

A DATASET OF DYNAMIC REVERBERANT SOUND SCENES WITH DIRECTIONAL INTERFERERS FOR SOUND EVENT LOCALIZATION AND DETECTION

Archontis Politis¹, Sharath Adavanne¹, Daniel Krause¹, Antoine Deleforge²
Prerak Srivastava², Tuomas Virtanen¹,

¹ Audio Research Group, Tampere University, Tampere, Finland

² Universite de Lorraine, CNRS, Inria, LORIA, F-54000 Nancy, France

ABSTRACT

This report presents the dataset and baseline of Task 3 of the DCASE2021 Challenge on Sound Event Localization and Detection (SELD). The acoustical synthesis remains the same as in the previous iteration of the challenge, however the new dataset brings more challenging conditions of polyphony and overlapping instances of the same class. The most important difference is the introduction of directional interferers, meaning sound events that are localized in space but do not belong to the target classes to be detected and are not annotated. Since such interfering events are expected in every real-world scenario of SELD, the new dataset aims to promote systems that deal with this condition effectively. A modified SELDnet baseline employing the recent ACCDOA representation of SELD problems accompanies the dataset and it is shown to outperform the previous one. The new dataset is shown to be significantly more challenging for both baselines according to all considered metrics. To investigate the individual and combined effects of ambient noise, interferers, and reverberation, we study the performance of the baseline on different versions of the dataset excluding or including combinations of these factors. The results indicate that by far the most detrimental effects are caused by directional interferers.

Index Terms— Sound event localization and detection, sound source localization, acoustic scene analysis, microphone arrays

1. INTRODUCTION

Sound event localization and detection (SELD) is an audio processing task that aims to jointly detect temporally target classes of sound events and localize them in space when active. In that sense it differs from the classic sensor array task of sound source localization (SSL) which utilizes only spatial information to detect, localize, and track sources independently from their signal content [1]. It also differs from the popular sound event detection (SED) task which is focused on the temporal detection and classification part, omitting the spatial information of the scene. The spatiotemporal characterization of the scene produced by SELD makes it suitable for a range of applications such as robot audition and machine listening in general [2, 3], acoustic monitoring [4, 5], smart home environments [5, 6], improved human-machine interaction [7], speech recognition [8], and sonic information visualization [9], among others.

Research interest in SELD grew quickly during the last couple of years, with deep learning methods handling the task jointly [10], or fusing information from solving individual subtasks of SED and SSL [11, 12]. This interest culminated in the task becoming part of the DCASE Challenge in 2019, with participants bringing novel approaches to the problem, summarized in [13]. The dataset used in

the challenge [14] included sound scenes from two different array formats with sound events spatialized in both azimuth and elevation using spatial room impulse responses (SRIRs) of real rooms. Additionally, spatial ambient noise captured in situ was added to the recordings. For the next iteration of the task in the DCASE Challenge 2020, a new dataset was generated based on SRIRs from additional rooms with more realistic and challenging conditions beyond the limitations of the first one [14]. More specifically, the discrete grid of potential directions-of-arrival (DOAs) of the older dataset was replaced with continuous DOA trajectories and, apart from static events, moving sources using interpolated SRIRs were emulated at different speeds. Furthermore, the newer SRIRs were captured in rooms of more diverse acoustical properties and from a wider range of distances, resulting in longer reverberation times and more challenging direct-to-reverberant ratios (DRRs).

The second iteration of the SELD task in DCASE2020 brought additional innovations, with participants experimenting with homogeneous joint loss functions [15, 16], self-attention layers [16, 17], advanced spatial augmentation strategies [15, 17], combinations of model-based localization with learning-based SED [18, 19], data-based fusion of individual SSL and SED systems [18, 20], and event- or track-based prediction modeling, instead of class-based prediction [21, 19]. The latter development specifically tried to address the case of same-class events occurring simultaneously [12, 21, 22], a case that distinguishes the SELD task from SED and becomes possible mainly due to spatial information. Research following the DCASE2020 challenge investigated fusion of pre-trained SED and SSL models [20], or parameter sharing between joint, semi-joint SELD models, and models fusing SSL and SED subsystems [22].

This report introduces the new **TAU-NIGENS Spatial Sound Events 2021**¹ dataset and the baseline² of the SELD challenge task in DCASE2021³. The major difference of this dataset with the previous one is the introduction of localized interfering events outside of the target classes. This condition, naturally encountered in a real environment, introduces new challenges to the task. Apart from the dataset and baseline description, we present an extensive evaluation of the baseline on different versions of the dataset with and without the presence of ambient noise, directional interferers, and reverberation.

¹<https://doi.org/10.5281/zenodo.4844825>

²<https://github.com/sharathadavanne/seld-dcase2021>

³<http://dcase.community/challenge2021/task-sound-event-localization-and-detection>

2. DATASET

Similarly to the dataset of the previous iteration, the current one consists of 800 one-minute spatial recordings, of which 600 constitute the development set of the dataset, and the other 200 the evaluation set. The recordings are sampled at 24kHz, and they are offered in two 4-channel spatial audio formats, the raw signals of a tetrahedral microphone array and first-order Ambisonics, abbreviated as MIC and FOA for the rest of the paper. Detailed descriptions of the formats in terms of their directional encoding properties can be found in the previous challenge dataset report [14].

2.1. Sound events

The sound event samples are sourced from the *NIGENS general sound events database* [23], which consists of 14 classes of specific sound types, and an additional general one with disparate sounds not belonging to any of the other classes. We use the sounds in the 12 classes *alarm*, *crying baby*, *crash*, *barking dog*, *female scream*, *female speech*, *footsteps*, *knocking on door*, *male scream*, *male speech*, *ringing phone*, *piano* as target events, and the sounds in the classes *running engine*, *burning fire* and the general class as directional interferers. This division results in about 500 distinct sound samples distributed across the target events of the dataset, and about 400 across the interfering events.

2.2. Dataset synthesis

The synthesis of the spatial sound recordings are based on a collection of SRIRs acquired continuously along measurement trajectories inside 13 enclosures of Tampere University. The RIR collection and synthesis process is described in more detail in [14]. We summarize briefly the acoustical properties of the dataset. SRIRs are extracted along the measurement trajectories with an approximate resolution of 1 degree, resulting on about 1184 to 6480 possible RIRs/DOAs per room, depending on the type (circular/linear) and number of measurement trajectories. Events added in a single recording can be static or moving. The source position for a static event is drawn randomly from the pool of SRIRs of a single room used in that recording, while moving events are synthesized for one of the measured trajectories in the room. Moving events are synthesized to have an approximate speed of 10°/sec, 20°/sec, or 40°/sec, drawn randomly. The dataset is split into 8 folds with distinct rooms and samples in each of them. Distinct rooms result in different reverberation conditions, and even though similar ranges of DOAs may occur between rooms, the source distance, DRR, and reverberation conditions are distinct between folds for a certain DOA.

The events are laid out in layers in each recording, with the total number of layers determining the maximum polyphony possible. The parameter determining the density of events per layer and, hence, the average per-frame polyphony is the total gap time distributed between events in each layer. A larger gap time results in fewer events per layer and a lower average polyphony, while a smaller gap time results in higher event density and average polyphony. The last event per layer is truncated to fit the total 1 minute duration. For the present dataset there are three layers of target events and an additional layer of interfering events, resulting in a total maximum polyphony of 4. In addition to the spatialized reverberant events, multichannel ambient noise that was collected in each room with the same recording setup as the SRIRs is truncated to 1 minute segments and added to the event mixtures. The noise is scaled to result in signal-to-noise ratios (SNRs) drawn uniformly

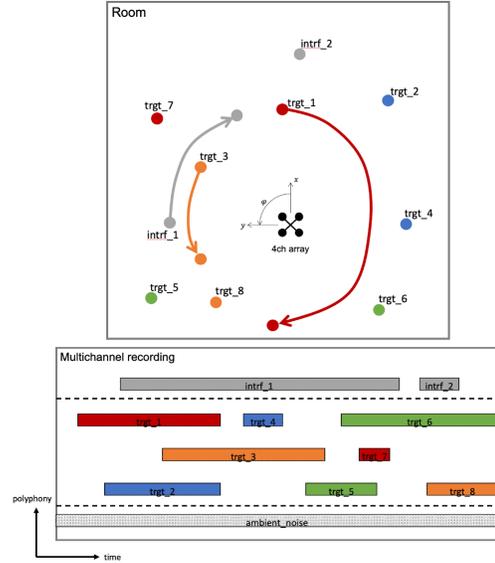


Figure 1: A graphic depiction of an emulated recording, with colored objects indicating target classes, gray objects indicating interferers and ambient noise, and arrows indicating moving events.

from noiseless (30dB) to noisy (6dB) conditions, with respect to the total energy of the target events excluding silences. An example of the layering of events in one recording is shown in Fig. 1.

2.3. Differences with DCASE2020 task 3 dataset

Even though the acoustical and synthesis characteristics of the new dataset are similar to the dataset of the previous DCASE2020 Challenge, the following differences make it more challenging:

1. Directional interferers, out of the target classes of a detection system, are common in real conditions and they add to the challenge by forcing a strong joint modeling and training strategy that can learn to ignore them.
2. The overall maximum polyphony is increased from 2 to 3 target events.
3. The recordings are not anymore divided into recordings with no overlap (polyphony 1), and recordings with two simultaneous events (polyphony 2). Instead all recordings have the maximum level of polyphony, with all intermediate levels (from silence to 3 simultaneous target events + interference) varying during the duration of the recording. This choice reflects more natural recording conditions in a real dataset.
4. Even though the dataset of DCASE2020 had instances of the same class occurring at the same time, such occurrences were fairly rare. In the present dataset, these occurrences have been increased in order to give a clear advantage to systems that can resolve this difficult but realistic case.

3. BASELINE

Similar to the previous iterations of the challenge, we adopt a modified version of SELDnet [10] as the baseline method, due to its conceptual simplicity. Its architecture remains a convolutional recur-

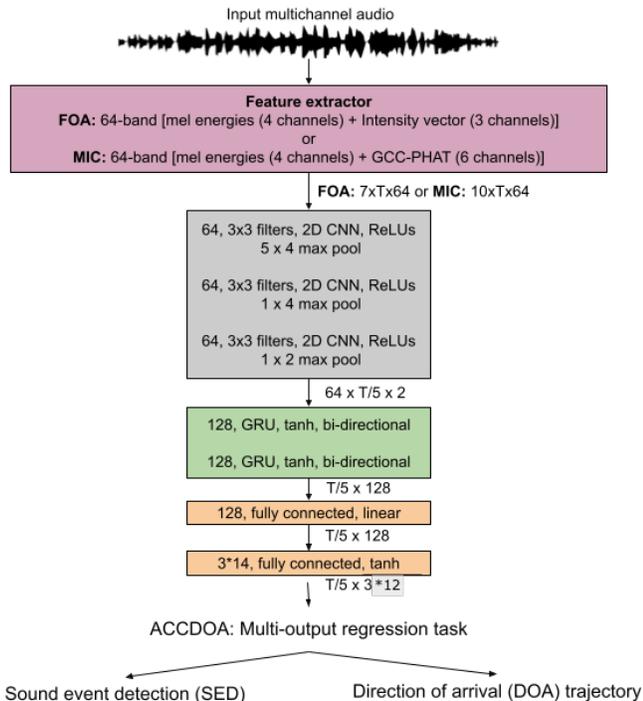


Figure 2: Convolutional recurrent neural network with ACCDOA loss for SELD.

rent neural network (CRNN) receiving multichannel log-mel spectrograms as inputs, together with acoustic intensity vectors [24] for the FOA dataset, and generalized cross-correlation (GCC-PHAT) sequences for the MIC dataset, added as extra channels. The baseline implementation extracts log-mel spectrograms in 64 mel-bands from 1024-point FFTs, using a 40 ms window and 20 ms hop length at 24kHz. The intensity vectors are similarly extracted for every FFT bin and aggregated in the same number of mel-bands as the spectrograms, while the GCC sequences are also truncated to the same number of lag values as the mel-bands, adopted from [11]. More details on the architecture and features can be found in [14].

The only difference of the current SELDnet baseline with respect to the previous DCASE challenge iteration is the output format and the respective loss function. The original SELDnet architecture employs separate output branches for detection and localization, with as many classification outputs and as many localization regressors as the number of classes. In the current baseline, we adopt the *activity-coupled cartesian direction of arrival* representation (ACCDOA) introduced in DCASE2020 Challenge by Shimada et al. [15], which unifies the SED and SSL losses into a single homogeneous regression loss, simplifying the overall architecture while simultaneously improving its performance. Using the ACCDOA representation, the network receives a sequence of T STFT frames of multichannel features and outputs $T/5 \times 3$ Cartesian vector coordinates for each of the target classes, with the direction of each vector indicating DOA and the vector length indicating class activity probability. A value of 0.5 on the length is used as the class activity threshold. The reduction in temporal resolution is intended to match the 100 ms resolution of annotations in the challenge. A block diagram of the current baseline is shown in Fig. 2.

4. EVALUATION

The dataset and baseline are delivered to the challenge participants at the commencement of the challenge, along with the development set of the dataset consisting of the 6 first folds, while the last two folds are made available during the evaluation phase of the challenge. Participants are required to report results on the test set of the development set using the predefined split of Table 1, so that conclusions can be drawn among the submissions on the same configuration. The evaluation split on Table 1, however, applies only to the evaluation results of the baseline presented herein, since during the evaluation phase participants have to report results on the testing folds (7–8) using the development dataset (folds 1–6) for training and validation in any way they see fit.

Table 1: Evaluation setup

Dataset	Splits		
	Training	Validation	Testing
Development	1,2,3,4	5	6
Evaluation	2, 3, 4, 5, 6	1	7, 8

The submissions are evaluated using the same combination of joint detection/localization metrics studied in [25, 13] and introduced in DCASE2020. Closer to SED evaluation, the localization-dependent error rate (ER_X) and F1-score (F_X) express detection performance but they penalize correct detections that occur further from the reference than some threshold distance X . On the other hand, the class-dependent localization error (LE_{CD}) and localization recall (LR_{CD}) are inspired by classical localization metrics, but are computed for each class individually before being averaged. The LE_{CD} is a mean angular localization error after pairing the predicted DOAs to their closest reference DOAs, while LR_{CD} is a simple recall metric on the detected localized events without any spatial threshold. Since in the SELD case there can be multiple simultaneous references of the same class, the detection metrics are modified to consider multiple instances of the same class and penalize cases where, e.g., only one of the predictions belong to that class. For the exact formulation of the metrics the reader is referred to [13]. The submissions are first ranked for each of the four metrics individually, and the final rank of each system is determined by the sum of the four individual ranks.

5. RESULTS

In order to evaluate the performance of the new baseline utilizing the ACCDOA loss, we compare it against the previous SELDnet baseline of DCASE2020, on the development set of DCASE2020 and the current one. Table 2 shows a clear improvement of the ACCDOA version in all metrics. Especially in the more challenging new dataset, the ACCDOA loss brings large gains in detection and improves localization accuracy by about 25%. A significant decrease of performance for both methods is also observed from the DCASE2020 dataset to the DCASE2021 dataset. This suggests that the new dataset is more challenging, as intended.

To get a more detailed picture on the effect of the various components in the scene, namely reverberation, ambient noise, and directional interferers, we generate various versions of the dataset including those components in various combinations. More specifically, the *targets*, *targets+ambience*, *targets+interferers*, *targets+ambience+interferers* develop from the presence of targets

Table 2: Comparison between the DCASE2020 baseline (2020-multi) and the current one (2021-accdoa) on the development set of DCASE2020 and the current development set.

	FOA				MIC			
	$ER_{20^\circ} \downarrow$	$F_{20^\circ} \uparrow$	$LE_{CD} \downarrow$	$LR_{CD} \uparrow$	$ER_{20^\circ} \downarrow$	$F_{20^\circ} \uparrow$	$LE_{CD} \downarrow$	$LR_{CD} \uparrow$
DCASE2020 development set								
2020-multi	0.70	44.4%	24.3°	61.9%	0.71	40.4%	25.4°	55.4%
2021-accdoa	0.60	51.9%	17.9°	59.8%	0.61	48.5%	19.3°	55.2%
DCASE2021 development set								
2020-multi	0.77	24.7%	32.1°	44.8%	0.81	19.1%	41.6°	47.4%
2021-accdoa	0.73	30.7%	24.5°	40.5%	0.75	23.4%	30.6°	37.8%

Table 3: Performance of the DCASE2021 baseline for different versions of the dataset with increasingly adverse conditions. The highlighted row corresponds to the version of the dataset used in the challenge.

	Development set								Evaluation set							
	FOA				MIC				FOA				MIC			
	$ER_{20^\circ} \downarrow$	$F_{20^\circ} \uparrow$	$LE_{CD} \downarrow$	$LR_{CD} \uparrow$	$ER_{20^\circ} \downarrow$	$F_{20^\circ} \uparrow$	$LE_{CD} \downarrow$	$LR_{CD} \uparrow$	$ER_{20^\circ} \downarrow$	$F_{20^\circ} \uparrow$	$LE_{CD} \downarrow$	$LR_{CD} \uparrow$	$ER_{20^\circ} \downarrow$	$F_{20^\circ} \uparrow$	$LE_{CD} \downarrow$	$LR_{CD} \uparrow$
Non-reverberant results																
targets	0.49	62.0%	16.3°	65.7%	0.54	55.4%	20.8°	63.7%	0.45	64.7%	15.8°	69.2%	0.50	58.7%	19.0°	67.0%
targets+ambience	0.49	61.2%	16.4°	65.6%	0.57	51.2%	20.8°	58.9%	0.47	62.7%	16.4°	67.6%	0.49	59.6%	18.6°	66.1%
targets+interferers	0.69	36.9%	24.1°	45.2%	0.72	27.7%	30.5°	42.2%	0.61	45.8%	21.2°	53.2%	0.69	31.5%	27.8°	44.9%
targets+ambience+interferers	0.66	40.3%	22.7°	46.9%	0.73	26.7%	30.4°	42.5%	0.63	44.5%	22.0°	52.8%	0.70	32.1%	28.2°	46.6%
Reverberant results																
targets	0.55	53.7%	19.9°	61.3%	0.59	47.0%	22.0°	57.3%	0.52	56.4%	19.6°	64.7%	0.53	54.6%	20.8°	61.2%
targets+ambience	0.57	50.3%	20.2°	59.3%	0.62	44.2%	22.8°	53.6%	0.52	56.4%	19.0°	62.7%	0.56	51.6%	21.0°	58.5%
targets+interferers	0.71	32.7%	26.7°	44.2%	0.76	24.0%	32.6°	39.4%	0.69	34.8%	24.7°	43.7%	0.75	24.8%	32.7°	40.4%
targets+ambience+interferers	0.73	30.7%	24.5°	40.5%	0.75	23.4%	30.6°	37.8%	0.67	37.2%	23.9°	45.8%	0.73	27.1%	30.8°	40.6%

only, to the inclusion of ambient noise or interferers separately, to the full dataset combining all components. Excluding the effect of reverberation is less straightforward due to the use of real SRIRs for the synthesis. In order to generate reverberation-free versions of the dataset, the sound events for each recording in the original dataset are spatialized with anechoic IRs of the same Eigenmike spherical microphone array used to capture the SRIRs. The anechoic array IRs are computed for the same measurement trajectories and DOAs as the measured SRIRs in each room, and stored in a similar data structure. Additionally, each IR is delayed and scaled according to the source distance of the respective measured SRIR, following an inverse distance law and a speed of sound of $c = 343m/sec$. Delaying and scaling ensures that the events between the reverberant and non-reverberant versions are approximately time-aligned and with comparable distance-dependent attenuation. The Eigenmike responses were measured in an equirectangular grid of 5° azimuth and 5° elevation in the large anechoic chamber of Aalto University, as described in [26]. Since the DOAs in the reverberant dataset do not necessarily coincide with the measurement grid of the array, array response interpolation is performed to recover anechoic IRs at the DOAs of the measured SRIRs, based on a spherical harmonic expansion of the array steering vectors, as in [26].

The results are presented in Table 3. As expected, reverberation affects negatively all combinations, increasing error rates and decreasing F-scores and localization recall in a consistent manner between the same scenarios. Additionally, it decreases localization accuracy by 2° – 4° . Inclusion of the ambient noise has a small but noticeable effect when added to the targets, without interferers. The small effect may be due to the large range of possible positive SNRs (6–30dB) distributed uniformly across the recordings. Interestingly, together with directional interference, inclusion of ambient noise seems to improve certain results slightly. This may be due to potential regularization effects of noise and is worth further investigation.

The most detrimental effects happen with the inclusion of the

directional interferers, proving that this challenging case will need to be taken into account for future SELD systems. Error rate ER increases up to about 40% in the non reverberant case for the FOA recordings, and up to about 33% for the MIC recordings. Similarly, in reverberant scenarios, the ER increases up to about 28% for both FOA and MIC formats. F-scores decrease by up to 40% on the FOA dataset and up to 50% on the MIC dataset, for both anechoic and reverberant conditions. The localization recall (LR) also drops by about 30% for both formats and both anechoic and reverberant conditions. Finally, localization errors increase by up to about 7° in the case of FOA recordings and up to 10° in the case of MIC recordings, for both anechoic and reverberant conditions. In general, the MIC dataset exhibits a worse performance than FOA in all cases. This fact may be attributed to the input features employed in the baseline for each format. GCC sequences for the MIC format may become very noisy in complex scenes with multiple simultaneous events, while the intensity vectors of the FOA format can potentially retain robustness due to their narrowband nature and sparsity of the event signals in the time-frequency domain.

6. CONCLUSIONS

In this report we describe the new dataset and baseline for the SELD task of the DCASE2021 challenge. The differences with the dataset of DCASE2020 challenge are highlighted; namely, inclusion of directional interferers, higher polyphony, and more frequent simultaneous same-class event occurrences. The evaluation task setup is also described, including a predefined fixed split on the development data for straightforward comparison of submissions. The new baseline using the ACCDOA representation shows improved performance compared to the previous one. A detailed analysis of the baseline on various versions of the dataset shows that between reverberation, ambient noise, and directional interferers, the latter has the most detrimental effect in all evaluation metrics.

7. REFERENCES

- [1] C. Evers, H. W. Löllmann, H. Mellmann, A. Schmidt, H. Barfuss, P. A. Naylor, and W. Kellermann, “The LOCATA challenge: Acoustic source localization and tracking,” *IEEE/ACM Trans. Audio, Speech, and Language Proc.*, vol. 28, pp. 1620–1643, 2020.
- [2] J.-M. Valin, F. Michaud, J. Rouat, and D. Létourneau, “Robust sound source localization using a microphone array on a mobile robot,” in *IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, vol. 2, 2003, pp. 1228–1233.
- [3] W. He, P. Motlicek, and J.-M. Odobez, “Joint localization and classification of multiple sound sources using a multi-task neural network,” in *Interspeech*, 2018, pp. 312–316.
- [4] G. Valenzise, L. Gerosa, M. Tagliasacchi, F. Antonacci, and A. Sarti, “Scream and gunshot detection and localization for audio-surveillance systems,” in *IEEE Conf. on Advanced Video and Signal Based Surveillance*, 2007, pp. 21–26.
- [5] H. M. Do, M. Pham, W. Sheng, D. Yang, and M. Liu, “RiSH: A robot-integrated smart home for elderly care,” *Robotics and Autonomous Systems*, vol. 101, pp. 74–92, 2018.
- [6] O. Brdiczka, J. L. Crowley, and P. Reignier, “Learning situation models in a smart home,” *IEEE Trans. Systems, Man, and Cybernetics*, vol. 39, no. 1, pp. 56–63, 2008.
- [7] J. Cech, R. Mittal, A. Deleforge, J. Sanchez-Riera, X. Alameda-Pineda, and R. Horaud, “Active-speaker detection and localization with microphones and cameras embedded into a robotic head,” in *IEEE/RAS Int. Conf. on Humanoid Robots*, 2013, pp. 203–210.
- [8] W. He, P. Motlicek, and J.-M. Odobez, “Deep neural networks for multiple speaker detection and localization,” in *IEEE Int. Conf. on Robotics and Automation (ICRA)*, 2018, pp. 74–79.
- [9] Y. G. Matsinos, A. D. Mazaris, K. D. Papadimitriou, A. Mniestriz, G. Hatzigiannidis, D. Maioglou, and J. D. Pantis, “Spatio-temporal variability in human and natural sounds in a rural landscape,” *Landscape ecology*, vol. 23, no. 8, pp. 945–959, 2008.
- [10] S. Adavanne, A. Politis, J. Nikunen, and T. Virtanen, “Sound event localization and detection of overlapping sources using convolutional recurrent neural networks,” *IEEE J. Selected Topics in Sig. Proc.*, vol. 13, no. 1, pp. 34–48, 2018.
- [11] Y. Cao, Q. Kong, T. Iqbal, F. An, W. Wang, and M. Plumbley, “Polyphonic sound event detection and localization using a two-stage strategy,” in *Work. on Detection and Classification of Acoustic Scenes and Events (DCASE)*, 2019.
- [12] T. N. T. Nguyen, D. L. Jones, and W.-S. Gan, “A sequence matching network for polyphonic sound event localization and detection,” in *IEEE Int. Conf. on Acoustics, Speech and Sig. Proc. (ICASSP)*, 2020, pp. 71–75.
- [13] A. Politis, A. Mesaros, S. Adavanne, T. Heittola, and T. Virtanen, “Overview and evaluation of sound event localization and detection in DCASE 2019,” *IEEE/ACM Trans. Audio, Speech, and Language Proc.*, 2020.
- [14] A. Politis, S. Adavanne, and T. Virtanen, “A dataset of reverberant spatial sound scenes with moving sources for sound event localization and detection,” in *Work. on Detection and Classification of Acoustic Scenes and Events (DCASE)*, 2020, pp. 165–169.
- [15] 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 *IEEE Int. Conf. on Acoustics, Speech and Sig. Proc. (ICASSP)*, 2021, pp. 915–919.
- [16] H. Phan, L. Pham, P. Koch, N. Q. K. Duong, I. McLoughlin, and A. Mertins, “On multitask loss function for audio event detection and localization,” in *Work. on Detection and Classification of Acoustic Scenes and Events (DCASE)*, 2020, pp. 160–164.
- [17] Q. Wang, J. Du, H.-X. Wu, J. Pan, F. Ma, and C.-H. Lee, “A four-stage data augmentation approach to resnet-conformer based acoustic modeling for sound event localization and detection,” *arXiv preprint arXiv:2101.02919*, 2021.
- [18] T. N. T. Nguyen, D. L. Jones, and W. S. Gan, “Ensemble of sequence matching networks for dynamic sound event localization, detection, and tracking,” in *Work. on Detection and Classification of Acoustic Scenes and Events (DCASE)*, 2020, pp. 120–124.
- [19] A. Pérez-López and R. Ibáñez-Usach, “Papafil: A low complexity sound event localization and detection method with parametric particle filtering and gradient boosting,” in *Work. on Detection and Classification of Acoustic Scenes and Events (DCASE)*, 2020, pp. 155–159.
- [20] T. N. T. Nguyen, N. K. Nguyen, H. Phan, L. Pham, K. Ooi, D. L. Jones, and W.-S. Gan, “A general network architecture for sound event localization and detection using transfer learning and recurrent neural network,” in *IEEE Int. Conf. on Acoustics, Speech and Sig. Proc. (ICASSP)*, 2021, pp. 935–939.
- [21] Y. Cao, T. Iqbal, Q. Kong, Y. Zhong, W. Wang, and M. D. Plumbley, “Event-independent network for polyphonic sound event localization and detection,” in *Work. on Detection and Classification of Acoustic Scenes and Events (DCASE)*, 2020, pp. 11–15.
- [22] Y. Cao, T. Iqbal, Q. Kong, F. An, W. Wang, and M. D. Plumbley, “An improved event-independent network for polyphonic sound event localization and detection,” in *IEEE Int. Conf. on Acoustics, Speech and Sig. Proc. (ICASSP)*, 2021, pp. 885–889.
- [23] I. Trowitzsch, J. Taghia, Y. Kashef, and K. Obermayer, “The NIGENS general sound events database,” *arXiv preprint arXiv:1902.08314*, 2019.
- [24] V. Pulkki, A. Politis, M.-V. Laitinen, J. Vilkkamo, and J. Ahonen, “First-order directional audio coding (DirAC),” in *Parametric Time-Frequency Domain Spatial Audio*, V. Pulkki, S. Delikaris-Manias, and A. Politis, Eds. John Wiley & Sons, 2017, pp. 89–138.
- [25] A. Mesaros, S. Adavanne, A. Politis, T. Heittola, and T. Virtanen, “Joint measurement of localization and detection of sound events,” in *IEEE Work. on Applications of Sig. Proc. to Audio and Acoustics (WASPAA)*, 2019, pp. 333–337.
- [26] S. Tervo and A. Politis, “Direction of arrival estimation of reflections from room impulse responses using a spherical microphone array,” *IEEE/ACM Trans. Audio, Speech, and Language Proc.*, vol. 23, no. 10, pp. 1539–1551, 2015.