

# FREQUENCY & CHANNEL ATTENTION FOR COMPUTATIONALLY EFFICIENT SOUND EVENT DETECTION

*Hyeonuk Nam, Seong-Hu Kim, Deokki Min, Yong-Hwa Park*

Korea Advanced Institute of Science and Technology, South Korea  
{frednam, seonghu.kim, minducky, yhpark}@kaist.ac.kr

## ABSTRACT

We explore on various attention methods on frequency and channel dimensions for sound event detection (SED) in order to enhance performance with minimal increase in computational cost while leveraging domain knowledge to address the frequency dimension of audio data. We have introduced frequency dynamic convolution (FDY conv) in a previous work to release the translational equivariance issue associated with 2D convolution on the frequency dimension of 2D audio data. Although this approach demonstrated state-of-the-art SED performance, it resulted in a model with 150% more trainable parameters. To achieve comparable SED performance with computationally efficient methods for practicality, we explore on lighter alternative attention methods. In addition, we focus on attention methods applied to frequency and channel dimensions. Joint application Squeeze-and-excitation (SE) module and time-frame frequency-wise SE (tfwSE) to apply attention on both frequency and channel dimensions shows comparable performance to SED model with FDY conv with only 2.7% more trainable parameters compared to the baseline model. In addition, we performed class-wise comparison of various attention methods to further discuss various attention methods' characteristics.

**Index Terms**— sound event detection, computationally efficient, attention, frequency dimension, channel dimension

## 1. INTRODUCTION

Sound event detection (SED), which aims to recognize a target sound event class and corresponding time localization within a given audio clip, has potential to be applied in various applications such as automation, robotics and monitoring [1, 2, 3]. In order to recognize and locate sound events, we need strong pattern recognition tools. Recent advances in deep learning (DL) methods brought significant progress in SED [2, 3]. While most works directly applied DL methods from other domains to SED without modification, few works adapted DL methods to SED by thoroughly analysing unique characteristics of audio data and sound events.

Frequency dimension has to be carefully considered when applying DL methods on audio-related DL applications. It is shown by previous works that methods considering frequency dimension significantly improved SED performance [4, 5, 6]. SED has been heavily relying on convolutional recurrent neural networks (CRNN)

based architectures [2, 3]. 2D convolution in CRNN assumes shift-invariance on both time and frequency dimensions thus enforces translational equivariance on both dimensions [4]. However, frequency is a shift-variant dimension where the same pattern sounds different when translated along the frequency dimension. At the same time, frequency exhibits loose shift-invariance within short frequency range thus slight pitch-shift does not harm auditory perception much. Thus frequency dimension is a delicate yet essential component to be considered for audio domain.

In a previous study, we introduced frequency dynamic convolution (FDY conv) to release translational equivariance by 2D convolution on the frequency dimension of 2D audio data to consider its shift-variant characteristic [4]. While FDY conv showed impressive performance on SED, it added 150% more parameters to the model. However, in order to apply SED on various real applications, we might need to implement SED on devices with limited specifications. Thus, there is a need for computationally efficient SED methods which is lighter but as competent as current state-of-the-art models. To address this limitation and improve the practicality of SED models, we explore various lighter attention methods to enhance SED performance more efficiently. We aim to achieve this by addressing the frequency and channel dimensions, since those are two emphasized dimensions in audio domain [4, 6, 7]. Thus we experiment with various attention methods on frequency and channel dimensions. The main contributions of this work are:

1. We explore various alternative attention methods which are computationally efficient for practicality, while considering channel and frequency dimensions to consider unique characteristics of audio domain.
2. Joint application of squeeze-and-excitation (SE) and proposed time-frame frequency-wise SE (tfwSE) to re-weight both channel and frequency dimensions shows comparable performance to state-of-the-art method while only adding model parameters by 2.7%.
3. We discuss the characteristics of various attention methods on SED to provide further insights for practical implementation.

The official implementation code is available on GitHub<sup>1</sup>.

## 2. METHODS

While frequency dynamic convolution (FDY conv) showed state-of-the-art performance and have been widely adopted on SED [6, 8, 9, 10, 11, 12], it adds considerable number of trainable parameters to the networks due to multiple basis kernels [4, 13].

This work was supported by the Institute of Civil Military Technology Cooperation funded by the Defense Acquisition Program Administration and Ministry of Trade, Industry and Energy of Korean government under grant No. UM22409RD4, and Korea Research Institute of Ships and Ocean engineering a grant from Endowment Project of "Development of Open Platform Technologies for Smart Maritime Safety and Industries" funded by Ministry of Oceans and Fisheries(PES4880).

<sup>1</sup><https://github.com/frednam93/lightSED>

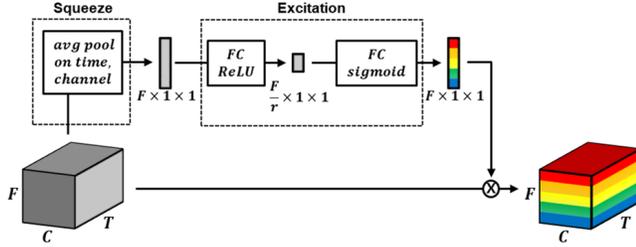


Figure 1: An illustration of frequency-wise Squeeze-Excitation.

Thus there remains a need for sufficiently well-performing model with fewer parameters for practical applications. Since FDY conv’s strength comes from attention mechanism which selectively focus on important elements of the input, we explore other alternative attention methods to achieve comparable performance.

### 2.1. Variants of Squeeze-and-Excitation

One alternative computationally efficient attention method widely used is squeeze-and-excitation (SE) [14]. It has been widely applied to various CNN-based models for its light yet powerful performance. SE module is composed of squeeze operation and excitation operation. Squeeze operation averages output of 2D convolution on two dimensions except channel to obtain squeezed intermediate representation. Excitation operation applies two successive fully connected (FC) layers to obtain attention weights representing relative importance of each channel. The channels of convolution output is re-weighted by multiplying the attention weight [14]. When applied to 2D audio data, squeeze operation is applied to the convolution output by:

$$z_c = \frac{1}{F \times T} \sum_{f=1}^F \sum_{t=1}^T x_{cft} \quad (1)$$

where  $z_c$  is intermediate representation after squeeze operation on  $c$ -th channel and  $x_{cft}$  is the output by preceding 2D convolution with channel index  $c$ , frequency index  $f$  and time index  $t$ .  $F$  and  $T$  are frequency and time dimension sizes of 2D convolution output. The excitation operation is composed of two FC layers as follows:

$$\mathbf{s} = \sigma(\mathbf{W}_2 \delta(\mathbf{W}_1 \mathbf{z})) \quad (2)$$

where  $\mathbf{s}$  is attention weight, also known as scale, which is multiplied to the output of preceding convolution.  $\mathbf{z}$  is the intermediate representation vector. Both span channel dimension of size  $C$ .  $\mathbf{W}_1$  and  $\mathbf{W}_2$  are FC layers,  $\delta$  refers to ReLU activation and  $\sigma$  refers to sigmoid function.

To apply attention-based re-weighting on frequency dimension, Thienpondt *et al.* proposed frequency-wise Squeeze-Excitation (fwSE) which applies SE on frequency dimension instead [15]. Thus, instead of pooling time and frequency dimensions, fwSE pools channel and time dimensions during squeeze operation as follows:

$$z_f = \frac{1}{C \times T} \sum_{c=1}^C \sum_{t=1}^T x_{cft} \quad (3)$$

The following excitation operation is the same as (2), just that two FC layers are applied on frequency dimensions instead. Then, obtained attention weight for each frequency bin is multiplied to corresponding frequency components of the preceding convolution output. Fig. 1 illustrates the fwSE mechanism.

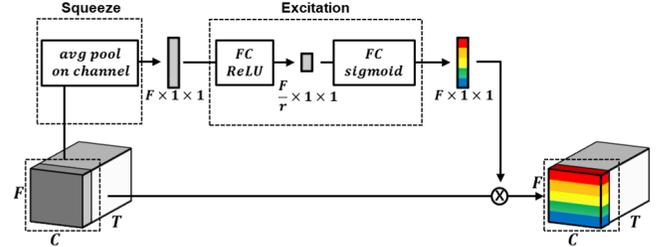


Figure 2: An illustration of time-frame frequency-wise Squeeze-Excitation on one time frame. tfwSE applies this procedure for every time frames.

Since frequency component varies over time, we propose time-frame fwSE (tfwSE) which applies fwSE on every time frames of input instead of time-averaged input. Thus, tfwSE only pools channel dimension in squeeze operation and then applies excitation operation on every time frames. The squeeze operation on time frame  $t$  can be expressed by following equation:

$$z_{ft} = \frac{1}{C} \sum_{c=1}^C x_{cft} \quad (4)$$

where  $z_{ft}$  is intermediate representation after squeeze operation. Then excitation is applied on frequency dimension on each time frame as follows:

$$\mathbf{s}_t = \sigma(\mathbf{W}_2 \delta(\mathbf{W}_1 \mathbf{z}_t)) \quad (5)$$

where  $\mathbf{s}_t$  is scale on time frame  $t$  and  $\mathbf{z}_t$  is the intermediate representation vector corresponding to time frame  $t$ , both spanning channel dimension. Opposed to fwSE by Thienpondt *et al.* which applies frequency-wise attention weights evenly over time-dimension by referring to representative averaged information of the convolution output, proposed tfwSE applies frequency-wise attention weights for each time frame by referring to each individual contents within corresponding time frame [15]. While this could increase computation of excitation operator (fwSE applies excitation on one  $\mathbf{z}$  per audio clip, while tfwSE applies excitation on  $T$   $\mathbf{z}_t$  per audio clip), it could help generalizing excitation operation on many time frames. The mechanism of tfwSE is illustrated in Fig. 2. This method was previously applied by our submission on detection and classification of acoustic scenes and events (DCASE) 2022 challenge task 3 as well, showing its performance [16]. Similarly, we could apply original SE on each time frame as well. We named it as time-frame SE (tSE). Note that this is not a time-wise version of SE like fwSE, as we do not apply SE by pooling channel and frequency dimensions during squeeze to leave time dimension. Instead, we pool frequency dimension only and apply SE on every time frame in similar way shown in Fig. 2.

### 2.2. Channel-Frequency Attention Methods

Li *et al.* [7] proposed C2D-Att for speaker verification which applies 2D convolution to obtain attention weights for both channel and frequency dimensions simultaneously. C2D-Att first pools time dimension by averaging, and then apply two consecutive 2D convolution modules to channel and frequency dimensions by introducing additional channel dimension which is increased to 8 and then back to 1. This results in channel-frequency attention weights

which are multiplied to channel and frequency dimensions of preceding 2D convolution output. C2D-Att improves the speaker verification performance compared to fwSE by re-weighting channel and frequency dimensions simultaneously [7].

However, considering that channel dimension in CNN is permutable dimension where the convolution module’s advantage capturing locality does not matter, it needs further verification if 2D convolution is the best option to apply channel and frequency attention on CNN. While CNN in C2D-Att applies 2D convolution kernel which finds local pattern across frequency and channel dimension, locality matters on frequency dimension only. Therefore, we experiment on joint application of attention on frequency and channel separately, without considering the locality of channel dimension using SE. To apply SE on two dimensions independently, we apply SE and tfwSE in series.

### 3. EXPERIMENTAL SETUPS

#### 3.1. Model Architecture

The model architecture is based on CRNN model, composed of seven convolution layers followed by two bidirectional gated recurrent unit (GRU) then a FC layer. On the strong predictions, we apply class-wise median filter as post processing. In this work, the model using FDY conv replaced all 2D convolution except the first one. SE and C2D-Att modules are inserted after the activation and before the average pooling within the convolution blocks. They are applied on all convolution layers except the last layer in this work. It is because Hu *et al.* has shown that SE module applies almost constant attention weights at the last layer thus it merely affects the model [14].

#### 3.2. Implementation Details

The overall implementation details follow the previous work [4], which could be referred on the official implementation code of which link is provided in the section 1. The experiments in this work are based on domestic environment sound event detection (DESED) dataset [3]. DESED is composed of synthesized strongly labeled dataset, real weakly labeled dataset and real unlabeled dataset for training and validation. For test, real validation dataset, which is strongly labeled, is used. We do not use any external dataset. We trained each model with single NVIDIA RTX Titan GPU. For the results listed in this paper, the metrics are based on the best score among total 24 models from 12 separate training runs.

DESED is composed of 10 second audio data with 16 kHz sampling rate. We extract mel spectrogram as the input feature for SED model. The settings for mel spectrograms are as follows: 2048 points for number of fft, 256 points for hop length, Hamming window for windowing function, and 128 mel bins. Data augmentation methods applied are frame shift [3], mixup [17], time masking [18] and FilterAugment [5]. Applying heavy data augmentation is crucial for training SED where real strongly labeled data is scarce [19]. As we use three levels of datasets, strongly labeled/weakly labeled/unlabeled dataset, we apply mean teacher to leverage unlabeled dataset [3, 20]. We apply FilterAugment with different random parameters on student and teacher model to train SED model robust against FilterAugment.

Table 1: Performance and computational cost comparison between the baseline, frequency dynamic convolution and various frequency and channel attention methods on DESED real validation dataset.

models	params	time	PSDS1	PSDS2	CB-F1
baseline	4.428M	3h 34m	0.409	0.641	0.520
+FDYconv	11.061M	6h 08m	<b>0.446</b>	<b>0.673</b>	0.525
+SE	4.537M	3h 49m	<b>0.435</b>	<b>0.654</b>	0.525
+tSE	4.537M	3h 52m	0.416	0.643	0.526
+fwSE	4.439M	3h 49m	0.411	0.634	0.522
+tfwSE	4.439M	3h 50m	0.415	0.638	0.509
+C2D-Att	4.429M	3h 53m	0.434	<b>0.659</b>	0.539
+tfwSE +SE	4.548M	4h 04m	0.437	0.650	0.532
+SE +tfwSE	4.548M	4h 06m	<b>0.442</b>	<b>0.657</b>	0.526

#### 3.3. Evaluation Metrics

Main evaluation metric employed in this study is the polyphonic sound detection score (PSDS) [21], which considers the intersection between predictions and ground truth to decide if prediction is correct. PSDS also accounts for cross triggers induced by other sound events in the audio. PSDS utilizes area under curve (AUC) - receiver operating characteristic (ROC) curves, enabling comparison of sound event detection (SED) performances without the need for threshold optimization. In DCASE Challenge 2021, 2022 and 2023 Task 4, two variations of PSDS (PSDS1 and PSDS2) are utilized to evaluate SED systems [3]. PSDS1 places emphasis on precise time localization by limiting tolerance for intersection criteria, while PSDS2 prioritizes accurate classification by penalizing cross triggers more. Additionally, we use collar-base F1 score (CB-F1) [22] for class-wise performance comparison, as PSDS cannot be obtained for single sound event. Both PSDS and CB-F1 are ranged between zero and one, and value closer to one indicates better SED performance.

## 4. RESULTS AND DISCUSSION

#### 4.1. Comparison of Attention Modules

Table 1 shows performance and computational cost of SED models with various frequency-wise and channel-wise attention methods. Computational costs are described by the number of trainable parameters representing model size and training time representing computational efficiency. Note that we aim to achieve computational efficiency as close to the baseline as possible and much less than FDY conv. For comparison, SED model with FDY conv is listed in table 1 as well. Note that the results for FDY-CRNN differ from the results in previous paper due to minor changes in setting. When we compare the performance of SED model with SE variants, we can observe that conventional SE definitely outperforms the baseline. On the other hand, fwSE only slightly outperform the baseline for PSDS1 while their PSDS2 is worse than the baseline. Considering that SE is proposed to re-weight channel dimension and each channel is independent from each other while frequency depends on other frequency bins, re-weighting appears to be more effective on channel dimension than on frequency dimension. In addition, considering the parameter increase in the model, SE has increased model size significantly more thus it involved more computational resource to the model. While SE has increased model size by  $\sim 2.5\%$ , fwSE has increased the model size by  $\sim 0.25\%$ . Proposed tfwSE is slightly better than fwSE in terms

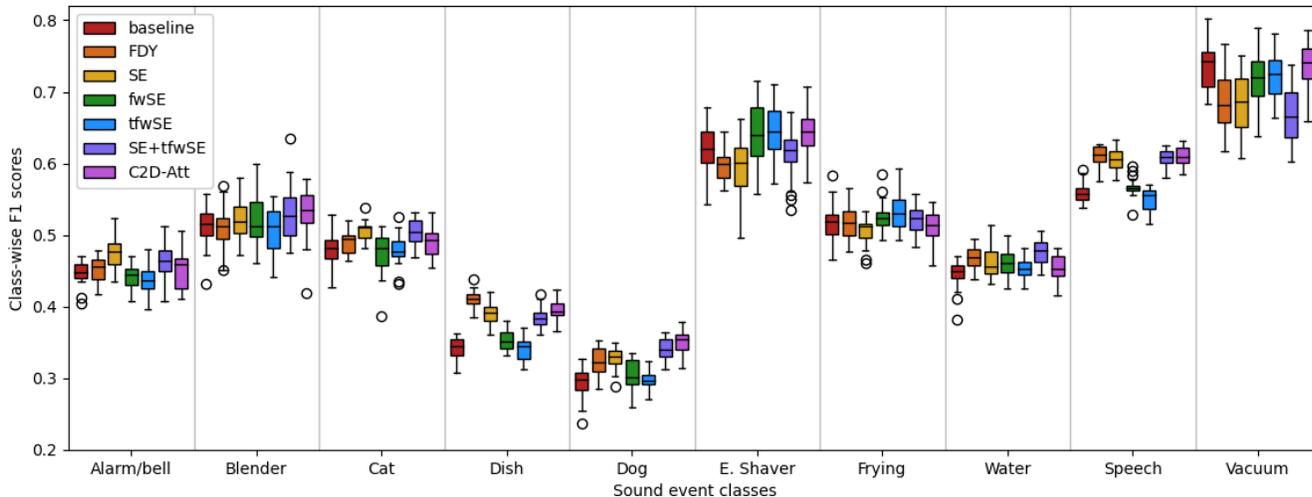


Figure 3: Box-plot of class-wise collar-base F1 scores by multiple models on DESED real validation dataset.

of both PSDS1 and PSDS2. On the other hand, tSE only slightly outperforms the baseline and performs worse than SE. While re-weighting frequency dimension on each time frame has improved frame-wise sound event classification of SED, this effect seems to be not so significant. One explanation to this could be the effect of bi-GRU which processes time-varying information. Likewise, temporal dynamic convolution which applies time-adaptive kernel performed worse than FDY conv on SED [23, 4]. On the other hand, tSE failed to improve SE. Re-weighting a dimension separately on each time frame was not as effective on channel dimension.

Results for methods applying attention simultaneously on channel and frequency dimensions, C2D-Att and joint applications of SE and tfwSE, are also listed in Table 1. C2D-Att shows descent performance comparable to SE, with less parameters compared to SE. In addition, joint application of SE followed by tfwSE shows improvement over SE. While joint application of SE after tfwSE shows similar performance to SE, we could still conclude that application of attention methods simultaneously on channel and frequency dimensions are effective. Furthermore, the combination of SE and tfwSE achieves comparable results to FDY conv in terms of PSDS1, reaching 99.1% of the PSDS1 by the model with FDY conv. Considering that high PSDS2 scores can be easily achieved using weakSED [19], we could regard that this model performs nearly as well as model with FDY conv. An interesting discovery is that while tfwSE degrades PSDS2 for the baseline model, the joint application of tfwSE after SE enhances PSDS2 compared to the model with SE alone. Moreover, considering that SE + tfwSE outperforms C2D-Att for PSDS1, 2D convolution considering locality of 2-dimensional patterns along frequency-channel dimensions is not as effective as separate consideration of channel and frequency dimensions. However, C2D-Att has advantage over SE + tfwSE in terms of the number of parameters which is increased by very small amount.

#### 4.2. Class-wise Performance Comparison

In Fig. 3, class-wise collar-based F1 scores on multiple models are shown as box-plot. Each box-plot is composed of class-wise F1 scores by 24 models from 12 separate training runs. Consistent to table 1, SE performs better than fwSE and tfwSE on many classes

in Fig. 3 as well. SE performed better than fwSE and tfwSE did on alarm/bell ringing, cat, dish, dog and speech while it performed worse on electric shaver, frying and vacuum cleaner. It seems that SE is stronger on transient and non-stationary sound events while it is weaker on quasi-stationary sound events, similar to FDY conv [4]. That is to say, while fwSE and tfwSE re-weight frequency dimension to address frequency dimension, they are stronger on quasi-stationary sound events than on non-stationary sound events. SE + tfwSE shows similar tendency with SE, but slightly better performance in general. Thus SE + tfwSE perform relatively better on non-stationary sound events and relatively worse on quasi-stationary sound events as well. C2D-Att also shows similar tendency with SE but it shows better performance on electric shaver and vacuum cleaner. Note that PSDS is an intersection-based score while the box-plots are based on collar-based score, there are slight discrepancy between table 1 and Fig. 3.

## 5. CONCLUSION

In conclusion, we experimented on various frequency and channel attention methods to enhance SED performance while minimizing computational cost. The study addressed the challenge of effectively addressing the frequency dimension of audio data by leveraging attention methods. The attention methods demonstrated comparable performance to the previous approach of FDY conv, while reducing the computational cost and improving practicality. In addition, we performed class-wise performance of the attention methods to further analyze the characteristics of SED models with different attention methods. Future research could aim to optimize the proposed attention methods by applying them jointly with FDY conv either to push the performance even more or to find balance between computational cost and the performance.

## 6. ACKNOWLEDGMENT

We would like to thank Junhyeok Lee from Supertone Inc. for valuable discussions.

## 7. REFERENCES

- [1] T. Virtanen, M. D. Plumbley, and D. Ellis, *Computational Analysis of Sound Scenes and Events*, 1st ed. Springer Publishing Company, Incorporated, 2017, pp. 3–11, 71–77.
- [2] E. Çakır, G. Parascandolo, T. Heittola, H. Huttunen, and T. Virtanen, “Convolutional recurrent neural networks for polyphonic sound event detection,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 25, no. 6, pp. 1291–1303, 2017.
- [3] N. Turpault. Dcase2021 task4 baseline. GitHub. Available: [https://github.com/DCASE-REPO/DESED\\_task](https://github.com/DCASE-REPO/DESED_task). [Online]. Available: [https://github.com/DCASE-REPO/DESED\\_task](https://github.com/DCASE-REPO/DESED_task)
- [4] H. Nam, S.-H. Kim, B.-Y. Ko, and Y.-H. Park, “Frequency Dynamic Convolution: Frequency-Adaptive Pattern Recognition for Sound Event Detection,” in *Proc. Interspeech*, 2022.
- [5] H. Nam, S.-H. Kim, and Y.-H. Park, “Filteraugument: An acoustic environmental data augmentation method,” in *International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2022.
- [6] S. Xiao, X. Zhang, and P. Zhang, “Multi-dimensional frequency dynamic convolution with confident mean teacher for sound event detection,” in *International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2023.
- [7] J. Li, Y. Tian, and T. Lee, “Convolution-based channel-frequency attention for text-independent speaker verification,” in *International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2023, pp. 1–5.
- [8] K. He, X. Shu, S. Jia, and Y. He, “Semi-supervised sound event detection system for dcase 2022 task 4,” DCASE2022 Challenge, Tech. Rep., 2022.
- [9] S. Suh and D. Y. Lee, “Data engineering for noisy student model in sound event detection,” DCASE2022 Challenge, Tech. Rep., 2022.
- [10] S. Xiao, “Pretrained models in sound event detection for dcase 2022 challenge task4,” DCASE2022 Challenge, Tech. Rep., 2022.
- [11] T. Khandelwal, R. K. Das, A. Koh, and E. S. Chng, “Leveraging audio-tagging assisted sound event detection using weakified strong labels and frequency dynamic convolutions,” *arXiv preprint arXiv:2304.12688*, 2023.
- [12] L. Xu, L. Wang, S. Bi, H. Liu, and J. Wang, “Semi-supervised sound event detection with pre-trained model,” in *International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2023.
- [13] Y. Chen, X. Dai, M. Liu, D. Chen, L. Yuan, and Z. Liu, “Dynamic convolution: Attention over convolution kernels,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, June 2020.
- [14] J. Hu, L. Shen, and G. Sun, “Squeeze-and-excitation networks,” in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018.
- [15] J. Thienpondt, B. Desplanques, and K. Demuynck, “Integrating frequency translational invariance in tdnns and frequency positional information in 2d resnets to enhance speaker verification,” in *Proc. Interspeech*, 2021, pp. 2302–2306.
- [16] B.-Y. Ko, H. Nam, S.-H. Kim, D. Min, S.-D. Choi, and Y.-H. Park, “Data augmentation and squeeze-and-excitation network on multiple dimension for sound event localization and detection in real scenes,” DCASE2022 Challenge, Tech. Rep., 2022.
- [17] H. Zhang, M. Cisse, Y. N. Dauphin, and D. Lopez-Paz, “mixup: Beyond empirical risk minimization,” in *International Conference on Learning Representations*, 2018.
- [18] 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,” in *Proc. Interspeech*, 2019, pp. 2613–2617.
- [19] 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., 2021.
- [20] A. Tarvainen and H. Valpola, “Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results,” in *Advances in Neural Information Processing Systems*, vol. 30, 2017.
- [21] Ç. Bilen, G. Ferroni, F. Tuveri, J. Azcarreta, and S. Krstulović, “A framework for the robust evaluation of sound event detection,” in *International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2020, pp. 61–65.
- [22] A. Mesaros, T. Heittola, and T. Virtanen, “Metrics for polyphonic sound event detection,” *Applied Sciences*, vol. 6, no. 6, 2016.
- [23] S.-H. Kim, H. Nam, and Y.-H. Park, “Temporal dynamic convolutional neural network for text-independent speaker verification and phonemetic analysis,” in *International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2022.