos, and the sound should be both temporally and content-wise aligned with visual signals. This task is extremely challenging because some sounds generated *outside* a camera can not be inferred from video content. The model may be forced to learn an incorrect mapping between visual content and these irrelevant sounds. To address this challenge, we propose a framework named REGNET. In this framework, we first extract appearance and motion features from video frames to better distinguish the object that emits sound from complex background information. We then introduce an innovative audio forwarding regularizer that directly considers the real sound as input and outputs bottlenecked sound features. Using both visual and bottlenecked sound features for sound prediction during training provides stronger supervision for the sound prediction. The audio forwarding regularizer can control the irrelevant sound component and thus prevent the model from learning an incorrect mapping between video frames and sound emitted by the object that is out of the screen. During testing, the audio forwarding regularizer is removed to ensure that REGNET can produce purely aligned sound only from visual features. Extensive evaluations based on Amazon Mechanical Turk demonstrate that our method significantly improves both temporal and content-wise alignment. Remarkably, our generated sound can fool the human with a 68.12% success rate. Code and pre-trained models are publicly available at <https://github.com/PeihaoChen/regnet>.

**Index Terms**—video sound generation, visually aligned sound, audio forwarding regularizer.

## I. INTRODUCTION

VARIOUS visual events in our daily life are usually accompanied by different sounds. Because visual events and sounds are correlated, a person can instinctively infer sounds by observing visual events. In this paper, we will address the task of deriving sound from silent videos, which is beneficial to many real-world applications, such as video editing automation, generating sound for silent film, and assistance for people with visual impairment.

The *alignment* to the corresponding video is an essential characteristic for plausible generated sound. Specifically, a visually aligned sound should have two aspects. The first aspect

is the *temporal alignment*, *i.e.*, synchronization. A temporally-aligned approach should produce a sound event at exactly the time or soon after the time when the corresponding visual event occurs. For example, a dog barking should be synchronized with the visual event in which the dog opens its mouth and leans forward. The second aspect of the alignment is the *content-wise alignment*. A sound that is content-wise aligned should produce only the types of sound that correspond to what is happening in the video. For example, when a dog opens its mouth, we would expect to hear a dog barking instead of a human utterance.

Unfortunately, despite various recent research efforts that apply deep generative models to produce visually aligned sound that is conditional on visual features [1], [2], the alignment issue remains challenging. However, alignment is a considerably lesser concern for other sound generation tasks, including speech synthesis and voice conversion. This leads to the following research question: why is alignment especially challenging for generating visually aligned sounds?

In this paper, we will uncover an important explanation from an information perspective. Unlike speech synthesis, where text and speech have significant information correspondence, in the task of generating sound from video, visual and sound information do not have a strict correspondence. In most cases, sound can be decomposed into two parts. The first part is visually *relevant* sound, which can be inferred from video content, *e.g.*, when we see a dog open its mouth and lean forward, we will naturally associate this video clip with the sound of barking. The second part is visually *irrelevant* sound that cannot be inferred from video content, *e.g.*, people talking behind the camera. In the training phase of the existing sound generation paradigm, only video frames are taken as input and the loss is calculated between the predicted sound and the ground truth sound which contains both visually relevant and irrelevant information. This *information mismatch* issue in the existing training paradigm may confuse the model. To minimize the loss, the model is forced to learn an incorrect mapping between video frames and visually irrelevant sound, which cripples the alignment performance.

To generate visually aligned sound from videos, we propose a spectrogram-based sound generation model named REGNET. Specifically, there are two mechanisms to ensure visual alignment. First, REGNET introduces a time-dependent appearance and motion feature extraction module, which provides sufficient visual information for generating the temporal and content-wise aligned sounds. Second and more critically, to resolve the information mismatch issue induced by visually irrelevant sound information, we introduce a novel mechanism referred to as *audio forwarding regularizer*. The audio forwarding regularizer passes the ground truth sound through an encoder with a bottleneck and applies the resulting bottlenecked sound

This work was partially supported by Science and Technology Program of Guangzhou, China under Grants 202007030007, National Natural Science Foundation of China (NSFC) 61836003 (key project), Program for Guangdong Introducing Innovative and Entrepreneurial Teams 2017ZT07X183, Fundamental Research Funds for the Central Universities D2191240.

P. Chen is with the School of Software Engineering, South China University of Technology, Guangzhou 510640, China, and also with the Pazhou Laboratory, Guangzhou, 510335, China (e-mail: phchencs@gmail.com).

M. Tan, H. Xiao, and D. Huang are with the School of Software Engineering, South China University of Technology, Guangzhou 510641, China (e-mail: mingkuitan@scut.edu.cn; xiaohongdonghd@gmail.com; im.huangdeng@gmail.com).

Y. Zhang and C. Gan are with the MIT-IBM Watson AI Lab, Cambridge, MA 02142, USA (e-mail: yang.zhang2@ibm.com; ganchuang1990@gmail.com).

P. Chen and Y. Zhang contributed equally to this paper.

M. Tan is the corresponding author.

features with the visual feature to predict the ground truth sound itself. We will show in the experiment part that this mechanism is helpful for visually aligned sound generation. During training, the audio forwarding regularizer can control for the irrelevant sound component to ensure that the audio generation module can learn a correct mapping between visual features and the relevant sound component. During testing, the audio forwarding regularizer is removed, so that REGNET is able to generate purely aligned sound only from visual feature. In addition, we introduce generative adversarial networks (GANs) [3] to improve the quality of the generated sound. We have assembled a video dataset that contains distinctive video and sound events (and thus alignment between sound and video is crucial). Extensive evaluations on Amazon Mechanical Turk (AMT) demonstrate that REGNET significantly improves over existing methods in terms of the temporal and content-wise alignment of the generated sound.

Our main contributions are listed as follows:

- We explore the challenges of generating visually-aligned sound in the aspect of information mismatch between visual features and sounds. To solve these challenges, we propose a novel audio forwarding regularizer to provide missing information during the training phase.
- We propose REGNET, which is an algorithm capable of generating visually aligned sound in high temporal and content-wise alignment.

## II. RELATED WORKS

**Visually aligned sound synthesis.** Synthesizing audio for video has recently attracted considerable attention [4], [5], [2], [3], [1]. Owens *et al.* [4] propose the task of predicting the sound emitted by hitting and scratching objects with a drumstick. Chen *et al.* [5] exploit conditional generative adversarial networks [3] to achieve cross-modal audio-visual generation of musical performances. Chen *et al.* [1] propose to generate sound considering the sound class and use perceptual loss to align semantic information. Zhou *et al.* [2] collect an unconstrained dataset (VEGAS) that includes 10 types of sound recorded in the wild and propose a SampleRNN-based method to directly generate a waveform from a video. However, these studies disregard the information mismatch issue between video and sound, which may cause poor temporal synchronization and content-wise alignment.

**Visual-sound learning.** The synchronized visual and sound information contained in most videos provide supervision for each other [6], [7]. Based on visual supervision, Harwath *et al.* [8] present a network to learn associations between natural image scenes and audio captions. Arandjelovic *et al.* [9] and Aytar *et al.* [10] learn the visual-audio correlation via unlabeled videos in an unsupervised manner. Several recent works [11], [12], [13] perform visual sound separation using visual-sound representation. Also, Gan *et al.* [14] present a moving vehicle tracking system based on visual-sound relationships. Some other interesting visual-sound researches include audio-visual co-segmentation [15] and audio-visual navigation [16]. We also learn visual-sound representation and transform information from the video domain to the audio domain.

**Text to speech (TTS).** TTS is to synthesize audio speak from text. Ling *et al.* [17] use neural networks to predict sound features from pre-tokenized text features and then generate a waveform from these features. Van *et al.* [18] synthesize impact sounds from physical simulations. Van *et al.* [19] introduce WaveNet to generate a raw audio waveform. Mehri *et al.* [20] propose a model for unconditional audio generation, and Sotelo *et al.* [21] present Char2Wav for speech synthesis. Shen *et al.* [22] propose to predict mel spectrogram from text and then transform it to a waveform using WaveNet. Our task is closely related to TTS, but we consider visual information the input and synthesize various kinds of visually aligned sounds in daily life.

**Generative adversarial network.** GANs [3] have achieved considerable success in generating high-quality images [23], [24], [25], [26], [27]. However, directly adapting image GANs architectures, such as CGANs [28], DCGANs [29] and WGANs [30], may be problematic for audio generation. Inspired by Reed *et al.* [31], which generate images conditioned on text captions, Chen *et al.* [5] design conditional GANs to achieve cross-modal audio-visual generation. By investigating waveform and spectrogram strategies, Donahue *et al.* [32] attempt to apply GANs to synthesize one-second slices of raw-waveform audio in an unsupervised manner. To generate high-fidelity and locally coherent audio, Engel *et al.* [33] apply GANs to model log magnitudes and instantaneous frequencies with sufficient frequency resolution in the spectral domain. In this paper, we apply GANs on a spectrogram to improve the sound quality.

## III. REGNET FRAMEWORK

In this section, we will introduce the challenge of generating visually aligned sound, together with the description about our proposed REGNET framework. Then, we will discuss why the REGNET is capable of solving this challenge and learn a correct mapping between video frames and visually relevant sound.

### A. Problem formulation

First, we will mathematically formulate the problem and discuss its challenges. For the remainder of this section, we will use the upper-case letters  $X$  or  $\mathbf{X}$  to denote random variables (unbolded) or random vectors (bolded), respectively, and the lower-case letters,  $x$  or  $\mathbf{x}$  to denote deterministic values.  $\mathbb{E}[\cdot]$  denotes the expectation.

Denote  $(\mathbf{V}(t), \mathbf{S}(\tau))$  as a visual-sound pair.  $\mathbf{V}(t)$  represents the visual signal (vectorized) at each frame  $t$ .  $\mathbf{S}(\tau)$  represents the sound representation (waveform or spectrogram) at each frame  $\tau$ . We use different frame indexes  $t$  and  $\tau$  because visual and sound signals have different sampling rates.

We introduce our formal modeling of the challenge that sound has a visually irrelevant component that can not be inferred from video content, such as people talking behind a camera. Assume that the audio can be decomposed into a relevant signal and an irrelevant signal:

$$\mathbf{S}(\tau) = \mathbf{S}_r(\tau) + \mathbf{S}_i(\tau), \quad (1)$$

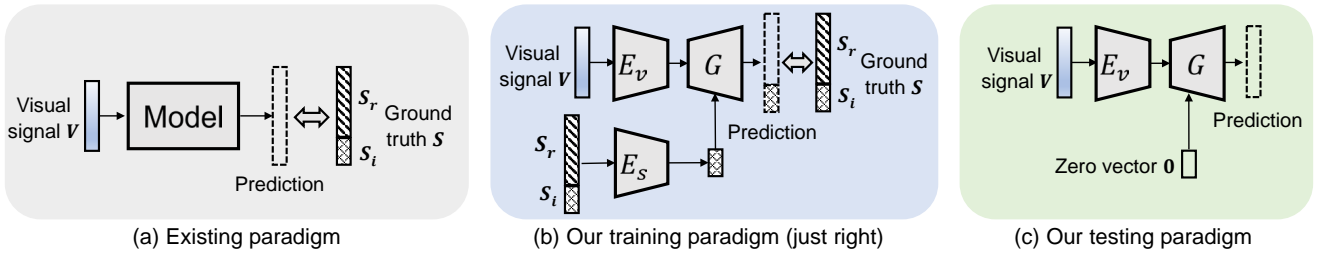


Fig. 1: Comparisons between the existing paradigm and our training and testing paradigm. (a) For the existing paradigm, the model is forced to learn an incorrect mapping between a visual signal and visually irrelevant sound. (b) We avoid this situation by incorporating an audio forwarding regularizer. (c) During the testing phase, the visually relevant sound is predicted by removing the audio forwarding regularizer.

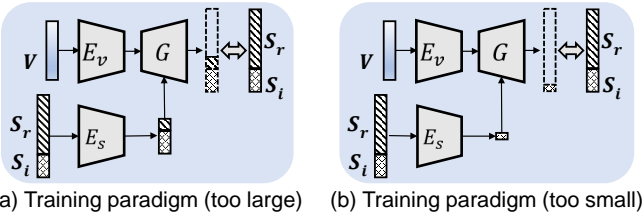


Fig. 2: Influence of different information capacities of audio forwarding regularizer. (a) Because the capacity is too large, the generator perfectly reconstructs the sound but has a weak dependence on the visual signal  $V$ . (b) Because the capacity is too small, the generator is forced to learn an incorrect mapping between the visual signal  $V$  and the irrelevant sound  $S_i$ .

where the subscript  $r$  denotes “relevant”; the subscript  $i$  denotes “irrelevant”. We further assume that a relation *only* exists between the video and relevant sound, which is denoted as  $f(\cdot)$ . Irrelevant sound is independent of both the relevant sound and visual features. Formally,

$$S_r(\tau) = f(V(t)), \quad S_i(\tau) \perp S_r(\tau), \quad S_i(\tau) \perp V(t), \quad (2)$$

where  $\perp$  denotes independence.

Our goal is to generate the visually relevant component  $S_r$ <sup>1</sup> from the visual signal  $V$ . However, the visually relevant component  $S_r$  is not directly available during training. Instead, we only have the sound representation  $S$ , which is a mixture of  $S_r$  and  $S_i$ . As shown in Figure 1 (a), the existing models take the visual signal  $V$  as input and regard the sound representation  $S$  as target. In this setting, only the visual information is provided but the model is forced to predict both visually relevant and irrelevant sound which is independent of the visual signal. With insufficient information, the model can only learn an incorrect mapping between  $V$  and  $S_i$ , which produces artifacts and misalignment in generated sound. We refer to it as an information mismatch challenge.

### B. REGNET framework with audio forwarding regularizer

We propose a REGNET to resolve the information mismatch challenge by introducing an audio forwarding regularizer. As

<sup>1</sup>The time argument is removed to represent that it is a collection of the signal at ALL frames.

shown in Figure 1 (b), REGNET consists of three modules: the visual encoder  $E_v(\cdot)$ , the audio forwarding regularizer  $E_s(\cdot)$  and the generator  $G(\cdot, \cdot)$ .

The visual encoder  $E_v(\cdot)$  takes the visual signal  $V$  as the input and outputs a set of frame features. The audio forwarding regularizer  $E_s(\cdot)$  takes the sound signal  $S$  as the input and outputs what we refer to as the regularizer output. The generator  $G(\cdot, \cdot)$  then predicts  $S$  with frame features and the regularizer output as input. There are two different types of predictions, with or without the audio forwarding. The prediction with audio forwarding, denoted as  $\hat{S}_a$ , is predicted from both the frame features and the regularizer output. The prediction without audio forwarding, denoted as  $\hat{S}_0$ , is predicted by removing audio forwarding regularizer. As shown in Figure 1 (c), in this work, we replace the regularizer output by a zero vector when predicting  $\hat{S}_0$ . Formally,

$$\hat{S}_a = G(E_v(V), E_s(S)), \quad \hat{S}_0 = G(E_v(V), \mathbf{0}). \quad (3)$$

The predicted sound with audio forwarding or without audio forwarding, respectively, *i.e.*,  $\hat{S}_a$  and  $\hat{S}_0$ , will be utilized during the training and testing phases. We exploit the adversarial training mechanism to render the predicted sound more realistic. The discriminator  $D(\cdot, \cdot)$  tries to discriminate whether the input sound is real or fake, conditional on the visual signal, *i.e.*,  $D(S', V)$ , where  $S'$  can be either generated sound or real-world sound.

From Equation (3), the significant difference between REGNET and the existing paradigm is the introduction of the audio forwarding regularizer, which can be regarded an information bottleneck to control for the audio forwarding. As will be discussed in the next subsection, with careful bottleneck tuning in the training phase, audio forwarding regularizer can provide supplementary visually irrelevant sound information for the generator. With sufficient information in the training phase, including visual signal and irrelevant sound information, the information mismatch challenge is resolved. The generator can learn the correct mapping between a visual signal and visually aligned sound. During the testing phase, the interference of the irrelevant sound is removed, and the generator can generate visually relevant sound using only visual signals.

**Training details:** During training, the generator tries to minimize the following loss, which involves the prediction

with audio forwarding:

$$L_{\text{rec}} + L_G = \mathbb{E}[\|\hat{S}_a - S\|_2^2] + \mathbb{E}[\log(1 - D(\hat{S}_a, \mathbf{V}))], \quad (4)$$

where the first term is the L2 reconstruction error, and the second term is the adversarial loss.

On the other hand, the discriminator tries to discriminate the real sound from the predicted sound with audio forwarding, which minimizes the following standard adversarial loss:

$$L_D = -\mathbb{E}[\log(D(S, \mathbf{V}))] - \mathbb{E}[\log(1 - D(\hat{S}_a, \mathbf{V}))]. \quad (5)$$

**Testing details:** In the testing phase, we will not use any regularizer output that involves the ground truth sound. Instead, we will consider the  $\hat{S}_0$  defined in Equation (3) the predicted results.

As readers may have noticed, REGNET introduces two relatively unintuitive mechanisms. First, the audio forwarding regularizer sends information about  $S$  to predict  $S$ . This process resembles cheating during training. Second, a mismatch exists between training and testing. During training,  $\hat{S}_a$  is employed, whereas  $\hat{S}_0$  is employed during testing. In the next subsection, we will show that these mechanisms can produce the desirable outcome.

### C. Why does REGNET work?

In a visually sound generation task, we expect the generator to infer visually relevant sound from a visual feature. During training, however, only a mixture of visually relevant and irrelevant sounds is available to be regarded the target. The generator is forced to infer the irrelevant sound to minimize the gap between the prediction and the target. Considering the visual feature the input, the generator cannot reasonably infer irrelevant sound because they are independent of each other. To resolve this conflict, some information about visually irrelevant sound should also be considered the input. In this paper, we introduce an audio forwarding regularizer to provide this information. Two sources of information are available to predict  $S$ : the visual feature and the regularizer output. By carefully tuning the information capacity of the audio forwarding channel, the regularizer output can take the responsibility of predicting visually irrelevant sound, which enables the generator to learn the correct mapping between a visual feature and the desired sound. In the following, we analyze the influence of information capacity in three cases.

**Too large capacity:** if the information capacity is too large (as shown in Figure 2 (a)), then the audio forwarding channel can send all information about  $S$  for prediction, which enables a perfect prediction of  $S$  but a lack of dependence on the visual feature. As the capacity starts to decrease, the audio forwarding channel will prioritize sending the information about  $S_i$  and enable the visual feature to address the information about  $S_r$ , because the regularizer output is the only source of sound  $S_i$ .

**Too small capacity:** if the information capacity is too small (as shown in Figure 2 (b)), then some information about  $S_i$  will not be accounted for. To minimize the reconstruction loss between the prediction and the sound  $S$ , the model is forced to learn incorrect mapping between a visual feature and the irrelevant sound  $S_i$ , resulting in artifacts and misalignment

in the generated sound, which is the case for most existing training paradigms.

**Just right capacity:** if the information capacity is just right (as shown in Figure 1 (b)), the audio forwarding channel sends all information about  $S_i$ , while the visual feature takes responsibility for the  $S_r$  part. The generator is allowed to learn a solid mapping between a visual feature and the relevant sound  $S_r$ . During testing, the audio forwarding regularizer is removed and the REGNET can infer the desired visually aligned sound from video content, which eliminates the inference of visually irrelevant sound.

Therefore, the key to the success of REGNET is to obtain proper control for the information capacity of an audio forwarding regularizer.

## IV. REGNET ARCHITECTURE

In this section, we will illustrate the architecture of each module in REGNET and provide details of the vocoder, which is used to convert the predicted spectrogram into a waveform. The schematic of REGNET is shown in Figure 3.

**Visual encoder.** We design a time-dependent visual encoder to extract appearance and motion features for visually aligned sound generation. Specifically, a BN-Inception [34] model is utilized as a feature extractor to explore RGB and flow features at  $T$  different time steps, which are the visual signal in Section (III). The BN-Inception model is pretrained on ImageNet and frozen in the training phase. The concatenation of  $T$  RGB and flow features is processed by the visual encoder, which consists of three 1D convolutional layers and a two-layer bidirectional LSTM (Bi-LSTM) [35]. Each convolutional layer is followed by a batch normalization (BN) [34] layer and a rectified linear unit (ReLU) [36] activation function. The outputs in the forward and backward paths of Bi-LSTM at each time step are concatenated as  $T$  encoded visual features.

**Audio forwarding regularizer.** We design an audio forwarding regularizer to provide supplementary visually irrelevant sound information for sound prediction during training. The audio forwarding regularizer takes as input the ground truth spectrogram, which is the mixture of visually relevant and irrelevant sounds. The ground truth spectrogram is processed by a two-layer Bi-LSTM with the cell dimension  $D$ . The concatenation of the outputs in the two paths of Bi-LSTM is a  $2D \times T'$  feature map, where  $T'$  is the time dimension of the input spectrogram. Then, we uniformly downsample this feature map by  $S$ , which generates a  $2D \times (T'/S)$  feature map with some information eliminated. To match the temporal dimension of encoded visual features, we upsample it by replication and generate the  $2D \times T$  regularizer output. Note that the Bi-LSTM can be regarded as an information bottleneck. We control for its information capacity by changing its output dimension  $D$  and adjusting the downsampling rate  $S$ . The audio forwarding regularizer with a larger output dimension and smaller downsampling rate passes through richer sound information. The intuitive influence and empirical influence of the information capacity are discussed in subsection III-C and V-H, respectively.

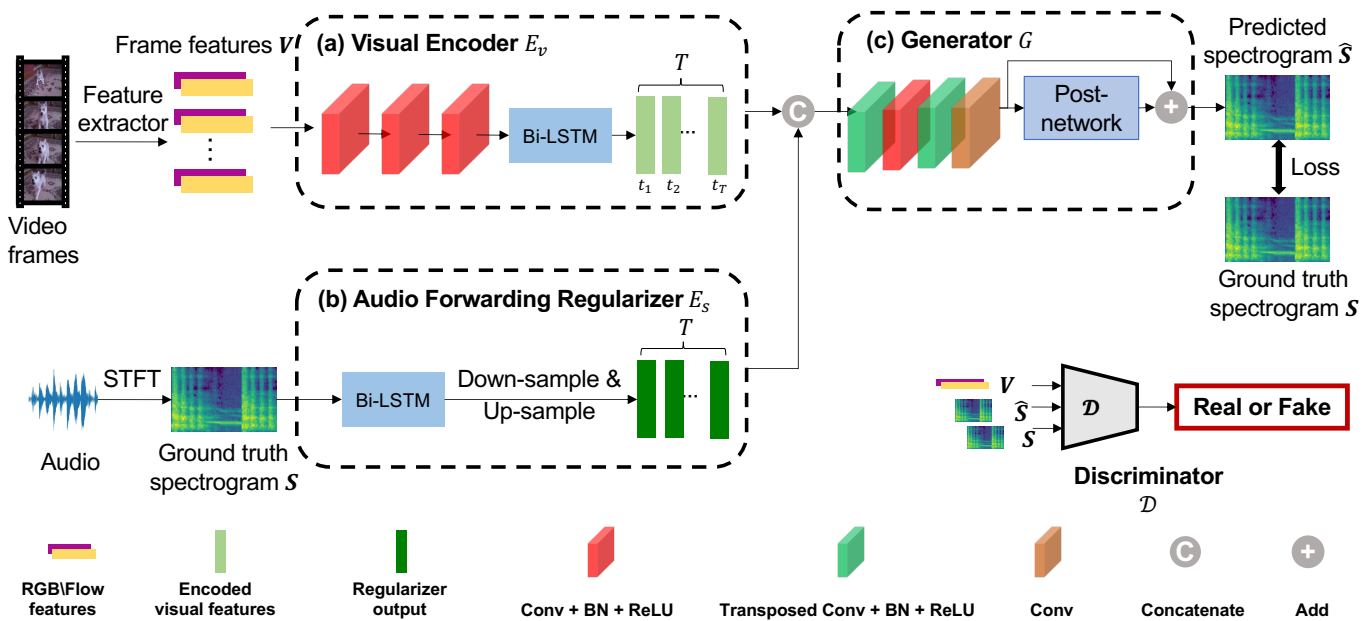


Fig. 3: Schematic of our REGNET. (a) The visual encoder summarizes the extracted RGB and flow frame features as time-dependent visual features. (b) The audio forwarding regularizer takes the real sound as input and outputs bottlenecked sound features. (c) The generator considers visual features and bottlenecked real sound information to reconstruct the real sound (spectrogram) in the training phase.

**Generator.** We design a generator to predict sound (spectrogram  $\hat{S}^2$ ) from the concatenation of an encoded visual feature and regularizer output. First, we predict the initial spectrogram  $\hat{I}$  by two 1D convolutional layers and two 1D transposed convolutional layers. With the exception of the last layer, each layer is followed by a BN layer and a ReLU activation function. Then, a post-network [22], which consists of five 1D convolutional layers, is introduced upon  $\hat{I}$  to add the fine structure  $\hat{R}$  in a residual way. Specifically,  $\hat{S} = \hat{I} + \hat{R}$ , where  $\hat{I}$  and  $\hat{R}$  are the input and output of the post-network.

In addition to the supervision on the final predicted spectrogram  $\hat{S}$ , we also add the L2 constraint on the initial predicted spectrogram  $\hat{I}$ , which is expressed as follows:

$$L'_{\text{rec}} = \mathbb{E}[\|\hat{I} - S\|_2^2]. \quad (6)$$

The loss in Equation (4), (5) and (6) can be summarized as follows:

$$L_{\text{total}} = L_{\text{rec}} + \alpha L'_{\text{rec}} + \beta(L_G + L_D). \quad (7)$$

**Discriminator.** We introduce the discriminator  $D$  for adversarial training. The discriminator  $D$  takes as input the extracted frame feature and a spectrogram and distinguishes whether the spectrogram is derived from real video or generated by the proposed REGNET. The PatchGANs [37] is applied to preserve the high-frequency structure in the local area. Specifically, we first separately process the input frame feature and spectrogram with two 1D transposed convolutional layers and a 1D convolutional layer. Then, we concatenate them and use four 1D convolutional layers to predict a score that indicates whether the input spectrogram is real or fake.

<sup>2</sup>We use  $\hat{S}$  to represent the spectrogram with or without audio forwarding.

**Vocoder.** Given a synthesized spectrogram, we apply WaveNet [19] to convert it to a waveform. Instead of predicting the discretized value, we estimate a 10-component mixture of logistic distributions (MoL) to generate 16-bit samples at 22,050 Hz. We independently train the WaveNet model for each sound class.

## V. EXPERIMENTS

In this section, first, we describe the datasets in experiments and then provide implementation details. Second, we compare the proposed REGNET with a state-of-the-art method in terms of the alignment and reality. Last, we conduct ablation studies to explore the influence from the audio forwarding regularizer in REGNET.

### A. Datasets

To evaluate the temporal alignment of our generated sound, we collect 8 video categories that have high synchronization between visual and audio contents. Such synchronization is evaluated according to the “visual relevant” experiment introduced in [2]. Specifically, we exchange the sound between two videos in the same sound class and then test the turkers on Amazon Mechanical Turk (AMT) to determine whether they perceive this change. If the turkers are sensitive to this change, it means the sound is highly synchronized with the visual content. According to the results, we select 4 sound classes from the VEGAS [2] dataset, namely **Fireworks, Dog, Drum, Baby crying**, and 4 sound classes from the AudioSet [38] dataset, namely **Cough, Hammer, Gun, Sneeze**. We denote them as *Sub-VEGAS* and *Sub-AudioSet* dataset respectively. On average, each class contains 1626 videos, and the length

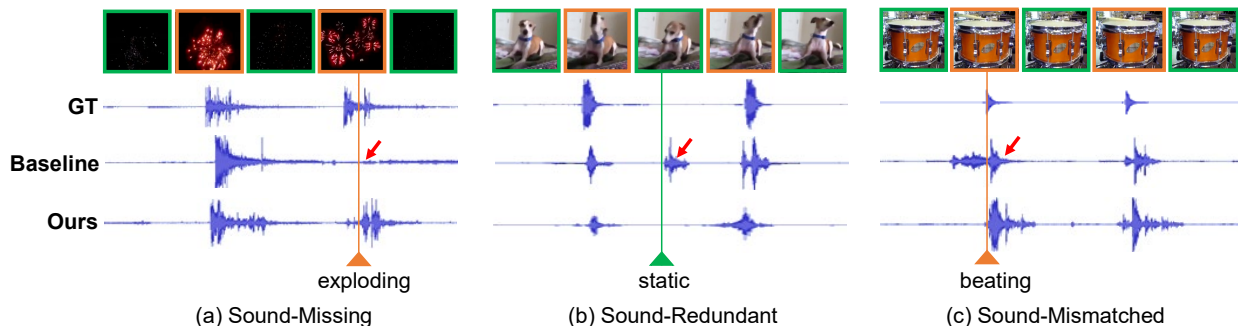


Fig. 4: Evaluation on alignment with 3 misaligned cases (from left to right), including *Sound-Missing*, *Sound-Redundant*, and *Sound-Mismatched*.

TABLE I: Evaluation of alignment. We report the percentage of the human judgments, where the results are preferred over the other methods based on three criteria.

	Baseline	Ours
<i>Sound-Missing</i>	39.38%	<b>60.62%</b>
<i>Sound-Redundant</i>	40.67%	<b>59.33%</b>
<i>Sound-Mismatched</i>	43.23%	<b>56.77%</b>

for each video is 6.73 seconds. The last 128 videos from each class are selected to build a testing set.

### B. Implementation details

**Data preprocessing.** The training videos are at 21.5 frames per second (fps) and are padded to 10 seconds by duplicating and concatenating. We extract the RGB and flow features after the last average pooling layer in the BN-Inception network. The sampling rate of audio is 22,050 Hz. The raw waveform is converted to a spectrogram via a Short Time Fourier Transform (STFT). The hop size of the STFT is set to 256 with a window length of 1,024. We transform the spectrogram on the mel scale and set the number of mel frequency bins to 80. The first 220,160 points in the raw waveform are selected for the STFT, and thus, the time dimension of the spectrogram is 860.

**Training details.** The REGNET is implemented in PyTorch. The loss is illustrated in Equation (7) and we set  $\alpha$  and  $\beta$  to 1 and 10,000, respectively. We train the proposed REGNET for 1,000 epochs and the Adam optimization method [39] is adopted with an initial learning rate of 0.0002. Each category is independently trained. The output dimension  $D$  and downsampling rate  $S$  for the audio forwarding regularizer are set to 32 and 860, respectively, by default.

**Baseline.** Visual2Sound [2] is a state-of-the-art method for generating sound given visual input, which applies SampleRNN to directly generate a waveform based on the encoded visual feature. We consider it our baseline and use the generated results provided by the authors of [2] for performance comparisons.

**Evaluation metrics.** We design several metrics to evaluate the alignment and quality of the generated sound. We describe the details of each metric with its abbreviation. 1) *Sound-Missing*: the moment when an action (or event) may reveal

TABLE II: Results of the “real-or-fake” task. We show the percentages that indicate the frequency that a video is determined to be real.

	Real sound	Ours
Dog	85.54%	<b>71.09%</b>
Drum	86.71%	<b>67.04%</b>
Fireworks	91.01%	<b>86.33%</b>
Baby	91.33%	<b>72.59%</b>
Cough	82.03%	<b>55.98%</b>
Gun	91.19%	<b>70.13%</b>
Hammer	90.36%	<b>68.57%</b>
Sneeze	80.21%	<b>53.24%</b>
Average	88.65%	<b>68.12%</b>

sound but the model fails to generate it. 2) *Sound-Redundant*: the moment when no action (or event) may reveal sound occurs, but the model generates redundant sound. 3) *Sound-Mismatched*: the content of the generated sound is mismatched with the video content. 4) *Real-or-fake*: Whether the generating sound is able to fool the human into thinking that it is a real sound. For the first three metrics, the turkers are asked to select the video with **fewer** misaligned moments, which indicates better alignment performance. For the *real-or-fake* task, we present the percentage that indicates the frequency that a video is determined to be real. Each HIT is conducted by three turkers, and we aggregate their votes. The first two metrics and the third metric evaluate temporal alignment and content-wise alignment, respectively. The fourth metric measures general alignment and sound quality.

### C. Evaluation on alignment

**Setup.** A visually aligned sound should be temporally and content-wise aligned to the video. For temporal alignment, we use *Sound-Missing* and *Sound-Redundant* as evaluation metric; for content-wise alignment, we use *Sound-Mismatched* as evaluation metric. In the AMT test, the subjects are shown a paired video with the same visual content but different sounds, which are generated by REGNET and the baseline, respectively. The subjects are required to choose the video with the better performance on alignment. The experiment is conducted on *Sub-VEGAS* because the authors of [2] only provide the results in four video types in *Sub-VEGAS*.

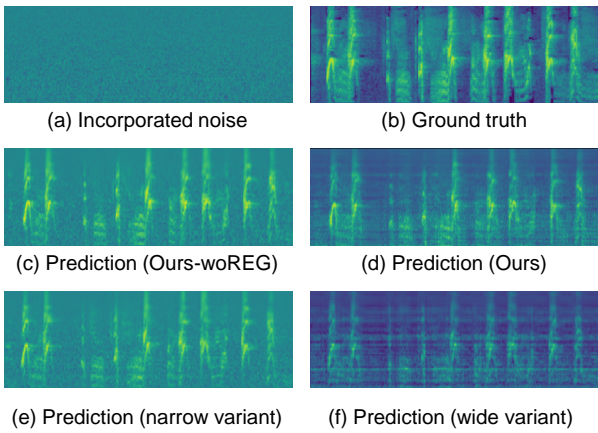


Fig. 5: Spectrograms in the experiment using Dog-gauss data. We mix the Gaussian noise (a) into the ground truth audio (b) to build the Dog-gauss simulated training sample. (c), (d), (e) and (f) are the predicted spectrograms from different models, in which the one from our REGNET (d) is able to get rid of the influence from the incorporated noise and preserve the sound architecture.

**Results.** We compare our method with baseline [2] and report the results in Table I. REGNET outperforms the baseline by 21.24% and 18.66% in terms of *Sound-Missing* and *Sound-Redundant*, respectively, which demonstrates that the proposed REGNET can generate sound with better temporal alignment. Specifically, approximately 60% of people think that the sound generated from our REGNET contains less missing sound moment and less redundant sound. This finding indicates that our REGNET can better capture the overlap between the visual information and the corresponding sound, even though the overlap is small. Besides, REGNET outperforms baseline [2] by 13.54% on *Sound-Mismatched*, which indicates that our REGNET learns a better mapping between visual content and sound. The sound generated by REGNET is better consistent with video content.

**Qualitative visualization.** We visualize the waveform of real sound with the sound generated by baseline [2] and our REGNET in Fig. 4. In the first case (a), where the fireworks are exploding, the baseline fails to generate any sound. The second case (b) shows that the baseline generates numerous redundant sounds when the dog is still. In the third case (c), where a man is hitting a drum, the baseline contains a metal hitting sound, which is not content-wise aligned. On the other hand, the sounds generated by REGNET are more plausible both in temporal alignment and content-wise alignment. We strongly encourage readers to view the demo videos in the supplementary materials.

#### D. Evaluation on real-or-fake task

**Setup.** We conduct a “real-or-fake” task to evaluate whether our REGNET can generate plausible sound. Specifically, we randomly combine a video with the generated sound or the real sound, and then require the turkers to make their own decisions based on the general alignment and sound quality.

TABLE III: L2 distances ( $\times 10^{-2}$ ) between the predicted and the ground truth spectrograms using simulated data.

Training data	Ours-woREG	Ours-VS	Ours
Dog-gauss	3.63	2.03	<b>1.96</b>
Dog-fireworks	3.02	<b>1.82</b>	1.87

TABLE IV: Performance comparisons between REGNET and the counterpart without audio forwarding regularizer in the “real-or-fake” task using simulated data.

Training data	Ours-woREG	Ours
Dog	68.75%	<b>71.09%</b> (+2.34%)
Dog-gauss	58.30%	<b>67.70%</b> (+9.40%)
Dog-fireworks	63.70%	<b>68.80%</b> (+5.10%)

**Results.** From Table II, 68% of sounds generated from REGNET success to fool human into thinking that these sounds are real, which is close to the success rate of the real sound itself. Note that the videos in the dataset have discriminative sound (as described in subsection V-A), such that the human is sensitive to the alignment of these sounds. Fooling human by the generated sound is difficult. To fool a human, the generated sound should be temporally and content-wise aligned with the visual content. Considering the Drum video for example, when the man is hitting a drum, if the model fails to generate any sound or generates a drum hitting sound soon at this time, the human may distinguish that the sound is fake. Also, if the generated sound does not sound similar to drum hitting, the human will easily perceive it.

#### E. Comparison with a simple method

As we train REGNET individually for each type of sound, one simple method for visually aligned sound generation is to detect the occurrence of specific actions and then play the predetermined audio clip. Specifically, we label a time step as an action occurrence if both the sound and optical flow amplitudes are higher than a threshold. We use the visual encoder described in section IV to encode visual features. Then, we apply a multi-layer perceptron to perform binary classification. During testing, we manually crop several audio clips with different lengths from training data. For each detected action occurrence, we play the predetermined audio clip with a suitable length.

The “real-or-fake” task on dog video shows that only 43.65% generated sounds can fool humans, which is significantly worse than REGNET (71.09%). This result demonstrates that it is non-trivial to generate plausible sound. One possible reason is that such a simple method can not handle diverse sounds in one video (e.g., dog barking from young and old dogs). Besides, it is hard to define action occurrences without annotations.

#### F. Effectiveness of audio forwarding regularizer

We aim to evaluate the performance of our REGNET with and without audio forwarding regularizer. We also evaluate the necessity of feeding both visual and audio information

TABLE V: Cosine similarity between the regularizer output from Dog-fireworks sound and other sounds. The data with the same background sound generate more similar regularizer output.

(Similar background)		(Different background)	
Baby-fireworks	Dog-fireworks	Dog-gauss	Baby-gauss
0.48	0.59	0.26	0.26

TABLE VI: L2 distances ( $\times 10^{-2}$ ) on spectrogram to visually relevant and irrelevant sounds when setting visual feature to zero and feeding input data to regularizer only. (The irrelevant sounds indicate Gaussian, fireworks and drum sounds, respectively for three different input data.)

Input data	Dog (relevant sound)	Gauss. or fireworks or drum (irrelevant sound)
Dog-gauss	4.36	<b>0.84</b>
Dog-fireworks	3.69	<b>1.97</b>
Dog-drum	4.53	<b>2.29</b>

to the audio forwarding regularizer to learn irrelevant sound information.

**Simulated data.** Based on the Dog data in *Sub-VEGAS* dataset, we mix extra noise (visually irrelevant sound) into the original corresponding audio. With more complicated visually irrelevant sound in the training sample, it is more difficult for the model to learn the correct mapping between visual contents and visually relevant sounds. Besides, in this case, the incorporated noise becomes the main source of visually irrelevant sound, and we can approximately regard the original corresponding audio as the ground-truth visually relevant sound. Specifically, we separately incorporate two kinds of noise into the original corresponding audio, including zero-mean Gaussian noise with a standard deviation of 0.1 and the sound from fireworks videos, which generates **Dog-gauss** and **Dog-fireworks** simulated data.

**Setup.** In addition to the proposed REGNET, we construct a counterpart in which we remove the audio forwarding regularizer in both the training phase and testing phase, and refer to it as **Ours-woREG**. Also, we input both visual feature and spectrogram into the audio forwarding regularizer to construct a variant, namely **Ours-VS**. We train these models using simulated data to evaluate whether our REGNET is robust to different kinds of visually irrelevant sound in the training samples. We report the success rate for a “real-or-fake” task because this metric is sensitive to the quality of the generated sound, especially the alignment and reality. Besides, considering the original audio in the simulated data as ground truth, we use the L2 distance to evaluate the performance of REGNET in generating visually relevant sound.

**Objective results.** In Table III, we show the L2 distance between the predicted spectrogram and the ground truth spectrogram in the experiments using simulated data. Note that we consider the original audio the ground truth. With two kinds of simulated data, the predicted spectrogram from our REGNET and Ours-VS is closer to the ground truth compared with Ours-woREG counterpart (1.96, 2.03 v.s. 3.63 and 1.87,

TABLE VII: L2 distances ( $\times 10^{-2}$ ) between the predicted and the background sound spectrograms when feeding into zero vector or regularizer output. (The regularizer output indicates the regularizer output from Dog-gauss, Dog-fireworks and Dog-drum, respectively when comparing with different background sound spectrogram.)

Background sound	Zero vector	Regularizer output
Gaussian	5.05	<b>2.72</b>
Fireworks	5.56	<b>2.78</b>
Drum	6.11	<b>2.79</b>

1.82 v.s. 3.02, respectively). This finding indicates that the audio forwarding regularizer enables REGNET to perform better in generating visually relevant sound. Also, it demonstrates that the audio forwarding regularizer does not need to take as input visual content explicitly. We argue that because the audio forwarding regularizer is trained for each type of video individually, the video category information is available for the network as one kind of visual information. Such category information is enough for regularizer to learn visually irrelevant sound. In Fig. 5, we visualize four spectrograms, including the incorporated Gaussian noise, original audio (ground truth), and predicted spectrograms from our REGNET and Ours-woREG counterpart. The Ours-woREG counterpart has learned some incorrect mappings between a visual feature and visually irrelevant noise, while our REGNET focuses only on generating the visually relevant component.

**Subjective results.** In Table IV, we show the performance of our REGNET and Ours-woREG counterpart in the “real-or-fake” task. Our REGNET consistently outperforms the Ours-woREG counterpart, and achieves performance gains of 2.34%, 9.40%, and 5.10% on three training data. This finding demonstrates that the audio forwarding regularizer is important for generating realistic sound. Note that the incorporated irrelevant sounds (*i.e.*, Gaussian noise and fireworks sound) dramatically deteriorate the performance and cause a 10.45% and 5.05% decrease in performance, while our REGNET is robust to the irrelevant sounds. One possible reason is that the proposed audio forwarding regularizer is helpful for eliminating the negative influence from irrelevant sounds in the training phase, which enables the model to learn proper mapping between visual features and visually relevant sounds.

### G. Exploration on regularizer output

We aim to explore what the audio forwarding regularizer has learned. Specifically, we will answer the following three questions: 1) is the regularizer output from spectrograms with similar background sounds more similar than those with different background sounds? 2) what information can be reconstructed from regularizer output? 3) can the generator mix in sounds from regularizer output during testing?

**Setup.** To answer these questions, we generate several types of simulated data, namely Dog-gauss, Dog-fireworks, Dog-drum, Baby-gauss, and Baby-fireworks. The term behind



TABLE VIII: Performance of REGNET variants. The L2 distance ( $\times 10^{-2}$ ) is evaluated on the Dog-gauss data, while the other metrics are evaluated on both *Sub-VEGAS* data and *Sub-AudioSet* data.

		L2 distance (Dog-gauss)	<i>real-or-fake</i>	<i>Sound-Missing</i> ( <i>Sub-VEGAS</i> and <i>Sub-AudioSet</i> )	<i>Sound-Redundant</i>	<i>Sound-Mismatched</i>
Narrow variant	$S_{860}D_8$	3.18	70.18%	3.16	3.01	2.95
Wide variant	$S_{860}D_{1024}$	3.19	62.50%	2.96	3.20	2.92
	$S_{32}D_{32}$	3.08	66.67%	2.94	3.22	2.93
REGNET	$S_{860}D_{32}$	<b>1.96</b>	<b>71.09%</b>	<b>3.20</b>	<b>3.35</b>	<b>3.02</b>

the hyphen (-) can be regarded as visually irrelevant sound. Each type of simulated data includes 128 videos for testing.

For the first question, we train four audio forwarding regularizers individually on Dog-fireworks, Dog-gauss, Baby-fireworks, and Baby-gauss data. Then, we input these four types of sound to the corresponding trained audio forwarding regularizer to generate four types of regularizer outputs. We use cosine distance to measure the similarity between regularizer outputs generated from different sounds.

For the second question, we study what information can the model reconstruct when it only takes regularizer output as input. Specifically, we train three REGNET individually on Dog-gauss, Dog-fireworks, and Dog-drum data. For each trained REGNET, we set the visual feature as zero and input the corresponding regularizer output to generate spectrogram. We measure the similarity between the generated spectrogram and visually relevant (or irrelevant) sound.

For the third question, we train a REGNET on dog videos and try to mix in background sounds (*i.e.*, Gaussian, fireworks, or drum) by replacing zero vector with different representations (*i.e.*, the regularizer output from Dog-gauss, Dog-fireworks, or Dog-drum). To measure whether the generated spectrogram include background sound, we compute the L2 distance between spectrogram of generated and background sounds.

**Results.** Table V shows the cosine similarity between regularizer output from Dog-fireworks and other simulated data. The data with fireworks background sound (*i.e.*, Baby-fireworks and Dog-fireworks) have higher similarity values compared with those with other background sounds. This indicates the regularizer passes through the background sound information ignoring the visually relevant sound. For the second question, in Table VI, the generated spectrogram have smaller L2 distances to visually irrelevant sound compared to relevant sound (0.84 *v.s.* 4.36 and 1.97 *v.s.* 3.69 and 2.29 *v.s.* 4.53). This indicates the regularizer is good at reconstructing the irrelevant sound component (*i.e.*, Gaussian, fireworks and drum sounds) compared to the relevant part (*i.e.*, dog sound). We believe the reason is that the regularizer output mainly contains irrelevant sound information instead of the relevant one. For the third question, the L2 distances between the spectrogram of generated and background sounds are shown in Table VII. When we feed in regularizer output instead of zero vector, the distances are relatively smaller (2.72 *v.s.* 5.05 and 2.78 *v.s.* 5.56 and 2.79 *v.s.* 6.11), indicating that the background sound information contained in regularizer output is successfully mixed in the output.

#### H. Exploration on REGNET variants

As intuitively explained in subsection III-C, with appropriate control for the capacity of the audio forwarding regularizer, the REGNET can predict visually relevant sound. Next, we will show how we control for the capacity of the audio forwarding regularizer and how it affects the prediction  $\hat{S}_0$ .

**REGNET variants.** We control for the capacity of the audio forwarding regularizer by changing its downsampling rate  $S$  and output dimension  $D$ . By default,  $S$  and  $D$  are set to 860 and 32 in our REGNET, and thus, we denote it as  $S_{860}D_{32}$ . With a larger output dimension or by sampling the regularizer output more densely, more information can be passed through the audio forwarding regularizer. Thus, we design two wide REGNET variants  $S_{860}D_{1024}$  and  $S_{32}D_{32}$ , whose audio forwarding regularizer has a larger output dimension of 1024 or a smaller downsampling rate of 32. Conversely, we decrease  $D$  to 8 and construct the narrow variant  $S_{860}D_8$ .

**Setup.** We train the REGNET variants with different capacities on Dog-gauss data and leverage the L2 distance to evaluate their performance in generating visually relevant sound. Besides, we conduct experiments on *Sub-VEGAS* and *Sub-AudioSet* and leverage the “real-or-fake” task with *Sound-Missing*, *Sound-Redundant* and *Sound-Mismatched* metrics to evaluate the quality and alignment of the generated sound from different REGNET variants. To make comparisons among several variants, the three alignment metrics are evaluated by the mean opinion score (MOS) test. Specifically, we provide two reference videos for the turkers and score them as 1 and 5. The higher score indicates better alignment. The turkers are asked to assign a score of 1-5 according to the corresponding requirements.

**Objective results.** In Table VIII, we show the L2 distance in the experiment with Dog-gauss data. Compared with our REGNET, both the “narrow variant” and the “wide variant” have a larger L2 distance between the prediction and the ground truth sound (1.96 *v.s.* 3.18, 3.19, 3.08). This finding indicates that the optimal results are achieved only when the capacity of the audio forwarding regularizer is appropriately controlled. We further visualize the predicted spectrogram  $\hat{S}_0$  in Fig. 5. For the “wide variant”, when the audio forwarding regularizer is shut down, the generator misses some of the meaningful spectrogram structure. This is not surprising because the audio forwarding regularizer with a large capacity is able to provide sufficient information about the target spectrogram during the training phase. The generator tends to use this information for spectrogram reconstruction instead of learning mapping between a visual feature and visually relevant sound.

For the “narrow variant”, the predicted spectrogram contains considerable noise because the model is forced to learn some incorrect mapping between a visual feature and noise.

**Subjective results.** In Table VIII, our REGNET performs the best for both the “real-or-fake” task and three alignment metrics. Specifically, 71.09% generated sounds from REGNET are considered real sounds, compared with 62.50% and 66.67% generated sounds for the “wide variant” and 70.18% for the “narrow variant”. Our REGNET also achieves the highest scores with all three alignment metrics. Note that the “wide variant” achieves low scores for the *Sound-Missing* term, which demonstrates that the model fails to generate sound when action occurs. The reason is that the audio forwarding regularizer with excessively large capacity prevents the model from learning necessary mapping between a visual feature and sound.

### I. Exploration on discriminator

As there are multiple possible sounds that are visually aligned with the visual content, L2 reconstruction loss can only supervise the network to learn an overall sound structure. This may result in a smooth spectrogram and low-fidelity results. Adversarial training mechanism has been proven to be effective to generate high-fidelity and coherent audio in several works [33], [32], [5]. To study the necessity of discriminator in our REGNET, we train a REGNET without adversarial loss and conduct a “real-or-fake” human study on *Sub-VEGAS* and *Sub-AudioSet* datasets. Ours-woGAN fools 64.04% human on average, which is worse than REGNET (68.12%). This result demonstrates that the discriminator is helpful to generate high-fidelity sounds in REGNET.

## VI. CONCLUSION

We explain from an information mismatch perspective why the temporal and content-wise alignment remains challenging for generating sound from videos, and propose an audio forwarding regularizer to solve this challenge. With the audio forwarding regularizer, our REGNET can leverage sufficient information for prediction during the training phase and learn a concrete correspondence between visual information and sound information. Experiments show that the sound generated by our REGNET is more plausible and more aligned to video.

## REFERENCES

- [1] K. Chen, C. Zhang, C. Fang, Z. Wang, T. Bui, and R. Nevatia, “Visually indicated sound generation by perceptually optimized classification,” in *The European Conference on Computer Vision*, 2018, pp. 560–574.
- [2] Y. Zhou, Z. Wang, C. Fang, T. Bui, and T. L. Berg, “Visual to sound: Generating natural sound for videos in the wild,” in *The IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 3550–3558.
- [3] I. J. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. C. Courville, and Y. Bengio, “Generative adversarial nets,” in *Advances in Neural Information Processing Systems*, 2014, pp. 2672–2680.
- [4] A. Owens, P. Isola, J. McDermott, A. Torralba, E. H. Adelson, and W. T. Freeman, “Visually indicated sounds,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 2405–2413.
- [5] L. Chen, S. Srivastava, Z. Duan, and C. Xu, “Deep cross-modal audio-visual generation,” in *Proceedings of the on Thematic Workshops of ACM Multimedia*, 2017, pp. 349–357.
- [6] A. Owens, J. Wu, J. H. McDermott, W. T. Freeman, and A. Torralba, “Ambient sound provides supervision for visual learning,” in *The European Conference on Computer Vision*, 2016, pp. 801–816.
- [7] G. Monaci, P. Jost, P. Vanderheynt, B. Mailhe, S. Lesage, and R. Gribonval, “Learning multimodal dictionaries,” *IEEE Transactions on Image Processing*, vol. 16, no. 9, pp. 2272–2283, 2007.
- [8] D. F. Harwath, A. Torralba, and J. R. Glass, “Unsupervised learning of spoken language with visual context,” in *Advances in Neural Information Processing Systems*, 2016, pp. 1858–1866.
- [9] R. Arandjelovic and A. Zisserman, “Look, listen and learn,” in *The IEEE International Conference on Computer Vision*, 2017, pp. 609–617.
- [10] Y. Aytar, C. Vondrick, and A. Torralba, “Soundnet: Learning sound representations from unlabeled video,” in *Advances in Neural Information Processing Systems*, 2016, pp. 892–900.
- [11] H. Zhao, C. Gan, A. Rouditchenko, C. Vondrick, J. H. McDermott, and A. Torralba, “The sound of pixels,” in *European conference on computer vision*, 2018, pp. 587–604.
- [12] H. Zhao, C. Gan, W. Ma, and A. Torralba, “The sound of motions,” in *IEEE International Conference on Computer Vision*, 2019, pp. 1735–1744.
- [13] C. Gan, D. Huang, H. Zhao, J. B. Tenenbaum, and A. Torralba, “Music gesture for visual sound separation,” in *IEEE Conference on Computer Vision and Pattern Recognition*, 2020.
- [14] C. Gan, H. Zhao, P. Chen, D. Cox, and A. Torralba, “Self-supervised moving vehicle tracking with stereo sound,” in *IEEE International Conference on Computer Vision*, 2019, pp. 7052–7061.
- [15] A. Rouditchenko, H. Zhao, C. Gan, J. H. McDermott, and A. Torralba, “Self-supervised audio-visual co-segmentation,” in *IEEE International Conference on Acoustics, Speech and Signal Processing*, 2019, pp. 2357–2361.
- [16] C. Gan, Y. Zhang, J. Wu, B. Gong, and J. B. Tenenbaum, “Look, listen, and act: Towards audio-visual embodied navigation,” *IEEE International Conference on Robotics and Automation*, 2020.
- [17] Z. Ling, S. Kang, H. Zen, A. W. Senior, M. Schuster, X. Qian, H. M. Meng, and L. Deng, “Deep learning for acoustic modeling in parametric speech generation: A systematic review of existing techniques and future trends,” *IEEE Signal Processing Magazine*, vol. 32, no. 3, pp. 35–52, 2015.
- [18] K. van den Doel, P. G. Kry, and D. K. Pai, “FoleyAutomatic: physically-based sound effects for interactive simulation and animation,” in *Proceedings of the Annual Conference on Computer Graphics and Interactive Techniques*, 2001, pp. 537–544.
- [19] A. van den Oord, S. Dieleman, H. Zen, K. Simonyan, O. Vinyals, A. Graves, N. Kalchbrenner, A. W. Senior, and K. Kavukcuoglu, “Wavenet: A generative model for raw audio,” in *ISCA Speech Synthesis Workshop*, 2016, p. 125.
- [20] S. Mehri, K. Kumar, I. Gulrajani, R. Kumar, S. Jain, J. Sotelo, A. C. Courville, and Y. Bengio, “SampleRNN: An unconditional end-to-end neural audio generation model,” in *International Conference on Learning Representations*, 2017.
- [21] J. Sotelo, S. Mehri, K. Kumar, J. F. Santos, K. Kastner, A. C. Courville, and Y. Bengio, “Char2Wav: End-to-end speech synthesis,” in *International Conference on Learning Representations*, 2017.
- [22] J. Shen, R. Pang, R. J. Weiss, M. Schuster, N. Jaitly, Z. Yang, Z. Chen, Y. Zhang, Y. Wang, R. Ryan, R. A. Saurous, Y. Agiomyrgiannakis, and Y. Wu, “Natural TTS synthesis by conditioning wavenet on MEL spectrogram predictions,” in *IEEE International Conference on Acoustics, Speech and Signal Processing*, 2018, pp. 4779–4783.
- [23] F. Yu, X. Wu, J. Chen, and L. Duan, “Exploiting images for video recognition: Heterogeneous feature augmentation via symmetric adversarial learning,” *IEEE Transactions on Image Processing*, vol. 28, no. 11, pp. 5308–5321, 2019.
- [24] Y. Guo, Q. Chen, J. Chen, Q. Wu, Q. Shi, and M. Tan, “Auto-embedding generative adversarial networks for high resolution image synthesis,” *IEEE Transactions on Multimedia*, vol. 21, pp. 2726–2737, 2019.
- [25] X. Chen, C. Xu, X. Yang, L. Song, and D. Tao, “Gated-GAN: Adversarial gated networks for multi-collection style transfer,” *IEEE Transactions on Image Processing*, vol. 28, no. 2, pp. 546–560, 2019.
- [26] C. Wang, C. Xu, C. Wang, and D. Tao, “Perceptual adversarial networks for image-to-image transformation,” *IEEE Transactions on Image Processing*, vol. 27, no. 8, pp. 4066–4079, 2018.
- [27] C. Hsu, C. Lin, W. Su, and G. Cheung, “SiGAN: Siamese generative adversarial network for identity-preserving face hallucination,” *IEEE Transactions on Image Processing*, vol. 28, no. 12, pp. 6225–6236, 2019.
- [28] M. Mirza and S. Osindero, “Conditional generative adversarial nets,” *arXiv preprint arXiv:1411.1784*, 2014.

- [29] A. Radford, L. Metz, and S. Chintala, "Unsupervised representation learning with deep convolutional generative adversarial networks," in *International Conference on Learning Representations*, 2016.
- [30] M. Arjovsky, S. Chintala, and L. Bottou, "Wasserstein generative adversarial networks," in *Proceedings of the International Conference on Machine Learning*, 2017, pp. 214–223.
- [31] S. E. Reed, Z. Akata, X. Yan, L. Logeswaran, B. Schiele, and H. Lee, "Generative adversarial text to image synthesis," in *Proceedings of the International Conference on Machine Learning*, 2016, pp. 1060–1069.
- [32] C. Donahue, J. McAuley, and M. Puckette, "Adversarial audio synthesis," in *International Conference on Learning Representations*, 2019.
- [33] J. Engel, K. K. Agrawal, S. Chen, I. Gulrajani, C. Donahue, and A. Roberts, "GANSynth: Adversarial neural audio synthesis," in *International Conference on Learning Representations*, 2019.
- [34] S. Ioffe and C. Szegedy, "Batch normalization: Accelerating deep network training by reducing internal covariate shift," in *International Conference on Learning Representations*, 2015, pp. 448–456.
- [35] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural computation*, vol. 9, no. 8, pp. 1735–1780, 1997.
- [36] V. Nair and G. E. Hinton, "Rectified linear units improve restricted boltzmann machines," in *ICML*, 2010, pp. 807–814.
- [37] P. Isola, J. Zhu, T. Zhou, and A. A. Efros, "Image-to-image translation with conditional adversarial networks," in *The IEEE Conference on Computer Vision and Pattern Recognition*, 2017, pp. 5967–5976.
- [38] J. F. Gemmeke, D. P. Ellis, D. Freedman, A. Jansen, W. Lawrence, R. C. Moore, M. Plakal, and M. Ritter, "Audio set: An ontology and human-labeled dataset for audio events," in *The IEEE International Conference on Acoustics, Speech and Signal Processing*, 2017, pp. 776–780.
- [39] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *International Conference on Learning Representations*, 2015.



**Hongdong Xiao** received the B.E. degree in Computer Science and Technology from Nanchang University, China, in 2018. He is currently pursuing the M.E. degree at the School of Software Engineering, South China University of Technology, China. His research interests include deep learning in audio understanding.




**Deng Huang** received the B.E. degree in software engineering from South China University of Technology in 2019. He is currently pursuing the M.E. degree with South China University of Technology. His research interests include computer vision and deep learning.



**Peihao Chen** received the B.E. degree in Automation Science and Engineering from South China University of Technology, China, in 2018. He is working toward the M.E. degree in the School of Software Engineering, South China University of Technology, China. His research interests include Deep Learning in Video and Audio Understanding.



**Chuang Gan** is currently a Researcher with the MIT-IBM Watson AI Lab. His research interests mainly include multi-modality learning for video understanding.



**Yang Zhang** received his M.S. and Ph.D. degrees at the University of Illinois at Urbana-Champaign. He is currently a researcher with the MIT-IBM Watson AI Lab. His research interest includes deep learning for speech, audio, and time-series processing.



**Mingkui Tan** is currently a professor with the School of Software Engineering at South China University of Technology. He received his Bachelor Degree in Environmental Science and Engineering in 2006 and Master degree in Control Science and Engineering in 2009, both from Hunan University in Changsha, China. He received the Ph.D. degree in Computer Science from Nanyang Technological University, Singapore, in 2014. From 2014–2016, he worked as a Senior Research Associate on computer vision in the School of Computer Science, University of

Adelaide, Australia. His research interests include machine learning, sparse analysis, deep learning and large-scale optimization.