

DIVIDED SPECTRO-TEMPORAL ATTENTION FOR SOUND EVENT LOCALIZATION AND DETECTION IN REAL SCENES FOR DCASE2023 CHALLENGE

Technical Report

Yusun Shul¹, Byeong-Yun Ko², Jung-Woo Choi¹

Korea Advanced Institute of Science and Technology,
¹ School of Electrical Engineering, ² Dept. of Mechanical Engineering,
 Yuseong-gu, Daejeon, Republic of Korea
 {shulys, b.y.ko, jwoo}@kaist.ac.kr

ABSTRACT

Localizing sounds and detecting events in different room environments is a difficult task, mainly due to the wide range of reflections and reverberations. When training neural network models with sounds recorded in only a few room environments, there is a tendency for the models to become overly specialized to those specific environments, resulting in overfitting. To address this overfitting issue, we propose divided spectro-temporal attention. In comparison to the baseline method, which utilizes a convolutional recurrent neural network (CRNN) followed by a temporal multi-head self-attention layer (MHSA), we introduce a separate spectral attention layer that aggregates spectral features prior to the temporal MHSA. To achieve efficient spectral attention, we reduce the frequency pooling size in the convolutional encoder of the baseline to obtain a 3D tensor that incorporates information about frequency, time, and channel. As a result, we can implement spectral attention with channel embeddings, which is not possible in the baseline method dealing with only temporal context in the RNN and MHSA layers. We demonstrate that the proposed divided spectro-temporal attention significantly improves the performance of sound event detection and localization scores for real test data from the STARSS23 development dataset. Additionally, we show that various data augmentations, such as frameshift, time masking, channel swapping, and moderate mix-up, along with the use of external data, contribute to the overall improvement in SELD performance.

Index Terms— Sound Event Localization and Detection, Multi-Head Self-Attention, Convolutional Neural Network, Divided Spectral-Temporal Attention

1. INTRODUCTION

The sound event localization and detection (SELD) task in DCASE challenge task 3 involves the classification of 13 different sound events and the estimation of their direction of arrival (DoA) [1]. Given the diverse nature of audio signals, characterized by a wide range of reflections and reverberations, addressing the issue of overfitting deep neural networks (DNNs) to specific room environments within the training dataset is of utmost importance. In previous years till 2021, the DCASE challenge task 3 has not permitted the utilization of external data. However, the use of external data and the incorporation of various types of augmentation emerged as the key contributions to the challenge of the DCASE2022 challenge.

On the other hand, there is a growing trend of utilizing attention

mechanisms separately across different dimensions, as evidenced in studies such as [2], [3], and [4]. In the domain of video classification, [2] extensively investigated different forms of attention and concluded that allocating attention to both temporal and spatial aspects independently yields valuable results. Similarly, for speech enhancement tasks, DASformer [3] proposed the application of frame-wise spectral attention and band-wise temporal attention. Furthermore, DeFT-AN [4] introduced a block composed of three distinct layers: a channel-related dense block, a frequency-related conformer, and a temporal conformer, which were tailored for speech enhancement.

Hence, we adopt the concept of divided attention to the spectro-temporal dimension in the SELD task and demonstrate its improvement over the baseline model that solely considers temporal context. We also utilize both the external data synthesized with AudioSet [5] and augmentation methods such as frameshift [6], time masking [7], channel swap [8], and moderate mix-up [9], to improve the SELD performance along with the proposed DNN model.

2. PROPOSED METHOD

2.1. Implementation Details

Among the first-order Ambisonics (FoA) and tetrahedral microphone (MIC) audio formats available in STARSS23, the proposed model utilizes the FoA array signals. The audio signals having a sampling frequency of 24 kHz underwent a short-time Fourier transform with a hop length of 0.02s and a window size of 0.04s. Subsequently, the spectrograms were transformed into log-mel scales using 64 mel filter banks. Also, intensity vectors were calculated by multiplying the complex conjugation of the omnidirectional spectrograms, with the other directional array spectrograms, and taking the real part of the multiplication. These intensity vectors were then normalized using the total energy of the FoA signals. Therefore, the total number of input features was 7 channels consisting of four original spectrograms and three intensity components. Target DoA labels were converted into ACCDOA [10] and multi-ACCDOA [11] formats. Experimental evaluations were carried out for both target formats. The model was trained for 500 epochs with a batch size of 128. The Adam optimizer was employed, using a learning rate of 0.001, $\beta_1 = 0.9$, and $\beta_2 = 0.999$. The input time length for the model was set to 5 seconds.

Table 1: Features of audio, preprocessing, and input

Audio Features	
Audio Format	FoA (4 channels)
Sampling Frequency	24kHz
Preprocessing Features	
Hop Size	0.02s
Window Size	0.04s
Number of Mel Filter Banks	64
Target Label Resolution	0.1s
Input Features	
Features	Log-mel, Intensity Vectors
Batch Size (B)	128
Input Feature Length (T)	250
Input Channel (C)	7
Mel Frequency Bins (F)	64

2.2. External Data

STARSS23 [12] is an expanded version of STARSS22 [13] additionally including real recordings captured in new room environments. While this additional data introduce more room diversity to the development training (dev-train) dataset, the development test (dev-test) data also have more diversity. Consequently, it remains challenging to prevent the overfitting of a model to the room impulse responses (RIRs) present in the training dataset. Therefore, it is still necessary to use an external dataset to mitigate overfitting on the training data. In this regard, AudioSet [5] was utilized to synthesize external data through the convolution with RIRs in TAU-SRIR [14] using the generator from [15], in addition to the official synthetic dataset [16] based on FSD50k [17]. The balanced dataset from AudioSet [5] was utilized to ensure an equitable number of synthetic data across different classes. A total of 6 hours of the additional dataset with the maximum number of overlapping objects of two were synthesized and used for the training process along with the official synthetic data [16].

2.3. Data Augmentation

Since synthetic data often exhibit different data distributions compared to real recordings, it is beneficial to incorporate data augmentation along with external data. In light of this, we integrated various augmentations that have been proven to be effective for the SELD task, as demonstrated in [9]. These augmentations include frameshift [6], which involves shifting both the feature and label along the time axis using a randomly determined shifting parameter. Additionally, we employed time masking [7], randomly masking features and DoA labels along the time axis. Furthermore, channel swap [8] was utilized as a spatial augmentation method. Lastly, we incorporated moderate mix-up [9], an enhanced version of mix-up [18] specifically designed for the SELD task. This technique randomly selects a mixing ratio to blend two spectrograms and selects only one DoA with the larger mixing ratio as the new target label. All four augmentation methods were randomly selected and applied with random parameters generated for every batch and iteration.

2.4. Divided Spectro-Temporal (DST) Attention

The baseline model for DCASE2023 challenge task 3 [19] incorporates several key components: convolutional blocks, bi-

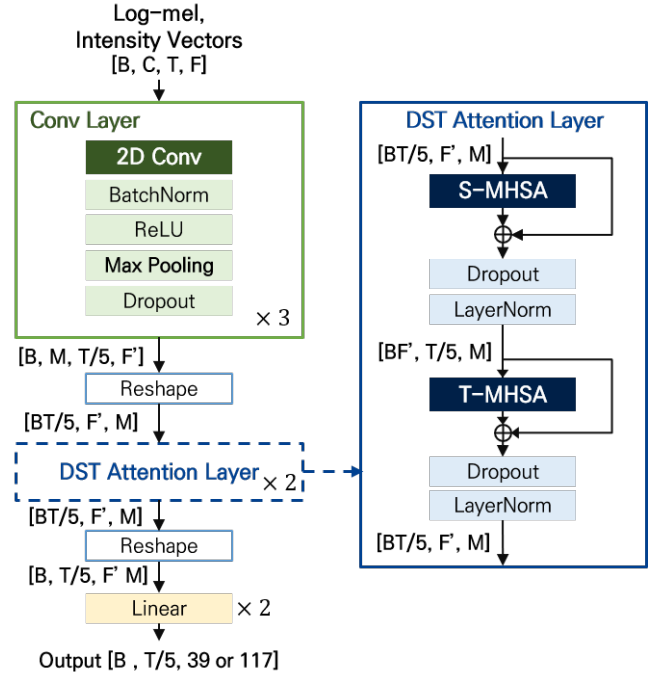


Figure 1: Architecture of Divided Spectro-Temporal Attention Model

directional gated recurrent units (GRU), and multi-head self-attentions (MHSA). The convolutional blocks are responsible for amplifying the input channels from seven to 64, while also employing max pooling after each convolutional operation to reduce the spectral dimension to two. Additionally, to align the temporal information with the number of target labels, time pooling is utilized. The size of the time pooling kernel for the three convolutional layers is set to $[5, 1, 1]$, respectively. The GRU and MHSA layers are specifically designed to capture and understand the temporal context from the spectral and channel embeddings.

As a modification to enable the model to learn spectral information from the embeddings, we propose divided spectro-temporal attention for SELD, as depicted in Fig. 1. This approach allows independent spectral attention from the channel embedding ($M = 64$), separate from temporal attention. We reduced the kernel size of the frequency pooling layer to retain spectral information that can be used in the attention layer. The output from the convolution block of size (batch, channel, time, frequency) = $[B, M, T/5, F']$ was then permuted for the divided spectro-temporal (DST) attention layer, such that the temporal dimension is merged to the batch dimension ($[BT/5, F', M]$). Then, spectral-MHSA (S-MHSA) was applied through attention maps of size $(F' \times F')$. The output from the S-MHSA layer was then merged with the input via the skip connection. The merged signal was then processed by the dropout layer and layer normalization. For the temporal-MHSA (T-MHSA), the frequency and time dimensions were swapped to yield the tensor size of $[BF', T/5, M]$. T-MHSA was then applied using attention maps of size $(T/5 \times T/5)$, followed by the addition of attention input, dropout, and layer normalization. The number of heads for MHSA was eight, and DST attention layers were repeated twice. The output of the final DST attention layer was reshaped to have

Table 2: Training Parameters. In frequency pooling, $[c_1, c_2, c_3]$ represents pooling sizes in three convolution layers, respectively.

Model	Frequency Pooling	Spectral Dim. (F')
Baseline	[4,2,2]	2
DST	[4,2,1]	8
Attention	[2,2,1]	16
	[2,1,1]	32
	[1,1,1]	64

the size $[B, T/5, F'M]$ and fed into the two linear layers. The linear layers then reduced the size of embedding from $F'M$ to that of the target embedding, which is 39 for ACCDOA and 117 for multi-ACCDOA. The size of the spectral attention affects the SELD performance as presented in Table 4. The size of spectral dimension was varied ($F' = 8, 16, 32, 64$) by reducing the frequency pooling size as presented in Table 2. The final model was determined to have $F' = 16$.

3. RESULTS

The investigation results on DST attention, augmentation methods, and utilization of external data are summarized in Table 3. The results demonstrate that incorporating DST attention significantly enhances the SELD performance compared to the baseline model for both the single-ACCDOA and multi-ACCDOA tasks. Moreover, the performance consistently improves when augmentation techniques are applied, and these gains are further amplified when external data is employed.

Additionally, as discussed in Section 2.4, the size of spectral attention has an impact on SELD performance (Table 4). When the spectral dimension size is increased, the SELD score is improved for all cases of multi-ACCDOA training and up to $F' = 16$ for single-ACCDOA training, compared to the baseline shown in Table 3. However, increasing the spectral dimension beyond $F' = 16$ for the single-ACCDOA yields negative effects. It is worth noting that excessive spectral dimension ($F' > 16$) reduces the performance.

In Table 4, we can see that utilizing multi-ACCDOA outperforms single-ACCDOA when no augmentation or external data are involved with the training. However, with augmentation and external data, employing single-ACCDOA yields better outcomes, as shown in Table 3. Consequently, the final submitted version is the model trained using single-ACCDOA, four types of augmentation techniques, six hours of additional external data, and a spectral dimension size of 16 for the DST attention model.

The SELD performances of the submitted models on the dev-test data are displayed in Table 5. Submission 2 is the model trained up to epoch 500, while submission 1 is a refinement of submission 2 through the training up to 700 epochs. On the other hand, submission 3 is the same model as submission 2 but is trained with different random initializations. It is noteworthy that all three submissions surpass the baseline model in terms of both event detection and localization, showing improved performances in all four metrics.

4. CONCLUSION

This technical report introduces DST attention for the SELD task, which demonstrates superior performance compared to the baseline

Table 3: Experimental results of the proposed DST attention model on dev-test data in comparison with the baseline model.

Model	ACC-DOA	Aug.	Ext. Data	SELD Score	ER	F-score	LE	LR
Baseline	Multi	-	-	0.4791	0.570	29.90	22.00	47.70
	Single	-	-	0.4642	0.615	33.62	22.88	54.91
DST	Multi	-	-	0.4345	0.580	39.50	20.03	55.83
Attention ($F' = 16$)	O	-	-	0.4215	0.525	40.45	17.31	53.06
		O	O	0.4329	0.555	41.48	16.92	50.26
	Single	-	-	0.4423	0.570	37.16	20.30	54.20
		O	-	0.4209	0.505	39.87	17.95	52.24
	O	O	0.4078	0.510	41.96	16.54	55.10	

Table 4: SELD performances of DST attention with various F' on dev-test data (augmentation and external data were not used)

ACCDOA	F'	SELD Score	ER	F-score	LE	LR
Multi	8	0.4422	0.595	37.42	19.48	56.01
	16	0.4345	0.580	39.50	20.03	55.83
	32	0.4517	0.570	35.86	18.24	50.60
	64	0.4523	0.600	37.55	17.89	51.48
Single	8	0.4474	0.590	37.20	19.67	53.78
	16	0.4423	0.570	37.16	20.30	54.20
	32	0.4758	0.600	33.41	18.55	46.60
	64	0.4804	0.590	34.08	19.79	43.75

Table 5: SELD performance of submitted models on dev-test data

Submission	SELD Score	ER	F-score	LE	LR
1	0.4023	0.495	42.68	16.70	55.19
2	0.4078	0.510	41.96	16.54	55.10
3	0.4149	0.515	41.24	17.68	54.12

model across all SELD metrics when combined with suitable frequency pooling. Additionally, the inclusion of augmentation methods such as frameshift, time masking, channel swap, and moderate mix-up effectively tackles the problem of overfitting to the training data. Furthermore, the integration of external data synthesized from the balanced AudioSet dataset proves to be beneficial in accounting for various recording environments.

5. REFERENCES

- [1] A. Politis, K. Shimada, P. Sudarsanam, S. Adavanne, D. Krause, Y. Koyama, N. Takahashi, S. Takahashi, Y. Mitsufuji, and T. Virtanen, "STARSS22: A dataset of spatial recordings of real scenes with spatiotemporal annotations of sound events," in *Proceedings of the 8th Detection and Classification of Acoustic Scenes and Events 2022 Workshop (DCASE2022)*, Nancy, France, November 2022, pp. 125–129. [Online]. Available: <https://dcase.community/workshop2022/proceedings>
- [2] G. Bertasius, H. Wang, and L. Torresani, "Is space-time attention all you need for video understanding?" in *Proceedings of the 38th International Conference on Machine Learning*, ser. Proceedings of Machine Learning Research, M. Meila and T. Zhang, Eds., vol. 139.

- PMLR, 18–24 Jul 2021, pp. 813–824. [Online]. Available: <https://proceedings.mlr.press/v139/bertasius21a.html>
- [3] S. Wang, X. Kong, X. Peng, H. Movassagh, V. Prakash, and Y. Lu, “Dasformer: Deep alternating spectrogram transformer for multi/single-channel speech separation,” in *ICASSP 2023 - 2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2023, pp. 1–5.
- [4] D. Lee and J.-W. Choi, “Deft-an: Dense frequency-time attentive network for multichannel speech enhancement,” *IEEE Signal Processing Letters*, vol. 30, pp. 155–159, 2023.
- [5] J. F. Gemmeke, D. P. W. Ellis, D. Freedman, A. Jansen, W. Lawrence, R. C. Moore, M. Plakal, and M. Ritter, “Audio set: An ontology and human-labeled dataset for audio events,” in *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2017, pp. 776–780.
- [6] 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,” <https://arxiv.org/abs/2107.03649>, 2021.
- [7] 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 *Interspeech 2019*. ISCA, sep 2019. [Online]. Available: <https://doi.org/10.21437%2Finterspeech.2019-2680>
- [8] L. Mazzon, Y. Koizumi, M. Yasuda, and N. Harada, “First order ambisonics domain spatial augmentation for dnn-based direction of arrival estimation,” in *Workshop on Detection and Classification of Acoustic Scenes and Events*, 2019.
- [9] 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,” <https://arxiv.org/abs/2206.12059>, 2022.
- [10] K. Shimada, Y. Koyama, N. Takahashi, S. Takahashi, and Y. Mitsufuji, “Accdoa: Activity-coupled cartesian direction of arrival representation for sound event localization and detection,” in *ICASSP 2021 - 2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2021, pp. 915–919.
- [11] K. Shimada, Y. Koyama, S. Takahashi, N. Takahashi, E. Tsunoo, and Y. Mitsufuji, “Multi-accdoa: Localizing and detecting overlapping sounds from the same class with auxiliary duplicating permutation invariant training,” in *ICASSP 2022 - 2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2022, pp. 316–320.
- [12] A. Politis, K. Shimada, P. Sudarsanam, A. Hakala, S. Takahashi, D. A. Krause, N. Takahashi, S. Adavanne, Y. Koyama, K. Uchida, Y. Mitsufuji, and T. Virtanen, “STARSS23: Sony-TAU Realistic Spatial Soundscapes 2023,” Mar. 2023. [Online]. Available: <https://doi.org/10.5281/zenodo.7880637>
- [13] A. Politis, Y. Mitsufuji, P. Sudarsanam, K. Shimada, S. Adavanne, Y. Koyama, D. A. Krause, N. Takahashi, S. Takahashi, and T. Virtanen, “STARSS22: Sony-TAU Realistic Spatial Soundscapes 2022 dataset,” May 2022. [Online]. Available: <https://doi.org/10.5281/zenodo.6600531>
- [14] A. Politis, S. Adavanne, and T. Virtanen, “TAU Spatial Room Impulse Response Database (TAU- SRIR DB),” Apr. 2022. [Online]. Available: <https://doi.org/10.5281/zenodo.6408611>
- [15] D. Krause and S. Adavanne. (2022) Dcase2022-data-generator. [Online]. Available: <https://github.com/danielkrause/DCASE2022-data-generator>
- [16] A. Politis, “[DCASE2022 Task 3] Synthetic SELD mixtures for baseline training,” Apr. 2022. [Online]. Available: <https://doi.org/10.5281/zenodo.6406873>
- [17] E. Fonseca, X. Favory, J. Pons, F. Font, and X. Serra, “Fsd50k: An open dataset of human-labeled sound events,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 30, pp. 829–852, 2022.
- [18] H. Zhang, M. Cisse, Y. N. Dauphin, and D. Lopez-Paz, “mixup: Beyond empirical risk minimization,” <https://arxiv.org/abs/1710.09412>, 2018.
- [19] S. Adavanne. (2023) seld-dcase2023. [Online]. Available: <https://github.com/sharathadavanne/seld-dcase2023>