

Frequency & Channel Attention Network for Small Footprint Noisy Spoken Keyword Spotting

Yuanxi Lin and Yuriy Evgenyevich Gapanjuk
Bauman Moscow State Technical University, Russia
linyux@student.bmstu.ru, gapyubmstu@bmstu.ru

Abstract—In this paper, we aim to improve the robustness of Keyword Spotting (KWS) systems in noisy environments while keeping a small memory footprint. We propose a new convolutional neural network (CNN) called FCA-Net, which combines mixer unit-based feature interaction with a two-dimensional convolution-based attention module. First, we introduce and compare lightweight attention methods to enhance noise robustness in CNN. Then, we propose an attention module that creates fine-grained attention weights to capture channel and frequency-specific information, boosting the model’s ability to handle noisy conditions. By combining the mixer unit-based feature interaction with the attention module, we enhance performance. Additionally, we use a curriculum-based multi-condition training strategy. Our experiments show that our system outperforms current state-of-the-art solutions for small-footprint KWS in noisy environments, making it reliable for real-world use.

I. INTRODUCTION

Keyword Spotting (KWS) [1] is a technology that detects specific keywords in audio streams. It aims to recognize trigger words or commands, allowing smart devices to work without physical contact. This technology is widely used in consumer electronics such as smart speakers, wearables, voice assistants, and Internet of Things devices [2]. We focus on small-footprint KWS in this paper, which requires a small memory footprint and low computational power. This kind of technology is practical because the KWS models are always powered on low-resource devices to provide a constant service.

Recent small-footprint KWS models based on the deep neural network (DNN) have achieved success with less noisy, close-talking audio datasets. Chen et.al [3] first utilize the DNN to regard keyword spotting as a classification problem. Tang and Lin [4] proposed a residual network (ResNet) based KWS system where they used dilated convolution to enlarge the size of the receptive field exponentially with the depth of the network. Choi et.al [5] propose a temporal convolutional neural network (TC-ResNet) for real-time KWS on mobile devices, with temporal convolution reducing computation and increasing performance in terms of accuracy compared to 2D convolutions. However, these systems face significant challenges in noisy environments because background noise can interfere with the detection of keywords, leading to false activations or missed detections.

Improving noise robustness in KWS systems is a crucial task to ensure reliable performance [6]. The motivation to handle this kind of environment is driven by the need for KWS

to function accurately in real-world settings where noise is often present, such as in busy households, public spaces, and industrial areas. Enhancing noise robustness will make KWS systems more effective and user-friendly, ensuring they can be relied upon in diverse and challenging auditory conditions.

To augment the noise robustness of acoustic models, curriculum learning has been proposed [7]. As a learning method, curriculum learning increases the difficulty of training samples progressively. It starts with initial training on easy, clean samples and then continues with difficult samples containing loud noise, similar to how a person studies according to a curriculum. Curriculum learning is widely used for noise-robust KWS [8].

Recent work on Ng et.al [9], [10] for small-footprint KWS addresses the challenge of noisy environments. They propose a new encoder architecture based on a mixer module. This mixer computes weighted feature interactions across global channels, improving information flow. The model ConvMixer shows good performance in noisy, far-field conditions. However, there is still room to improve with efficient attention modules to extract local features, such as SE [11] and ECA [12]. This improvement is important because local features can enhance the model’s ability to handle varying noise levels and complex audio inputs, leading to more accurate keyword spotting.

In this paper, we aim to enhance the robustness of KWS in noisy environments while ensuring a small memory footprint. We propose a novel network called FCA-Net combining feature interaction in ConvMixer with an efficient attention module. First, we introduce and compare various lightweight attention methods to improve the noise robustness of KWS systems. Then, we propose an efficient two-dimensional convolution-based attention module for noisy KWS. This module produces fine-grained attention weights to capture channel and frequency-specific information, enhancing the model’s representation ability for noise-robust speech. Additionally, we implement a curriculum-based multi-condition training strategy. Our experiments show that our system outperforms existing state-of-the-art solutions for small-footprint KWS in noisy conditions.

II. RELATED WORK

Small Footprint Keyword Spotting - With the widespread adoption of voice interfaces in smart consumer electronics, the application of small convolutional neural networks (CNNs) in

compact keyword spotting has become increasingly significant. Recent works investigated innovative convolution techniques to improve KWS performance. TC-ResNet [5], [13] applies 1D temporal convolution to enhance efficiency and accuracy. BC-ResNet [14] introduces broadcasted residual learning combining 1D and 2D convolutions. However, most studies do not consider noise robustness, which is a crucial factor for successful voice interfaces.

Noise robust Spoken Keyword Spotting - Data augmentation, especially noise mixture training, has been widely used to enhance the noise robustness of acoustic models [15]. Noise mixture training involves adding environmental noise to the training data, enabling the model to better handle noisy environments. However, small neural networks struggle to learn in high-noise conditions. To address this issue, curriculum learning has been introduced, which gradually increases the difficulty of training samples to improve the model’s robustness [7]. Although these methods have improved performance to some extent, they still fall short when dealing with high-noise environments.

Attention mechanisms in CNNs - Attention mechanism address the limitations of traditional CNNs in handling long-range dependencies and feature selection. It aims to enable the model to automatically focus on more important information regions. SE-Net [11], [16] introduces an effective mechanism for learning channel attention, achieving significant and promising performance outcomes for the first time. Subsequently, Woo et al. [17] proposed the Convolutional Block Attention Module (CBAM), which further enhanced model performance by combining channel attention and spatial attention. Attention mechanisms, while powerful, often come with significant computational and memory overheads. These overheads can hinder the deployment of models on edge devices. Consequently, recent research has focused on developing light attention mechanisms. ECA-Net [12] simplifies the attention mechanism by using an efficient channel attention module that avoids dimensionality reduction and performs local cross-channel interaction without significantly increasing the computational complexity. However, integrating these attention mechanisms with the noise-robust spoken keyword spotting systems is still under exploration.

III. METHODOLOGY

To overcome the difficulty of learning the noise-robustness property in small-footprint keyword spotting, we propose a joint framework that combines ConvMixer block and attention mechanism.

A. ConvMixer architecture

The ConvMixer model architecture is divided into three main sections: the pre-convolutional block, the ConvMixer block, and the post-convolutional block. Figure II provides an overview of this architecture.

The pre-convolutional and post-convolutional block includes 1D depthwise separable (DWS) convolution, batch normalization, and swish activation.

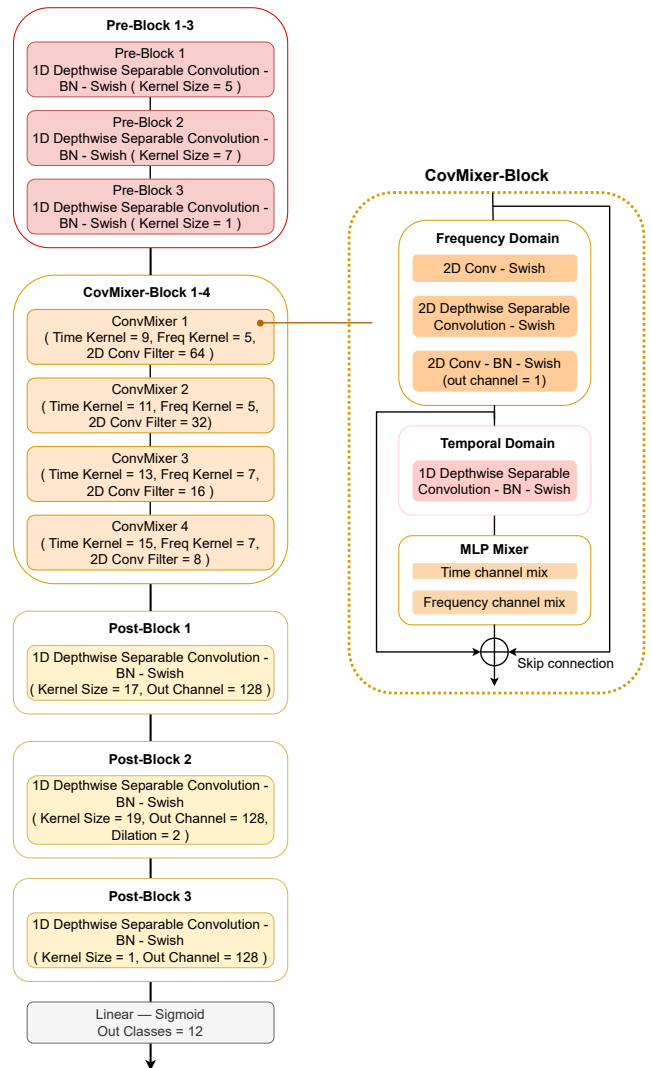


Fig. 1. Overview of ConvMixer model architecture.

The ConvMixer block integrates 2D and 1D DWS convolutions to extract features from both the frequency and temporal domains. It also incorporates a mixer layer to enhance information flow. The following equations describe the ConvMixer block:

$$z = \sigma \circ f_1(\sigma \circ f(x)) \quad (1)$$

In the ConvMixer block, f is a 2D-DWS convolution, and σ is the Swish activation function. First, the input x is convolved by the 2D-DWS convolution function f , and then the result is processed by the activation function σ to obtain the intermediate features. These intermediate features are convolved again by the 2D-DWS convolution f , and processed by the activation function σ , resulting in the frequency domain features z .

$$y_1 = \sigma \circ \text{BatchNorm}(f(z)) \quad (2)$$

Then the frequency domain features z is convolved by a




Fig. 2. An illustration of SE, ECA attention. GAP means global average pooling.

1D-DWS function f_2 . The result is processed by a batch normalization (BatchNorm) layer and finally activated by the function σ , resulting in the temporal domain features y_1 .

$$y_2 = \sigma \circ \text{BatchNorm}(f_2(y_1)) \quad (3)$$

The temporal domain features y_1 are processed by another 1D-DWS convolution function f_2 , then passed through a BatchNorm layer and activation, resulting in the processed features y_2 .

$$\tilde{y} = x + y_1 + f_3(y_2) \quad (4)$$

Finally, the input x , intermediate features y_1 , and processed features y_2 are added together to obtain the output \tilde{y} of the ConvMixer module. Eq (4) computes the processed features y_2 of the block with f_3 as the mixer layer. We propose using two types of multi-layer perceptrons (MLP): temporal channel mixing and frequency channel mixing, to induce interaction between the feature space. Each MLP mixing involves two linear layers and a GELU activation unit, independently processing each temporal and frequency channel.

B. Efficient attention module

1) *Squeeze-and-excitation block (SE)*: The Squeeze-and-Excitation (SE) block [11] is an efficient attention mechanism used in CNNs to improve performance. As shown in Figure 2 (a), the SE block has two main operations: Squeeze and Excitation. The Squeeze operation performs global average pooling on the feature maps across the channel dimension to create a compact intermediate representation. The Excitation operation uses two fully connected (FC) layers to assign importance weights to each channel. The SE module re-weights the output of each channel, highlighting important features and suppressing less significant ones. This helps the network focus on the most informative features, enhancing performance with minimal additional computational cost.

2) *Efficient channel attention (ECA)*: The Efficient Channel Attention (ECA) [12] block is a highly efficient channel attention mechanism for deep CNNs. As illustrated in Figure 3, the ECA block first uses global average pooling to aggregate features. It then applies a 1D convolution with an adaptively

determined kernel size, generating channel weights through a Sigmoid function. Unlike other methods, the ECA block avoids dimensionality reduction via fully connected layers, maintaining a direct correspondence between channels and their weights.

The ECA block captures local cross-channel interactions through 1D convolution, improving performance while maintaining efficiency. The kernel size k of this 1D convolution is adaptively determined based on the channel dimension C , ensuring that higher-dimensional channels interact more broadly, while lower-dimensional channels interact more narrowly. The ECA block achieves effective channel attention with low model complexity, significantly enhancing performance and demonstrating strong generalization in practical tasks.

3) *Convolution-based two-dimensional attention (C2D)*: The C2D block is integrated into CNNs to calculate fine-grained attention weights for various channel and frequency positions. This module enhances the efficiency and effectiveness of the attention mechanism by using convolution operations instead of the fully connected layers typically used in traditional attention mechanisms which shown in Figure 4.

The input feature map X of C2D block contains C channels, F frequency bins, and T time frames. This feature map is the output of a previous convolutional layer in the network. Then the Global Average Pooling (GAP) is applied along the time dimension of the input feature map X .

$$z(c, f) = \frac{1}{T} \sum_{t=1}^T X(c, f, t) \quad (5)$$

The result is a channel-frequency plane z with dimensions $C \times F$, where each element represents the average value of the corresponding channel and frequency bin across the time dimension. Next, the pooled channel-frequency plane z is passed through two consecutive 2D convolutional layers. These convolution operations are used to calculate the attention weights ω for different channel and frequency positions. The convolutional layers include BatchNorm and ReLU activation functions to enhance learning ability and stability. The attention weights ω represent the importance of each channel-frequency pair.




Fig. 3. The network structure of the C2D block.

$$\omega = \sigma(\text{Conv2}(\text{ReLU}(\text{BN}(\text{Conv1}(z)))))) \quad (6)$$

The calculated attention weights ω are then applied to the original input feature map X through element-wise multiplication. This step re-scales the input features, enhancing important features and suppressing less important ones as the following equation:

$$X' = \omega \odot X \quad (7)$$

This operation re-scales the input features, emphasizing important channel-frequency pairs and suppressing less important ones, thereby enhancing the representation capability of the network.

The C2D block calculates attention weights using convolution operations, which require fewer parameters and less computational cost compared to fully connected layers. By simultaneously considering both channel and frequency dimensions, the C2D block generates more detailed and specific attention weights, enhancing feature representation capability.

C. Proposed FCA-Net

Our proposed FCA-Net combines the ConvMixer with the C2D attention mechanism for keyword spotting. We add the C2D attention after each ConvMixer block, referring to this configuration as FCA-Net-all. Additionally, we explore other variants: FCA-Net-pre, which adds C2D attention only to the pre-convolutional blocks; FCA-Net-post, which adds C2D attention only to the post-convolutional blocks; and FCA-Net-final, which adds C2D attention before the final linear layer to study its impact on the final classification. We setup various experiments for different variants in the ablation study.

IV. EXPERIMENT SETTING

A. Dataset for Noisy Keyword Spotting

We evaluate our proposed system on the Google Speech Commands V2 dataset, which contains 105,000 utterances of 35 unique words, each 1 second long and sampled at 16 kHz. We use the official train, validation, and test split for the 12-label classification task. These labels include: ‘up’, ‘down’, ‘left’, ‘right’, ‘yes’, ‘no’, ‘on’, ‘off’, ‘go’, ‘stop’, ‘silence’, and ‘unknown’. The ‘unknown’ class includes the remaining words in the dataset.

To simulate a noisy far-field environment, we use noise samples from the MUSAN dataset, which has 930 files of

assorted noises, totaling about 6 hours, all sampled at 16 kHz. These noises include technical and non-technical sounds like DTMF tones, thunder, and car horns. We add these noises to our command samples to mimic various noisy conditions. This process is divided into four progressively harder steps. Initially, the model is trained on clean samples without noise. In the following three steps, noise is incrementally introduced to the fixed N samples in decrements of -5dB, with conditions in N samples uniformly distributed, i.e., [clean, 0], [clean, 0, -5], [clean, 0, -5, -10].

B. Implement details

We use Mel-Frequency Cepstral Coefficients (MFCC) with the following parameters: sample rate of 16 kHz, 40 MFCCs, and mel-spectrogram settings of $n_fft=400$, $hop_length=160$, $n_mels=64$, $f_min=20$, and $f_max=8000$. Each command is 1 second long; shorter commands are zero-padded to the right. During training, we perform data augmentation with time shifts of -100 to 100 ms and spectrogram masking with maximum lengths of 25 for both time and frequency. We generate noisy data with SNR of 0, -5, and -10 dB.

For stronger regularization, we use input mixup with a ratio of 0.2 on training samples. The model is trained with a batch size of 128 and an initial learning rate of 0.005, reduced by a factor of 0.85 every four epochs after the fifth epoch. We use the Adam optimizer and a learning rate scheduler with Cosine Annealing Warm Restarts. The model is trained for 200 epochs with early stopping if there is no improvement for 20 epochs. This study applied three augmentation methods: mixup [18], SpecAugment [19], and curriculum learning [7], [9].

C. Baselines

We also set up some SOTA keyword spotting models as the baseline to analyze our proposed model. Here are the models we used for experiments:

- **MHAtt-RNN:** This model integrates multi-head attention mechanisms with recurrent neural networks to enhance sequence learning and capture long-range dependencies.
- **KWT:** Keyword Transformer (KWT-1) uses transformer architecture to identify keywords, leveraging self-attention for efficient feature extraction. KWT-3 improves upon KWT-1 by incorporating more layers and sophisticated attention mechanisms for enhanced keyword detection.
- **AST-Tiny:** Audio Spectrogram Transformer-Tiny (AST-Tiny) employs a transformer architecture specifically designed for processing audio spectrograms, focusing on smaller model sizes.
- **BC-ResNet-6:** Broadcasted Residual Learning Network (BC-ResNet-6) incorporates broadcasted residual connections over 6 layers to enhance the learning of spatial features.
- **ResNet-15:** A variant of the ResNet architecture with 15 layers, ResNet-15 employs residual connections to facilitate deeper networks and improve feature learning.

TABLE I
COMPARISON WITH THE SOTA MODELS.

Model	Num. of Params (K)	MACs (M)	Acc. of V2-12 Official (%)	Accuracy of Noisy Test Command(%)				
				Clean	20 dB	0 dB	-5 dB	-10 dB
MHAtt-RNN	784	141.6	98.04	78.83	74.25	61.98	55.87	50.98
KWT-1	607	53.8	97.72	88.73	85.46	73.52	67.59	59.07
ResNet-15	238	961.2	96.48	89.45	87.34	79.00	73.58	66.73
MatchboxNet-6x2x64	140	36.8	97.60	87.34	85.19	75.58	70.06	62.35
BC-ResNet-6	202	55.4	98.20	93.67	93.30	87.48	83.06	79.40
ConvMixer	119	22.2	98.20	96.32	95.52	90.62	86.37	77.47
FCA-Net (Ours)	119	22.3	98.25	96.96	95.99	92.65	87.97	80.40
AST-Tiny	5,805	782.2	97.65	91.02	87.71	83.31	78.95	72.32
KWT-3	5,361	526.3	98.54	93.47	91.08	83.97	78.45	71.08

TABLE II
EXPLORATION OF DIFFERENT ATTENTION MODULES AND POSITIONS OF THE PERFORMANCE

Attention	Position	Accuracy of Noisy Test Command (%)				
		Clean	20 dB	0 dB	-5 dB	-10 dB
-	-	96.32	95.52	90.62	86.37	77.47
ECA	pre	95.23	94.49	89.37	85.24	76.97
	post	96.77	95.83	90.75	86.20	78.12
	all	96.73	95.68	91.86	86.60	80.00
	final	96.67	95.85	91.14	85.90	78.27
SE	pre	95.23	94.59	89.70	84.50	75.16
	post	94.70	94.22	89.64	84.89	76.85
	all	95.50	94.86	89.82	85.20	76.42
	final	95.20	94.10	89.97	85.51	76.75
C2D	pre	95.70	94.47	89.08	84.07	75.88
	post	95.27	94.65	89.02	84.00	75.56
	all	96.96	95.99	92.65	87.97	80.40
	final	96.92	96.09	91.18	86.82	78.42

- **MatchboxNet-6×2×64:** This model uses a combination of 6 blocks, each containing 2 convolutional layers with 64 filters, optimized for efficient keyword spotting.
- **ConvMixer:** ConvMixer combines convolution operations with mixer layers to extract and mix features across both frequency and temporal dimensions.

V. RESULTS

A. Compare study among the SOTA models

We compare the performance of FCA-Net with previously proposed SOTA models. Models are retrained from the official source code provided with our designed data environment. The results are shown in Table II. From the table, we observe that our proposed model achieved the highest accuracy among small models when tested on the official V2-12, with an accuracy improvement of approximately 2% to 3% over other small models like ConvMixer. Furthermore, FCA-Net shows a noticeable reduction in the number of model parameters and MACs, indicating lower memory and computation resource requirements. Most importantly, when evaluated in noisy conditions, FCA-Net demonstrates superior performance, achieving up to 7.4% better accuracy compared to other models with similar memory footprints. The proposed model is also competitive with larger transformer-based models (KWT-3, AST-Tiny) under challenging noisy conditions, showcasing its efficiency and reliability for real-world applications.

B. Ablation study

The results indicate that incorporating efficient attention modules significantly enhances the robustness of KWS systems in noisy environments. Among the different positions tested, the C2D attention module demonstrates superior performance across various noise levels. Notably, the C2D module achieves the highest accuracy when placed at all positions within the ConvMixer block, consistently outperforming the other attention methods, including ECA and SE. This setup proves to be the most effective strategy for enhancing noise robustness, making the C2D module particularly valuable for improving KWS system performance in challenging auditory conditions.

VI. CONCLUSIONS

In this paper, we proposed FCA-Net, a novel convolutional neural network designed to enhance the robustness of Keyword Spotting (KWS) systems in noisy environments while maintaining a small memory footprint. Our approach combines mixer unit-based feature interaction with an efficient two-dimensional convolution-based attention module, further strengthened by a curriculum-based multi-condition training strategy. Experimental results demonstrate that FCA-Net achieves state-of-the-art accuracy among small models and outperforms existing solutions in noisy conditions. Additionally, FCA-Net significantly reduces the number of parameters and computational resources required, making it both efficient and effective for real-world applications. Our findings highlight the potential of FCA-Net as a reliable and resource-efficient solution for robust KWS in diverse and challenging auditory environments.

REFERENCES

- [1] I. López-Espejo, Z.-H. Tan, J. H. L. Hansen, and J. Jensen, “Deep spoken keyword spotting: An overview,” *IEEE Access*, vol. 10, pp. 4169–4199, 2022.
- [2] L. Turchet, G. Fazekas, M. Lagrange, H. S. Ghadikolaei, and C. Fischione, “The internet of audio things: State of the art, vision, and challenges,” *IEEE Internet of Things Journal*, vol. 7, no. 10, pp. 10 233–10 249, 2020.
- [3] G. Chen, C. Parada, and G. Heigold, “Small-footprint keyword spotting using deep neural networks,” in *Proc. IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, IEEE, 2014, pp. 4087–4091.
- [4] R. Tang and J. Lin, “Deep residual learning for small-footprint keyword spotting,” in *Proc. IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, IEEE, 2018, pp. 5484–5488.
- [5] S. Choi, S. Seo, B. Shin, *et al.*, “Temporal convolution for real-time keyword spotting on mobile devices,” in *Proc. Interspeech 2019*, 2019, pp. 3372–3376.
- [6] Y. Lin, T. Zhou, and Y. Xiao, “Advancing airport tower command recognition: Integrating squeeze-and-excitation and broadcasted residual learning,” *arXiv preprint arXiv:2406.18313*, 2024.
- [7] S. Braun, D. Neil, and S.-C. Liu, “A curriculum learning method for improved noise robustness in automatic speech recognition,” in *Proc. European Signal Processing Conference (EUSIPCO)*, IEEE, 2017, pp. 548–552.
- [8] T. Higuchi, S. Saxena, M. Souden, T. D. Tran, M. Delfarah, and C. Dhir, “Dynamic curriculum learning via data parameters for noise robust keyword spotting,” in *Proc. IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, IEEE, 2021, pp. 6848–6852.
- [9] D. Ng, Y. Chen, B. Tian, Q. Fu, and E. S. Chng, “Conv-mixer: Feature interactive convolution with curriculum learning for small footprint and noisy far-field keyword spotting,” in *Proc. IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, IEEE, 2022, pp. 3603–3607.
- [10] D. Ng, Y. Xiao, J. Q. Yip, *et al.*, “Small footprint multi-channel network for keyword spotting with centroid based awareness,” in *Proc. Interspeech*, 2023, pp. 296–300.
- [11] J. Hu, L. Shen, and G. Sun, “Squeeze-and-excitation networks,” in *Proc. IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018, pp. 7132–7141.
- [12] Q. Wang, B. Wu, P. Zhu, P. Li, W. Zuo, and Q. Hu, “Eca-net: Efficient channel attention for deep convolutional neural networks,” in *Proc. IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020, pp. 11 534–11 542.
- [13] Y. Xiao, N. Hou, and E. S. Chng, “Rainbow Keywords: Efficient Incremental Learning for Online Spoken Keyword Spotting,” in *Proc. Interspeech*, 2022, pp. 3764–3768.
- [14] B. Kim, S. Chang, J. Lee, and D. Sung, “Broadcasted Residual Learning for Efficient Keyword Spotting,” in *Proc. Interspeech 2021*, 2021, pp. 4538–4542.
- [15] V. A. Trinh, H. S. Kavaki, and M. I. Mandel, “Importantaug: A data augmentation agent for speech,” in *Proc. IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, IEEE, 2022, pp. 8592–8596.
- [16] Y. Xiao and R. K. Das, “Dual knowledge distillation for efficient sound event detection,” *arXiv preprint:2402.02781*, 2024.
- [17] S. Woo, J. Park, J.-Y. Lee, and I. S. Kweon, “Cbam: Convolutional block attention module,” in *Proc. European Conference on Computer Vision (ECCV)*, 2018, pp. 3–19.
- [18] H. Zhang, M. Cisse, Y. N. Dauphin, and D. Lopez-Paz, “Mixup: Beyond empirical risk minimization,” in *Proc. International Conference on Learning Representations*, 2018.
- [19] D. S. Park, W. Chan, Y. Zhang, *et al.*, “SpecAugment: A simple data augmentation method for automatic speech recognition,” *Proc. Interspeech*, p. 2613, 2019.