

DETECT WHAT YOU WANT: TARGET SOUND DETECTION

Helin Wang^{1†}, Dongchao Yang^{1†}, Yuexian Zou^{1*}, Fan Cui², Yujun Wang²

¹ADSPLAB, School of ECE, Peking University, Shenzhen, China

² Xiaomi Corporation, Beijing, China

ABSTRACT

Human beings can perceive a target sound type from a multi-source mixture signal by the selective auditory attention, however, such functionality was hardly ever explored in machine hearing. This paper addresses the target sound detection (TSD) task, which aims to detect the target sound signal from a mixture audio when a target sound’s reference audio is given. We present a novel target sound detection network (TSDNet) which consists of two main parts: A conditional network which aims at generating a sound-discriminative conditional embedding vector representing the target sound, and a detection network which takes both the mixture audio and the conditional embedding vector as inputs and produces the detection result of the target sound. These two networks can be jointly optimized with a multi-task learning approach to further improve the performance. In addition, we study both strong-supervised and weakly-supervised strategies to train TSDNet and propose a data augmentation method by mixing two samples. To facilitate this research, we build a target sound detection dataset (*i.e.* URBAN-TSD) based on URBAN-SED and UrbanSound8K datasets, and experimental results indicate our method could get the segment-based F scores of 76.3% and 56.8% on the strongly-labelled and weakly-labelled data respectively.

Index Terms— target sound detection, conditional embedding, weakly supervised, data augmentation

1. INTRODUCTION

Human beings has the ability to focus auditory attention on a particular sound in a multi-source environment, however, there were few studies in this area for machine hearing. In this paper, we initially define a target sound detection (TSD) task, which aims to recognize and localize target sound source within a mixture audio given a reference audio or/and a sound label. For example, the violin sound can be detected within a concert recording and the talking sound can be detected in a noisy cafe environment. TSD has a lots of potential applications, such as noise monitoring for smart cities [1], bioacoustic species and migration monitoring [2] and large-scale multimedia indexing [3]. To the best of our knowledge, this paper is the first attempt that explicitly tackles this problem.

There is one similar task with TSD, *i.e.* sound event detection (SED). SED aims to classify and localize all pre-defined sound events (*e.g.*, train horn, car alarm) within an audio clip, which has

been widely studied [4, 5, 6, 7]. Compared to SED, TSD only focuses on detecting the event that we care about and ignores other events. Furthermore, TSD does not require to pre-define categories set, so it can be easy to extend to open domain detection. Other related tasks about extracting the target signal are speaker extraction [8, 9] which extracts the target speech from a mixture speech given a reference utterance of the target speaker, and acoustic events sound selection problems [10]. Different from them, our work focuses on the detection task, which is more suitable for many multimedia retrieval applications and the training data is easier to obtain.

To solve the TSD task, we propose a target sound detection network (*i.e.* TSDNet), and treat TSD as a binary classification problem for each frame of the audio, where the positive class is the sound event of interest, and the negative class is formed by the combination of all foreground and background interfering events and noises. By using reference audio, TSDNet can focus on the target sound and ignore other interference. More specifically, TSDNet is composed of a conditional network which is used to generate a sound-discriminative conditional embedding vector from the reference audio, and a detection network which is applied to obtain binary-classification results at each frame with the conditional embedding vector and mixture audio as inputs. In order to get more sound-discriminative conditional embedding vector for TSDNet, the conditional network is jointly optimized with both the TSD task and a sound event classification task. We further explore a data augmentation method for TSD, which randomly mixes two training samples to form a new training sample. In addition, we explore to train the TSDNet on both strong- and weakly-supervised TSD tasks. Here, weakly-supervised TSD task means the dataset only provides the presence or absence of target sound within the mixture sound, but not any timestamp information, which is more challenging.

Our contributions can be summarized as follows: (1) We propose a novel network for TSD that can be trained with both strong-labelled and weakly-labelled data. (2) We explore the jointly training method to get the robust conditional embedding vector, and propose a data augmentation method for the TSD task. (3) We establish a dataset for TSD, and our method achieves the segment-based F measures of 76.3% and 56.8% on the strongly-labelled and weakly-labelled data respectively.

2. PROPOSED METHOD

The architecture of our proposed network (TSDNet) is shown in Figure 1. TSDNet consists of two components: a conditional network which inputs the reference audio and outputs a conditional embedding vector, and a detection network which uses the conditional embedding vector and the mixture audio to get the detection results. In this section, we will describe the details of the whole network.

[†] Indicates equal contribution.

* Corresponding Author: zouyx@pku.edu.cn

This paper was partially supported by Shenzhen Science & Technology Research Program (No: GXWD20201231165807007-20200814115301001; No: JSGG20191129105421211) and NSFC (No: 62176008).




Figure 1: The architecture of our proposed TSDNet. Here, FC denotes the fully-connected layer. Reference signal indicates that the signal from the target class.

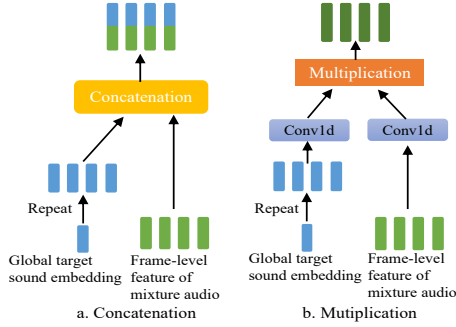


Figure 2: Two fusion modules.

2.1. Conditional Network

The purpose of the conditional network is to produce a global conditional embedding vector to represent the reference information. The input to the conditional network can be either a reference audio or a specific label, or both. In this paper, we focus on a reference audio as the input, for the reason that it is more challenging but easier to be transferred to new classes. Inspired by the powerful ability of extracting time-frequency robust features from audio with convolutional neural networks (CNNs) [11, 12, 13, 14], we apply a VGG-like CNN network for the conditional network, which uses the spectrogram feature as input and consists of 4 convolutional blocks with 64, 128, 256 and 512 output channels, respectively. Each convolutional block contains 2 convolutional layers with kernel size of 3×3 , followed by downsampling with average pooling size of 2×2 . Batch normalization [15] and ReLU function [16] are applied to all the convolutional layers. Global pooling layer [11] containing a global max-pooling function and a global average-pooling function is applied to get the global feature vector, which is then fed to a fully-connected layer to get the global conditional embedding vector with a fixed dimension of 128.

2.2. Detection network

We have built two types of detection network: strong- and weakly-supervised network. The details are given as follows.

Strong-supervised network The strong-supervised detection network is based on the state-of-the-art work for weakly supervised sound event detection by Dinkel *et al.* [4]. As shown in Figure 1, the network has two inputs: the conditional embedding vector and the mixture audio, which is trained to minimize the difference between the frame-level prediction results and the ground-truth labels. To be more specific, the network is composed of 4 convolutional layers, 1 Bi-GRU layer, and 2 fully-connected layers, each with a

LeakyReLU activation except the last layer, which has a sigmoid activation. Given the input feature of the mixture audio $\mathbf{x} \in \mathcal{R}^{t \times f}$, where t and f denote the number of frames and the dimension of each frame respectively, the network aims to predict frame-level probabilities $\hat{p}_i = \mathbb{P}(Y = k | X = x_i, \mathbf{e}; \phi)$ where ϕ denotes the trainable parameters of detection network, \mathbf{e} denotes the embedding from the conditional network and x_i denotes the i -th frame of the mixture audio \mathbf{x} . Here, the value of k is 0 or 1. Given the ground-truth label $p_i \in \{0, 1\}$ for each frame, the strong-supervised network can be optimized by minimize the binary cross-entropy (BCE) loss function:

$$\mathcal{L}_{sed} = \sum_{i=1}^t (-p_i \log \hat{p}_i - (1 - p_i) \log(1 - \hat{p}_i)) \quad (1)$$

where t indicates the number of frames.

Weakly-supervised network The difference between strong-supervised network and weakly-supervised network is that the latter needs a pooling layer to get the clip-level prediction. We add a linear softmax (LinSoft) pooling layer [17] after the last layer of the strong-supervised detection network. It aims to predict a clip-level probability $\hat{P} = f_{LSP}(\hat{p}_1, \hat{p}_2, \dots, \hat{p}_t)$ where $f_{LSP}(\cdot)$ denotes the linear softmax pooling function and T denotes the number of frames. Given the clip-level ground-truth label $P \in \{0, 1\}$, BCE loss is also applied as the loss function:

$$\mathcal{L}'_{sed} = -P \log \hat{P} - (1 - P) \log(1 - \hat{P}) \quad (2)$$

2.3. Fusion Module

As Figure 2 shows, we employ two fusion strategies to combine the conditional embedding and the feature of mixture audio. One is repeatedly concatenating the conditional embedding to the feature of mixture audio in each time frame. The other is projecting the conditional embedding and the feature of mixture audio to the same dimension by a 1-D convolutional layer, and then using multiplication operation to fuse them.

2.4. Mixup-TSD

Following the advanced data augmentation methods [18, 19, 20] for audio classification, we propose a data augmentation method based on the widely-used mixup [18] for TSD task. The core idea is to create a new training sample by mixing a pair of two training samples. More specifically, we can generate a new training sample $(M_{new}, R_{new}, y_{new})$ from the data and label pair (M_1, R_1, y_1) and (M_2, R_2, y_2) by the following equation.

$$X_{new} = \lambda X_1 + (1 - \lambda X_2) \quad (3)$$

$$R_{new} = \lambda R_1 + (1 - \lambda R_2) \quad (4)$$

$$y_{new} = \lambda y_1 + (1 - \lambda y_2) \quad (5)$$

where M_1 and M_2 are mixture audios, R_1 and R_2 are reference audios, y_1 and y_2 are the corresponding labels. Such Mixup-TSD works by generating lots of new training samples, and particularly, the target sound turns to be multi-label event instead of single-label event. We argue that this method can perform well at the early training stage, but needs a fine-tuning stage to fit for the single-label test. In practice, we set a ratio α to control whether to use Mixup-TSD while training, which means Mixup-TSD is applied with a probability of α . We set a linear decay for α from 0.3 to 0 during the whole training iterations.

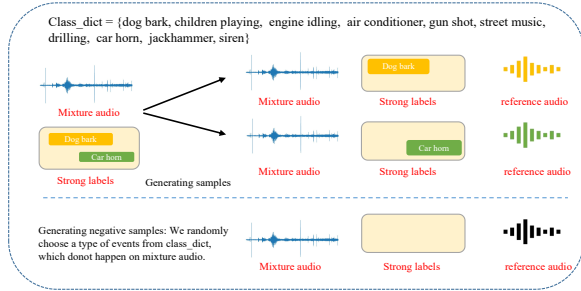


Figure 3: The process of data generation.

2.5. Jointly Training

To get a more sound-discriminative conditional embedding vector for TSDNet, we propose a jointly optimizing method for the conditional network via a multi-task manner. As show in Figure 1, the feature vector after the global pooling layer is fed to a fully-connected layer and a softmax nonlinearity to get the classification results. The conditional network can be optimized by minimizing the cross entropy loss between the predicted label and the ground truth label of the reference audio, along with the original detection task. The whole loss function is defined by formula (6).

$$\mathcal{L}_{total} = \mathcal{L}_{sed} + \mathcal{L}_{cls} \quad (6)$$

where \mathcal{L}_{sed} denotes the loss of the detection task, and \mathcal{L}_{cls} denotes the loss of the classification task, which is the cross-entropy loss between the ground-truth class labels and the predicted class labels.

3. URBAN-TSD DATASET

As there was no dataset for TSD task, we build a dataset called URBAN-TSD based on URBAN-SED [21] and UrbanSound8K datasets [22]. URBAN-SED is a sound event detection dataset with an urban setting, containing 10 event labels. This dataset’s source material is the UrbanSound8K dataset containing 27.8 hours of data split into about 4-second clip segments. The URBAN-SED dataset encompasses 10,000 soundscapes generated using the Sca-per soundscape synthesis library [21], which have been split into 6000 training, 2000 validation and 2000 test clips. The dataset contains mostly 10-second excerpts with strong labels, whereas each clip contains between one and nine events. We establish two types of dataset: strong-labelled dataset (URBAN-TSD-strong) and weakly-labelled dataset (URBAN-TSD-weak). The details of the number of samples are shown in Table 1.

3.1. URBAN-TSD-strong

We have built two strong datasets: URBAN-TSD-strong and URBAN-TSD-strong+. The details are given as follows.

URBAN-TSD-strong dataset: In this dataset, a sample includes three parts: a mixture audio, a reference audio and the strong label (the timestamp information of target sound). Mixture audios come from URBAN-SED dataset, and reference audio comes from UrbanSound8K. As shown in Figure 3, if there are N sound events in the mixture audio, we can generate N positive samples. For each positive sample, we randomly choose another sample that is in the same class from the UrbanSound8K as the reference audio.

URBAN-TSD-strong+ dataset: To further verify the ability of the model facing all negative frames, we generate samples that do not

Table 1: The number of strong- and weakly-labelled data.

Type	Strong	Strong+	Weak
Training	23106	29106	41059
Validation	7681	9681	13661
Test	7702	9702	13682

Table 2: The performance comparison between TSDNet and other SED methods. Note that multiplication fusion is applied.

Method	Segment-based F1
CRNN [23]	64.7
CDur [4]	64.8
CTrans [24]	64.51
SEDT-AQ-FT-P3 [25]	65.77
Ours	73.1

contain the target sound, which are called as negative samples. We add those samples because the mixture audio may not contain our target sound in the real world. The process of generating negative samples can summarize as: For any mixture audio from URBAN-SED, we randomly choose a reference audio whose sound events do not happen in the mixture audio. The timestamp information (label) of negative samples are all set as 0.

3.2. URBAN-TSD-weak

We also build a weakly-labelled dataset which is easier to obtain but more challenging. Comparing with strong-labelled dataset, there is no onset or offset time provided. The label is set as 1 (which indicates target sound happens in the mixture audio) or 0 (which indicates target sound does not happen in the mixture audio) for the whole audio clip. Similarly, for any mixture audio from URBAN-SED, if N events happen, we can generate N positive samples. For weakly-labelled dataset, negative samples are very important for training, and we generate the same number of negative samples as positive samples.

4. EXPERIMENTS

4.1. Experimental setups

Training the conditional network: We pre-train the conditional network with the classification task on the UrbanSound8K dataset. All the raw audios are down-sampled to 44.1kHz and applied a Short Time Fourier Transform (STFT) with a window size of 400 samples, followed by a Mel-scaled filter bank on perceptually weighted spectrogram. After that, we get log-mel spectrogram and MFCC feature, and concatenate them as the input, which is the state-of-the-art solution¹. The Adam optimizer [26] is used for a total of 50 epochs, with an initial learning rate of 1×10^{-3} .

Training the detection network: All the raw audios are down-sampled to 22.05kHz and applied a Short Time Fourier Transform (STFT) with a window size of 2048 samples, followed by a Mel-scaled filter bank on perceptually weighted spectrogram. This results in 64 Mel frequency bins and around 50 frames per second. The Adam optimizer [26] is used for 100 epochs, with an initial learning rate of 1×10^{-3} . Note that we only update the detection

¹<https://www.kaggle.com/adinishad/urbansound-classification-with-pytorch-and-fun/notebook>

Table 3: The segment-based F-measure (%) of TSDNet with different fusion strategies. We carry out the experiments for three times and report the mean and standard deviation values.

Method	strong	strong+	weak
Concatenation	73.1±0.61	67.2±0.41	52.8±0.53
Multiplication	73.1±0.19	69.3±0.23	53.0±0.32

Table 4: The segment-based F-measure (%) of different settings for Mixup-TSD. Here, multiplication fusion is applied.

Mixup-TSD setting	strong	strong+	weak
fixed ratio 0	73.1	69.3	53.0
fixed ratio 0.2	74.9	71.1	55.2
fixed ratio 0.5	74.2	70.8	54.7
fixed ratio 0.8	74.0	70.4	53.6
fixed ratio 1.0	73.8	70.2	53.2
linear decay ratio	75.1	71.6	55.8

network with the detection loss.

Mixup-TSD: For all the experiments, we use the Mixup-TSD method on the spectrogram level as we find spectrogram-level method works better than waveform-level method. Under otherwise stated, the linear decay is set for α from 0.3 to 0.

Metrics: We use the segment-based F-measure [27] as the evaluation metric, which is the most commonly used metrics for sound event detection. All the F-scores are macro-averaged.

4.2. The performance comparison between TSDNet and other SED methods

In this part, we conduct experiments on URBAN-TSD-strong dataset to validate the effectiveness of our proposed method (TSDNet). We choose four previous state-of-the-art SED methods on URBAN-SED dataset. As for the baselines, we follow the model architecture in [24] to build the the Transformer-based model, which is referred to as CTrans in this paper. CDur [4] is the backbone of our TSDNet’s detection network. Table 2 reports the results. We can see that our TSDNet significantly improve the performance due to introducing the reference audio.

4.3. Experimental results of different fusion strategies

We evaluate our baseline method with different fusion strategies. Table 3 reports the segment-based F-measure on the three datasets. The multiplication gets larger mean values and smaller standard deviation values in all situations, which is a more effective fusion method than the concatenation. In addition, TSDNet performs the worst on the URBAN-TSD-weak dataset, for the reason that no timestamp information is available during training.

4.4. Experimental results of Mixup-TSD

We further evaluate the Mixup-TSD method on three datasets, and report the performance in Table 4. If no augmentation is used, TSDNet can obtain the segment-based F-measures of 73.1%, 69.3% and 53.0% respectively. We can see that the Mixup-TSD method can significantly improve the performance over the baseline under all settings. More specifically, for a fixed ratio α , the performance tends to decrease with the ratio increasing, and the best performance

Table 5: The performance of jointly training. Note that multiplication fusion and Mixup-TSD are applied.

Jointly training	Strong	Strong+	Weak
×	73.1	69.3	53.0
✓	76.3	72.2	56.8

Table 6: Experimental results of open domain. Note that multiplication fusion is applied.

Jackhammer	Siren	Street_music	Average
40.9	42.2	48.0	43.7

can be achieved when the ratio is 0.2. We argue that the Mixup-TSD method can perform well at the early training stage, but it will change the detection target from single-label event to multi-label event so we need a fine-tuning stage to fit for the single-label test. With a linear decay ratio from 0.3 to 0, we can get the segment-based F-measures of 75.1%, 71.6% and 55.8% respectively.

4.5. Experimental results of jointly training

In addition, we explore whether jointly training can further improve the performance of target sound detection. Firstly, we initialize the TSDNet with a pre-trained model which is only trained by the detection task. Secondly, we fine-tune TSDNet by a multi-task manner introduced in Section 2.5. We set a learning rate of 1×10^{-4} for the whole model, and the Adam optimizer [26] is used for a total of 30 epochs. Experimental results are shown in Table 5. We can see that TSDNet with jointly training obtains 76.3%, 72.2% and 56.8% respectively, which brings 4.3%, 4.2% and 7.1% improvement over the baseline respectively. We find that the weak-supervised training still performs much worse than the strong-supervised training. We will study more effective weak-supervised methods in the future.

4.6. Open domain target sound detection

Table 6 shows the results of TSDNet evaluated on the open domain. We choose three events (jackhammer, siren, street_music) as new classes, which does not occur in the training set. Specifically, we first exclude all of the data include the three events in the URBAN-TSD-strong dataset, named URBAN-TSD-strong-7. Then we train TSDNet on URBAN-TSD-strong-7. We can see that TSDNet can be applied to the new class, *e.g.*, the F-score of street_music achieved 48.0% and the average F-score of the three new classes achieved 43.7%. We believe that adding more training data can further improve the performance of new class.

5. CONCLUSIONS

In this paper, we proposed a target sound detection network (TSDNet) which can be trained with both strong-supervised learning and weakly-supervised learning manners. In addition, our proposed Mixup-TSD data augmentation method and jointly training strategy further improve the performance of TSDNet. TSDNet is more applicable to real scenarios because it does not require prior knowledge about the number of sound events and easy to extend to the new events. In the future, we will explore more robust system. The source code and datasets of this work are released.²

²<https://github.com/yangdongchao/Target-sound-event-detection>

6. REFERENCES

- [1] J. P. Bello, C. Silva, O. Nov, R. DuBois, A. Arora, J. Salamon, C. Mydlarz, and H. Doraiswamy, “Sonyc: A system for the monitoring, analysis and mitigation of urban noise pollution,” *arXiv preprint arXiv:1805.00889*, 2018.
- [2] D. Stowell and D. Clayton, “Acoustic event detection for multiple overlapping similar sources,” in *2015 IEEE Workshop on Applications of Signal Processing to Audio and Acoustics (WASPAA)*. IEEE, 2015, pp. 1–5.
- [3] S. Hershey, S. Chaudhuri, D. Ellis, J. Gemmeke, A. Jansen, R. C. Moore, M. Plakal, D. Platt, R. A. Saurous, B. Seybold, et al., “CNN architectures for large-scale audio classification,” in *IEEE international conference on acoustics, speech and signal processing (ICASSP)*. IEEE, 2017, pp. 131–135.
- [4] H. Dinkel, M. Wu, and K. Yu, “Towards duration robust weakly supervised sound event detection,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 29, pp. 887–900, 2021.
- [5] L. Lin, X. Wang, H. Liu, and Y. Qian, “Specialized decision surface and disentangled feature for weakly-supervised polyphonic sound event detection,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 28, pp. 1466–1478, 2020.
- [6] Q. Kong, Y. Xu, W. Wang, and M. D. Plumbley, “Sound event detection of weakly labelled data with cnn-transformer and automatic threshold optimization,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 28, pp. 2450–2460, 2020.
- [7] Q. Kong, C. Yu, Y. Xu, T. Iqbal, W. Wang, and M. D. Plumbley, “Weakly labelled audioset tagging with attention neural networks,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 27, no. 11, pp. 1791–1802, 2019.
- [8] Q. Wang, H. Muckenhirn, K. Wilson, P. Sridhar, Z. Wu, J. R. Hershey, R. A. Saurous, R. J. Weiss, Y. Jia, and I. L. Moreno, “Voicefilter: Targeted voice separation by speaker-conditioned spectrogram masking,” *Proc. Interspeech*, pp. 2728–2732, 2019.
- [9] M. Ge, C. Xu, L. Wang, E. S. Chng, J. Dang, and H. Li, “Multi-stage speaker extraction with utterance and frame-level reference signals,” in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2021, pp. 6109–6113.
- [10] T. Ochiai, M. Delcroix, Y. Koizumi, H. Ito, K. Kinoshita, and S. Araki, “Listen to what you want: Neural network-based universal sound selector,” *Proc. Interspeech*, pp. 1441–1445, 2020.
- [11] Q. Kong, Y. Cao, T. Iqbal, Y. Wang, W. Wang, and M. Plumbley, “Panns: Large-scale pretrained audio neural networks for audio pattern recognition,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 28, pp. 2880–2894, 2020.
- [12] H. Wang, Y. Zou, D. Chong, and W. Wang, “Environmental sound classification with parallel temporal-spectral attention,” in *Proc. Interspeech*, 2020, pp. 821–825.
- [13] H. Wang, D. Chong, D. Huang, and Y. Zou, “What affects the performance of convolutional neural networks for audio event classification,” in *International Conference on Affective Computing and Intelligent Interaction Workshops and Demos (ACIIW)*. IEEE, 2019, pp. 140–146.
- [14] H. Wang, Y. Zou, and D. Chong, “Acoustic scene classification with spectrogram processing strategies,” in *Proceedings of 5th the Workshop on Detection and Classification of Acoustic Scenes and Events (DCASE)*, 2020, pp. 210–214.
- [15] S. Ioffe and C. Szegedy, “Batch normalization: Accelerating deep network training by reducing internal covariate shift,” in *International conference on machine learning*. PMLR, 2015, pp. 448–456.
- [16] V. Nair and G. Hinton, “Rectified linear units improve restricted boltzmann machines,” in *ICML*, 2010.
- [17] Y. Wang, J. Li, and F. Metze, “A comparison of five multiple instance learning pooling functions for sound event detection with weak labeling,” in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2019, pp. 31–35.
- [18] H. Zhang, M. Cisse, Y. N. Dauphin, and D. Lopez-Paz, “mixup: Beyond empirical risk minimization,” *arXiv preprint arXiv:1710.09412*, 2017.
- [19] D. S. Park, W. Chan, Y. Zhang, C.-C. Chiu, B. Zoph, E. D. Cubuk, and Q. V. Le, “SpecAugment: A simple data augmentation method for automatic speech recognition,” *Proc. Interspeech*, pp. 2613–2617, 2019.
- [20] Y. Tokozume, Y. Ushiku, and T. Harada, “Learning from between-class examples for deep sound recognition,” in *International Conference on Learning Representations*, 2018.
- [21] J. Salamon, D. MacConnell, M. Cartwright, P. Li, and J. P. Bello, “Scaper: A library for soundscape synthesis and augmentation,” in *IEEE Workshop on Applications of Signal Processing to Audio and Acoustics (WASPAA)*. IEEE, 2017, pp. 344–348.
- [22] J. Salamon, C. Jacoby, and J. P. Bello, “A dataset and taxonomy for urban sound research,” in *Proceedings of the 22nd ACM International Conference on Multimedia*, 2014, pp. 1041–1044.
- [23] I. Martín-Morató, A. Mesáros, T. Heittola, T. Virtanen, M. Cobos, and F. Ferri, “Sound event envelope estimation in polyphonic mixtures,” in *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2019, pp. 935–939.
- [24] K. Miyazaki, T. Komatsu, T. Hayashi, S. Watanabe, T. Toda, and K. Takeda, “Weakly-supervised sound event detection with self-attention,” in *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2020, pp. 66–70.
- [25] Z. Ye, X. Wang, H. Liu, Y. Qian, R. Tao, L. Yan, and K. Ouchi, “Sound event detection transformer: An event-based end-to-end model for sound event detection,” *arXiv preprint arXiv:2110.02011*, 2021.
- [26] D. Kingma and J. Ba, “Adam: A method for stochastic optimization,” in *ICLR*, 2015.
- [27] A. Mesáros, T. Heittola, and T. Virtanen, “Metrics for polyphonic sound event detection,” *Applied Sciences*, vol. 6, no. 6, p. 162, 2016.