



# RCT: Random Consistency Training for Semi-Supervised Sound Event Detection

Nian Shao<sup>1</sup>, Erfan Loweimi<sup>2</sup>, Xiaofei Li<sup>1\*</sup>

<sup>1</sup> Westlake University & Westlake Institute for Advanced Study, Hangzhou, China

<sup>2</sup> Centre for Speech Technology Research (CSTR), University of Edinburgh, Edinburgh, UK

shaonian@westlake.edu.cn, e.loweimi@ed.ac.uk, lixiaofei@westlake.edu.cn

## Abstract

Sound event detection (SED), as a core module of acoustic environmental analysis, suffers from the problem of data deficiency. The integration of semi-supervised learning (SSL) largely mitigates such problem. This paper researches on several core modules of SSL, and introduces a random consistency training (RCT) strategy. First, a hard mixup data augmentation is proposed to account for the additive property of sounds. Second, a random augmentation scheme is applied to stochastically combine different types of data augmentation methods with high flexibility. Third, a self-consistency loss is proposed to be fused with the teacher-student model, aiming at stabilizing the training. Performance-wise, the proposed modules outperform their respective competitors, and as a whole the proposed SED strategies achieve 44.0% and 67.1% in terms of the PSDS<sub>1</sub> and PSDS<sub>2</sub> metrics proposed by the DCASE challenge, which notably outperforms other widely-used alternatives.

**Index Terms:** semi-supervised learning, sound event detection, data augmentation, consistency regularization, hard mixup

## 1. Introduction

Sound conveys a substantial amount of information about the environment. The skill of recognizing the surrounding environment is taken for granted by humans while it is a challenging task for machines [1]. Sound event detection (SED) aims to detect sound events within an audio stream by labeling the events as well as their corresponding occurrence timestamps. Taking advantage of deep neural networks, promising results have been obtained for SED [2]. However, the high annotation cost poses obstacles on its further development.

Two widely applied solutions for such data deficiency problem are data augmentation (DA) and semi-supervised learning (SSL). SSL leverages abundant unlabelled data to improve the model generalization capacity while DA artificially enlarges the training dataset size in various forms like data warping, over-sampling, etc. [3]. For example, in computer vision (CV), different DA methods were proposed to transform the training images including rotation, noise injection, etc. [3, 4]. To combine multiple augmentation methods, *RandAugment* [5] presents a random strategy which arbitrarily selects one transformation in each training step. On the other hand, *mixup* [6, 7] conducts a linear interpolation of two data points to oversample the dataset, by which the mixed samples could push the decision boundaries into low-density regions. Essentially, DA regularizes the model by constraining the predicted labels to be invariant to any noise applied on the inputs. Such idea is known as *consistency regularization* (CR) in SSL [8]. When using CR for training, the model predictions are constrained to be invariant to any noise

not only on the inputs [9] but also on the hidden states [10, 11]. However, there exists a potential risk known as *confirmation bias* when the consistency loss is too heavily weighted in training [12]. To alleviate such risk, *MeanTeacher* [12] applies a consistency constraint in the model parameter space, which holds an exponential moving average (EMA) of the training student model to generate pseudo labels for unlabelled data. Other techniques such as *interpolation consistency training* (ICT) [9] and *unsupervised data augmentation* (UDA) [13] further combine mixup or *RandAugment* with *MeanTeacher* and CR, obtaining state-of-the-art SSL models.

The efforts on semi-supervised SED have achieved promising results [14–18], thanks to the Detection and Classification of Acoustic Scenes and Events (DCASE) challenge task 4<sup>1</sup>, which establishes a systematically organized semi-supervised dataset [2] containing a reasonable amount of weakly-labelled and unlabelled audio clips. The recent top-rank systems not only apply audio-specified DA methods including *SpecAugment* [19], time shift [20], pitch shift [21], etc., but also draw lessons from the SSL methods such as CR and *MeanTeacher*. The top-rank system in 2018 DCASE challenge [14] introduced *MeanTeacher* into the semi-supervised SED, which became a role system for many variant systems [15–18, 22]. Similar to ICT [9], shift consistency training (SCT) [18] proposes a way for unifying mixup, time shift and CR, obtaining a consistently better semi-supervised SED model. Utilizing SCT and ICT, Zheng et al. [15] achieved the top rank in DCASE 2021 challenge task 4. Although multiple audio-specified DA methods [18, 19, 21] are leveraged in these models, the proper adaptation and effective combination of these SSL techniques are not scrutinised to be optimally applied in audio processing.

In this work, concerned with the proper usages of SSL techniques in audio processing, we propose a SSL strategy for SED, which consists of three novel modules:

i) *Hard mixup*. Hard mixup investigates the proper usage of the mixup method in the audio processing tasks. The vanilla mixup [6] is trivially applied in many audio processing studies [23–25], while its adaptability for audio processing is not fully investigated. In the vanilla mixup, the mixed label is a soft convex combination of the labels from original samples. Instead, hard mixup preserves the hard label of the original samples, as the mixup of sound events, yielding concurrent sound events due to the additive property of audio signal. For example, the mixed label [0.4, 0, 0.6, 0] in the vanilla mixup would be presented as [1, 0, 1, 0] in the hard mixup.

ii) *Random warping*. Direct combination of various DA methods is not guaranteed to result in a performance gain, because of the complexity of finding an optimal set in a large hyperparameter searching space [3]. Although many efforts have

\* Corresponding author.

<sup>1</sup><https://dcase.community/challenge2022>

been made toward unifying various data augmentation schemes in CV [5, 26], they could not be trivially adapted for audio processing, since methods such as time shift [18] or pitch shift [21] are specifically designed for audio. Random warping is among the first attempts toward unifying data augmentation methods for audio processing. We perturb each training sample with a randomly selected transformation, which allows taking advantage of different types of augmentation techniques in a unified way. It is a simple yet effective policy, while finding the optimal magnitude for each transformation can be challenging. We empirically find the optimal magnitude values for three data augmentation methods and achieve a consistent performance gain.

iii) *Self-consistency loss*. One of the challenges in SSL is designing the unsupervised loss for unlabelled data [27]. In ICT [9] for mixup and in SCT [18] for time shift, the MeanTeacher model and student model respectively process the original and augmented samples. In this paper, we propose an additional self-consistency loss to the MeanTeacher loss, which constrains the student model to give consistent predictions for original and augmented samples. Such self-consistency constraint always holds regardless of the correctness of the predictions, and thus would stabilize the training procedure.

The combination of these three modules is referred to as *random consistency training* (RCT), and the flowchart is shown in Fig. 1. With better adaptability for corresponding audio signals, each proposed module experimentally outperform its competing methods in the literature. As a whole, the proposed SSL strategy notably outperforms its counterparts, and achieves top performance on the DCASE 2021 challenge dataset.

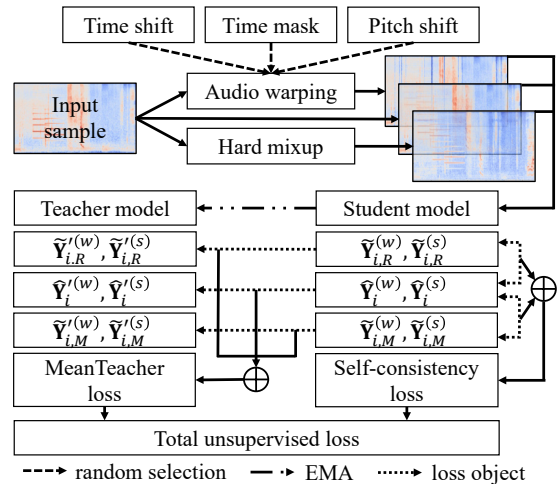
## 2. The Proposed Method

SED is defined as a multi-class detection problem where the onset and offset timestamps of multiple sound events should be recognized from the input audio clips. We denote the time-frequency domain audio clips as  $\mathbf{X}_i^{(l)} \in \mathbb{R}^{T \times K}$ , where  $T$  is the number of frames (same for all clips in experiment) and  $K$  is the dimension of LogMel filterbank features. Three types of data annotations are used for training data, i.e. weakly labelled, strongly labelled and unlabelled, which are indicated by superscript  $l \in \{w, s, u\}$ , respectively.  $i$  denotes the index of data sample among a total of  $N^{(l)}$  data points of one annotation type. Let  $C$  and  $\mathbf{Y}_i^{(l)}$  be the number of sound event classes and data labels, respectively. The weakly labelled and strongly labelled data have the clip-level and frame-level labels denoted by  $\mathbf{Y}_i^{(w)} \in \mathbb{R}^C$  and  $\mathbf{Y}_i^{(s)} \in \mathbb{R}^{T' \times C}$ , respectively. Since the required time resolution for SED is much lower than the one of sound frames, pooling layers are applied in the CNN, resulting in a coarser time resolution  $T'$  rather than  $T$  for the predictions.

The baseline model is a CRNN [14], consisting of a 7-layer CNN with Context Gating layers [28], cascaded by a 2-layer bidirectional GRU. An attention module is added at the end to produce different levels of predictions [14] and MeanTeacher [12] is employed for SSL. As shown in Fig. 1, a teacher model is obtained by an EMA of the student CRNN model to provide pseudo labels for unlabelled samples. The training loss is  $\mathcal{L} = \mathcal{L}_{\text{Supervised}} + \mathcal{L}_{\text{MeanTeacher}}$ , where  $\mathcal{L}_{\text{Supervised}}$  is the cross-entropy loss for the labelled data, and  $\mathcal{L}_{\text{MeanTeacher}}$  is the MeanTeacher mean square error (MSE) loss for the unlabelled data.

### 2.1. Random data augmentation for audio

RandAugment [5] is proposed as an efficient way of combining different types of image transformations. In RCT, we take such



**Figure 1:** Flowchart of RCT: both hard mixup and audio warping are first applied for data augmentation; MeanTeacher and self-consistency are used for SSL training. Subscripts R and M stand for Random warping and hard mixup, respectively.

idea to construct general audio warping methods for consistency regularization. The random data augmentation is accomplished by combining it with the proposed hard mixup, which means each training sample is augmented by both hard mixup and random warping. As shown in Fig. 1, the batch size is thus tripled in each training step.

**Hard mixup:** The vanilla mixup [6] conducts an interpolation of two data points belonging to different classes, aiming to smooth the decision boundary. Such operation is feasible for images, whereas the interpolation of two audio clips produces a new clip due to the additive property of the sounds. As a result, combinations of multiple audio clips could be regarded as a realistic sample containing concurrent sound events, and should be recognized as an audio clip involving all sound events. Thus, the audio mixup is proposed to directly add multiple samples together, and the mixture is labelled with all the classes in all original samples. Since each sound clip reflect the realistic energies of all sound sources, the energy of the mixed sample is remained unchanged. Moreover, we noticed that combining more than two audio clips brings extra benefits. That is, it further condenses the distribution of sound events and can help the model toward better discriminating the sound events. Therefore, we randomly add two or three samples together in the hard mixup.

**Random warping:** We use three audio warping methods in this work, and their warping magnitudes are set to  $d \in \{1, 2, \dots, 9\}$ . The only hyperparameter requires to be tuned is one unique  $d_{\text{max}}$  for all audio warping methods. For each mini batch, one warping method is randomly chosen with a random magnitude  $d$  uniformly distributed in  $[1, d_{\text{max}}]$ . This magnitude controls the audio warping intensities of the following methods:

- i) *Time shift* [18] circularly shifts each audio clip along time axis with a duration of  $1 \times d$  seconds.
- ii) *Time mask* [19] randomly selects  $5d$  intervals from the audio clip to be masked to 0. The length of each mask interval is set to 0.1 seconds.
- iii) *Pitch shift* [21] randomly raises or lowers the pitch of the audio clip by  $1/2 \times d$  semitones, where both pitches and formants are stretched.

The magnitude unit of each method (1 in time shift, 0.1 in time mask, 1/2 semitone in pitch shift) is empirically selected.

## 2.2. Self-consistency regularization

The MeanTeacher loss [12] used in SED [14] has already shown a notable capacity in mining information from unlabelled data. To further leverage the unlabelled data, we propose to apply *self-consistency* regularization in addition to the MeanTeacher loss. Let  $\hat{\mathbf{Y}}_i^{(w)} \in \mathbb{R}^C$  and  $\tilde{\mathbf{Y}}_i^{(w)} \in \mathbb{R}^C$  denote the weak (clip-level) predictions of original and augmented samples of the student model, respectively. Similarly let  $\hat{\mathbf{Y}}_i^{(s)} \in \mathbb{R}^{T' \times C}$  and  $\tilde{\mathbf{Y}}_i^{(s)} \in \mathbb{R}^{T' \times C}$  for the strong (frame-level) predictions. Self-consistency regulates the model by an extra MSE term added to the loss function

$$\begin{aligned} \mathcal{L}_{\text{SC}} = & r(\text{step}) \frac{1}{N^{(w)}C} \sum_i^{N^{(w)}} \|\mathcal{D}_{\text{aug}}^{(w)}(\hat{\mathbf{Y}}_i^{(w)}) - \tilde{\mathbf{Y}}_i^{(w)}\|_2^2 \\ & + r(\text{step}) \frac{1}{N^{(s)}CT'} \sum_i^{N^{(s)}} \|\mathcal{D}_{\text{aug}}^{(s)}(\hat{\mathbf{Y}}_i^{(s)}) - \tilde{\mathbf{Y}}_i^{(s)}\|_2^2, \end{aligned} \quad (1)$$

where  $\|\cdot\|_2$  denotes the Euclidean norm for a vector/matrix,  $r(\text{step})$  is a ramp-up function varying along the training step.  $\mathcal{D}_{\text{aug}}^{(w)}$  and  $\mathcal{D}_{\text{aug}}^{(s)}$  denote transformations applied on the predictions of original samples, as their labels should be correspondingly transformed according to the augmentation methods applied on the augmented samples. For the pitch shift and time mask, there is no need to change the labels. Time shift should accordingly shift the strong predictions along time axis. As for hard mixup, the mixed audio clip includes all the sound classes presented in the original audio clips. However, the labels for the combined sound classes cannot be trivially obtained by adding the predictions of original samples, since the summation of two soft-predictions/labels is meaningless. Instead, we define the label transformation for hard mixup as


$$\mathcal{D}_{\text{mixup}}^{(l)}(\hat{\mathbf{Y}}_i^{(l)}) = \vee_{i \in \mathcal{M}} \text{harden}(\hat{\mathbf{Y}}_i^{(l)}), \quad (2)$$

where  $\vee$  denotes element-wise OR operation,  $\mathcal{M}$  is an arbitrary set consists of two or three data samples used in the hard mixup, and  $\text{harden}(\cdot)$  is an element-wise binary hardening function which ceils (or floors) the matrix elements to 1 (or 0) if the elements are larger than 0.95 (or smaller than 0.05). This transformation first hardens the predictions of original samples, from which the active/inactive sound classes are combined.

The prediction of original and augmented samples ( $\hat{\mathbf{Y}}_i$  and  $\tilde{\mathbf{Y}}_i$ ) are both used for gradient updating, which means they are considered as the pseudo labels for each other during training. This is also one reason for choosing symmetric MSE function as an extra component in the loss function. The total loss  $\mathcal{L}$  used for training the CRNN model will be

$$\mathcal{L} = \mathcal{L}_{\text{Supervised}} + \mathcal{L}_{\text{MeanTeacher}} + \mathcal{L}_{\text{SC}}, \quad (3)$$

where  $\mathcal{L}_{\text{MeanTeacher}}$  is the average of MeanTeacher MSE losses for the original samples and two sorts of augmented samples (hard mixup and random warping), as shown in Fig. 1. The unsupervised loss is composed of the MeanTeacher loss and self-consistency regularization. Such self-consistency constraint between the original and augmented samples always holds regardless of the correctness of the predictions. This is different from ICT [9] that merges the MeanTeacher loss and the consistency loss as one single loss. In ICT, the MeanTeacher loss is set as the MSE loss between the prediction of augmented and original samples, where the predictions are obtained by the teacher model and student model, respectively. Such pseudo labels



**Figure 2:** The relative performance gain as a function of maximum transformation magnitude on the validation set ( $d_{\max}$ ). The transformation magnitude  $d \sim U[1, d_{\max}]$ . The relative performance gain is computed using the baseline performance (PSDS<sub>1</sub> = 34.7%, PSDS<sub>2</sub> = 53.7%). The markers and vertical lines represent the mean and standard deviation computed using three trials.

highly rely on the correctness of the MeanTeacher predictions, while incorrect pseudo labels may mislead the student model and consequently reduce the training efficiency.

## 3. Experimental Results and Discussion

Our codes for this work have been released on our website<sup>2</sup>. We use the baseline model [14] on DCASE 2021 Task 4 dataset<sup>3</sup> to test the performance of the proposed method. The dataset consists of 1578 weakly-labelled, 10000 synthesized strongly-labelled and 14412 unlabelled audio clips. Each 10-second audio clip is resampled to 16 kHz and frame blocked with a length of 128 ms (2048 samples) and a hop length of 16 ms (256 samples). After applying a 2048-point fast Fourier transform, 128-dimensional LogMel features are extracted for each frame, converting the 10-second audio clip into a  $626 \times 128$  spectrogram. All samples are normalized to  $[-1, 1]$  before training.

The batch size is set to 48, consisting of 12 weakly-labelled, 12 strongly-labelled and 24 unlabelled data points. The learning rate ramps up to  $10^{-3}$  until epoch 50 and is scheduled by Adam optimizer [29] until the end of training, i.e. 200 epochs. The weight for MeanTeacher and self-consistency losses,  $r(\text{step})$ , linearly ramps up from 0 to 2 at epoch 50 and is then kept unchanged. The system performance is evaluated through polyphonic sound detection scores (PSDSs) [30] according to the DCASE 2021 challenge guidelines. The metrics takes both response speed (PSDS<sub>1</sub>) and cross-trigger performance (PSDS<sub>2</sub>) into account; the larger the better for both metrics.


### 3.1. Ablation study on SED

In the random warping policy, the only hyperparameter that needs to be grid-searched is the maximum transformation magnitude  $d_{\max}$ . As shown in Fig. 2, the performance improves as the maximum transformation magnitude increases until about 5 or 6; accordingly, we set  $d_{\max}$  to 5.

Table 1 shows the result of ablation study, in which the proposed modules are individually added step by step. As seen, the proposed schemes including hard mixup, random warping and self-consistency, lead to noticeable positive contributions. To further investigate the efficacy of the proposed methods, we also conduct experiments to replace the hard mixup and

<sup>2</sup><https://github.com/Audio-WestlakeURCT>

<sup>3</sup><https://dcase.community/challenge2021/task-sound-event-detection-and-separation-in-domestic-environments>



**Figure 3:** Cross-entropy loss of strongly-supervised validation data when training with or without self-consistency loss, compared with the ICT scheme [9].

**Table 1:** Ablation study for RCT. Different modules are added step by step and each score is obtained by averaging three trials.

Model	PSDS <sub>1</sub> (%)	PSDS <sub>2</sub> (%)
Baseline	34.7	53.7
+ Vanilla mixup [6]	34.9	57.9
+ Hard mixup	36.4	57.4
+ Random warping	38.1	58.5
+ ICT consistency [9]	38.0	59.2
+ Self-consistency	40.1	61.4

self-consistency by the vanilla mixup [6] and ICT-like consistency [9], respectively. While the vanilla mixup is slightly better in cross-trigger (PSDS<sub>2</sub>), the hard mixup gives more significant gain in terms of the response time (PSDS<sub>1</sub>). The proposed self-consistency outperforms the ICT-like consistency in both metrics.

### 3.2. Comparison with other semi-supervised strategies

To evaluate the proposed SSL strategy collectively, we conduct a thorough comparison with other widely used SSL strategies, including ICT, SCT and their combination, employing the same baseline network used for the proposed strategy. Table 2 shows the comparison results, and Fig. 3 illustrates the validation loss vs epoch for the models with baseline MeanTeacher, ICT and self-consistency. Compared with the baseline MeanTeacher model, although ICT largely improves the performance by its proposed teacher-student consistency loss, the always-hold self-consistency loss further improves the performance with a lower validation loss curve shown in Fig. 3. SCT performance is not as good as ICT, indicating that the time shift is not as efficient as interpolation (mixup). Combining ICT and SCT does not outperform ICT alone, which implies that the naive addition is not an effective way to combine multiple different augmentations schemes. In contrast, as shown in Table 1, the proposed strategy is able to efficiently combine multiple different augmentations. Overall, the proposed method remarkably outperforms both ICT and SCT, due to the strength of the proposed modules and efficient stochastic combination of them.

### 3.3. Comparison with DCASE2021 submissions

To further assess the efficacy of the proposed method and conduct fair comparisons with heavily tuned DCASE 2021 submitted models, we also took advantage of some existing post-processing and ensembling techniques in our model. A temperature factor of 2.1 is used for inference temperature tuning as in [15]. Median filter is applied to frame-level predictions to

**Table 2:** Comparing the proposed SSL strategy with other alternatives. Each score is obtained by averaging three trials.

Model	PSDS <sub>1</sub> (%)	PSDS <sub>2</sub> (%)
Baseline [14]	34.7	53.7
SCT [18]	36.0	55.6
ICT [9]	37.7	57.7
ICT+SCT [18]	37.0	58.7
<b>RCT (proposed)</b>	<b>40.1</b>	<b>61.4</b>

**Table 3:** Comparing the proposed system with DCASE2021 top-ranked submissions. All models are named in the form of network architecture plus the SSL strategy.

Model	PSDS <sub>1</sub> (%)	PSDS <sub>2</sub> (%)
CRNN (baseline) [14]	34.7	53.7
FBCRNN+MLFL [24]	40.1	59.7
CRNN+IPL [17]	40.7	65.3
CRNN+DA [25]	41.9	63.8
CRNN+HeavyAug. [16]	43.4	63.9
RCRNN+NS [23]	45.1	67.9
SKUnit+ICT/SCT [15]	45.4	67.1
CRNN+ <b>RCT (proposed)</b>	<b>44.0</b>	<b>67.1</b>

smooth the possibly fluctuating original predictions, following the method proposed in [31]. Moreover, model ensembling was applied to fuse the predictions of multiple differently trained models. We trained eleven models with different variants of RCT: substituting time masking with frequency masking [19]; adding FilterAug [16] into audio warping choices; randomly selecting one or two methods in audio warping; and, reducing the weight of the MeanTeacher loss. We found that all different variants achieve reasonable performances. This demonstrates the flexibility of RCT and its capacity in incorporating new audio transformations into the framework with a low tuning cost overhead.

The proposed system is compared with DCASE2021 top-ranked submissions in Table 3. The scores of DCASE2021 submissions are directly quoted from the challenge results. The proposed system noticeably outperforms all other systems employing the baseline CRNN network, which verifies the superiority of the proposed RCT strategy. Furthermore, as seen the performance of the proposed system is very close to the two first-ranked submissions [15, 23]. They both use more powerful networks, i.e., SKUnit and RCRNN, which were able to largely improve the performance [15, 23]. The proposed framework is independent of the network and can be easily applied along with more advanced architectures to achieve higher performance.

## 4. Conclusion

In this paper, we developed a novel semi-supervised learning (SSL) strategy, named random consistency training (RCT), for sound event detection (SED) task. The proposed method improves several core modules of SSL, including unsupervised training loss and data augmentation schemes. It leads to achieving competitive performance on the DCASE 2021 challenge dataset. RCT provides a generic framework which can be effectively employed along with more advanced augmentation schemes and architectures. Besides, since RCT is not task-specific, it can potentially be applied in various audio and image processing tasks which is another broad avenue for future work.

## 5. References

- [1] T. Virtanen, M. D. Plumbley, and D. Ellis, *Computational Analysis of Sound Scenes and Events*. Springer, 2018.
- [2] N. Turpault, R. Serizel, A. Shah, and J. Salamon, "Sound event detection in domestic environments with weakly labeled data and soundscape synthesis," in *Acoustic Scenes and Events 2019 Workshop (DCASE2019)*, 2019, p. 253.
- [3] C. Shorten and T. M. Khoshgoftaar, "A survey on image data augmentation for deep learning," *Journal of Big Data*, vol. 6, no. 1, pp. 1–48, 2019.
- [4] F. J. Moreno-Barea, F. Strazzera, J. M. Jerez, D. Urda, and L. Franco, "Forward noise adjustment scheme for data augmentation," in *2018 IEEE symposium series on computational intelligence (SSCI)*. IEEE, 2018, pp. 728–734.
- [5] E. D. Cubuk, B. Zoph, J. Shlens, and Q. V. Le, "Randaugment: Practical automated data augmentation with a reduced search space," in *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*. IEEE Computer Society, 2020, pp. 3008–3017.
- [6] H. Zhang, M. Cisse, Y. N. Dauphin, and D. Lopez-Paz, "mixup: Beyond empirical risk minimization," in *International Conference on Learning Representations (ICLR)*, 2018.
- [7] Y. Tokozume, Y. Ushiku, and T. Harada, "Learning from between-class examples for deep sound recognition," in *International Conference on Learning Representations (ICLR)*, 2018.
- [8] M. Sajjadi, M. Javanmardi, and T. Tasdizen, "Regularization with stochastic transformations and perturbations for deep semi-supervised learning," *Advances in neural information processing systems (NIPS)*, vol. 29, pp. 1163–1171, 2016.
- [9] V. Verma, A. Lamb, J. Kannala, Y. Bengio, and D. Lopez-Paz, "Interpolation consistency training for semi-supervised learning," in *International Joint Conference on Artificial Intelligence (IJCAI)*, 2019, pp. 3635–3641.
- [10] P. Bachman, O. Alsharif, and D. Precup, "Learning with pseudo-ensembles," *Advances in Neural Information Processing Systems (NIPS)*, vol. 27, pp. 3365–3373, 2014.
- [11] L. Samuli and A. Timo, "Temporal ensembling for semi-supervised learning," in *International Conference on Learning Representations (ICLR)*, 2017.
- [12] A. Tarvainen and H. Valpola, "Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results," *Advances in Neural Information Processing Systems (NIPS)*, vol. 30, p. 1195–1204, 2017.
- [13] Q. Xie, Z. Dai, E. Hovy, T. Luong, and Q. Le, "Unsupervised data augmentation for consistency training," *Advances in Neural Information Processing Systems (NIPS)*, vol. 33, pp. 6256–6268, 2020.
- [14] L. JiaKai, "Mean teacher convolution system for dcase 2018 task 4," DCASE2018 Challenge, Tech. Rep., June 2018.
- [15] X. Zheng, H. Chen, and Y. Song, "Zheng ustc team's submission for dcase2021 task4 – semi-supervised sound event detection," DCASE2021 Challenge, Tech. Rep., June 2021.
- [16] H. Nam, B.-Y. Ko, G.-T. Lee, S.-H. Kim, W.-H. Jung, S.-M. Choi, and Y.-H. Park, "Heavily augmented sound event detection utilizing weak predictions," DCASE2021 Challenge, Tech. Rep., June 2021.
- [17] Y. Gong, C. Li, X. Wang, L. Ma, S. Yang, and Z. W. Wu, "Improved pseudo-labeling method for semi-supervised sound event detection," DCASE2021 Challenge, Tech. Rep., June 2021.
- [18] C.-Y. Koh, Y.-S. Chen, Y.-W. Liu, and M. R. Bai, "Sound event detection by consistency training and pseudo-labeling with feature-pyramid convolutional recurrent neural networks," in *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2021, pp. 376–380.
- [19] D. S. Park, W. Chan, Y. Zhang, C. Chiu, B. Zoph, E. D. Cubuk, and Q. V. Le, "SpecAugment: A simple data augmentation method for automatic speech recognition," in *Interspeech*, 2019.
- [20] L. Delphin-Poulat and C. Plapous, "Mean teacher with data augmentation for dcase 2019 task 4," DCASE2019 Challenge, Tech. Rep., June 2019.
- [21] B. McFee, E. Humphrey, and J. Bello, "A software framework for musical data augmentation," in *16th International Society for Music Information Retrieval Conference (ISMIR)*, 2015, pp. 248–254.
- [22] J. Yan, Y. Song, L.-R. Dai, and I. McLoughlin, "Task-aware mean teacher method for large scale weakly labeled semi-supervised sound event detection," in *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2020, pp. 326–330.
- [23] N. K. Kim and H. K. Kim, "Self-training with noisy student model and semi-supervised loss function for dcase 2021 challenge task 4," DCASE2021 Challenge, Tech. Rep., June 2021.
- [24] G. Tian, Y. Huang, Z. Ye, S. Ma, X. Wang, H. Liu, Y. Qian, R. Tao, L. Yan, K. Ouchi, J. Ebberts, and R. Haeb-Umbach, "Sound event detection using metric learning and focal loss for dcase 2021 task 4," DCASE2021 Challenge, Tech. Rep., June 2021.
- [25] R. Lu, W. Hu, D. Zhiyao, and J. Liu, "Integrating advantages of recurrent and transformer structures for sound event detection in multiple scenarios," DCASE2021 Challenge, Tech. Rep., June 2021.
- [26] E. D. Cubuk, B. Zoph, D. Mané, V. Vasudevan, and Q. V. Le, "Autoaugment: Learning augmentation strategies from data," in *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, 2019, pp. 113–123.
- [27] D.-H. Lee *et al.*, "Pseudo-label: The simple and efficient semi-supervised learning method for deep neural networks," in *Workshop on challenges in representation learning, ICML*, vol. 3, no. 2, 2013, p. 896.
- [28] M. Antoine, L. Ivan, and S. Josef, "Learnable pooling with context gating for video classification," *CoRR*, vol. abs/1706.06905, 2017. [Online]. Available: <http://arxiv.org/abs/1706.06905>
- [29] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *International Conference on Learning Representations (ICLR)*, 2015.
- [30] A. Mesaros, T. Heittola, and T. Virtanen, "Metrics for polyphonic sound event detection," *Applied Sciences*, vol. 6, no. 6, 2016. [Online]. Available: <https://www.mdpi.com/2076-3417/6/6/162>
- [31] Y. Liu, C. Chen, J. Kuang, and P. Zhang, "Semi-supervised sound event detection based on mean teacher with power pooling and data augmentation," DCASE2020 Challenge, Tech. Rep., June 2020.