



Spatial LibriSpeech: An Augmented Dataset for Spatial Audio Learning

Miguel Sarabia, Elena Menyaylenko, Alessandro Toso, Skyler Seto, Zakaria Aldeneh, Shadi Pirhosseinloo, Luca Zappella, Barry-John Theobald, Nicholas Apostoloff, Jonathan Sheaffer

Apple

miguelsdc@apple.com

Abstract

We present Spatial LibriSpeech, a spatial audio dataset with over 650 hours of 19-channel audio, first-order ambisonics, and optional distractor noise. Spatial LibriSpeech is designed for machine learning model training, and it includes labels for source position, speaking direction, room acoustics and geometry. Spatial LibriSpeech is generated by augmenting LibriSpeech samples with 200k+ simulated acoustic conditions across 8k+ synthetic rooms. To demonstrate the utility of our dataset, we train models on four spatial audio tasks, resulting in a median absolute error of 6.60° on 3D source localization, 0.43m on distance, 90.66ms on T30, and 2.74dB on direct-to-reverberant ratio estimation. We show that the same models generalize well to widely-used evaluation datasets, e.g., obtaining a median absolute error of 12.43° on 3D source localization on TUT Sound Events 2018, and 157.32ms on T30 estimation on ACE Challenge.

1. Introduction

Humans can infer a great amount of information from their acoustic environment, including determining the spatial location of sound sources, estimating the size of a room, and estimating the room reverberance [1, 2]. To develop machine learning algorithms with the same level of acoustic spatial awareness, we require a large and diverse dataset with annotations for each task. To the best of our knowledge, no publicly-available dataset has the scale, and diversity needed to train general models for multiple spatial audio tasks. This lack of large-scale datasets limits the applicability of modern machine learning techniques to spatial audio.

In this paper, we introduce Spatial LibriSpeech¹, a spatially augmented version of LibriSpeech [3] with optional noise from the Microsoft Deep Noise Suppression Challenge 2021 [4]. Spatial LibriSpeech augments the LibriSpeech and Microsoft Deep Noise Suppression samples by simulating how they would be perceived by a microphone array in various synthetic rooms. Spatial LibriSpeech contains over 650 hours of spatial audio with labels for source position, speaking direction, room acoustics, and room geometry (refer to Section 2 for a review of spatial audio tasks for which these labels are useful). Our goal is for Spatial LibriSpeech to be the main training dataset for spatial audio applications.

To create Spatial LibriSpeech, we first generated 8,952 synthetic rooms, which were used to obtain room impulse responses (RIRs) that were convolved and scaled with LibriSpeech samples. The RIRs are modeled on the Zylia microphone array, a 19-channel spherical microphone array. This ar-

ray facilitates a means for extracting a third-order ambisonics representation up to a frequency of 3080Hz [5]. For completeness, we also provide synthesized, full-bandwidth first-order ambisonics [6], which are aliasing free. Either of these formats may be used to simulate a wide variety of arrays [7]. We describe the dataset generation process in more detail in Section 3.1. A key advantage of convolving existing speech samples with synthetic RIRs for spatial audio dataset generation is that we can ensure our dataset spans a variety of acoustic room properties and room sizes. For instance, full-band T30 values range from 145ms to 2846ms, and room floor area ranges from 13.3m^2 to 277.4m^2 . Section 3.2 describes diversity statistics of Spatial LibriSpeech.

While Spatial LibriSpeech has the potential to be used to train models for multiple tasks, in this paper we focus on some of the most fundamental spatial audio detection tasks: i) 3D source localization, ii) source distance, iii) third-octave narrow-band direct-to-reverberant ratios (DRRs), and iv) third-octave narrow-band T30s. All models share the same architecture and training regime (presented in Section 4). When training spatial audio models, there is a choice of using the microphone array signals as inputs or converting the microphone array signals to ambisonics to obtain device-agnostic models. In our work, models consume first-order ambisonics.

Our models achieve a median absolute error of 6.60° in 3D source localization, 0.43m in distance estimation, 2.74dB in DRR estimation, and 90.66ms in T30 estimation. Since our training dataset is composed of simulated acoustics, we verify the transferability of our models by testing them on two evaluation datasets: TUT Sound Events 2018 [8], and ACE Challenge [9] (Section 4.1). We find that fine-tuning is beneficial for ACE Challenge but unnecessary for TUT Sound Events 2018, and we plot the responses of our models on both evaluation datasets showing that ACE responses are at the tail of the Spatial LibriSpeech responses (see Section 4.2). Lastly, we show that, for single-task models, using a version of Spatial LibriSpeech made by uniformly sampling 10% of the dataset yields the same performance as the full dataset (cf. Section 4.3).

2. Background

We start by reviewing *spatial audio* tasks to highlight the large number of potential use cases for Spatial LibriSpeech. We divide audio tasks into three categories: *source parameter estimation*, *environment parameter estimation*, and *spatial processing*. Source parameter estimation tasks are concerned with understanding audio sources. Examples include estimation of: source localization [25], source distance [26], and speaking direction [27]. Environment parameter estimation tasks involve understanding the environment where the audio is produced.

¹See github.com/apple/ml-spatial-librispeech.

Table 1: Comparison of existing datasets. ROOMS refers to the number of unique acoustic environments in the dataset, ROOM CONFIG. to the number of physical configurations per environment (e.g. different source receiver positions), CHANNELS to the availability of spatial information from a microphone array or other encoding: N - N channels mic array, B - Binaural, X^{th} OA - X^{th} order ambisonics, DATA-TYPES: $s\text{IR}$ - Simulated IR, $r\text{IR}$ - Recorded IR, $s\text{A}$ - Simulated Audio, $r\text{A}$ - Recorded Audio. LABELS: P - Position, SD - Speaking direction, R - Room acoustics, G - Room geometry, O - Other.

DATASET	ROOMS	ROOM CONFIG.	CHANNELS	DATA-TYPES	LABELS
dEchorate [10]	11	180	5	$s\text{IR}$, $r\text{IR}$	P
Arni [11]	1	21	3^{rd} OA, 4^{th} OA	$r\text{IR}$	P , R
GIR [12]	1	2,951	1	$r\text{IR}$	P , O
EasyCom [13]	1	50	4, B	$r\text{A}$	P , R , G , O
CoupledRooms [14]	2	101	4^{th} OA	$r\text{IR}$	P , R , G
DCASE2021 task 3 [15]	13	1,184–6,480	4, 1^{st} OA	$r\text{IR}$, $r\text{A}$	P
BUT ReverbDB [16]	8	155	1	$r\text{IR}$	P , R , G
SBSBRIR [17]	1	$180 \times 2^\circ$	B	$r\text{IR}$	P
BIRD [18]	12,500	8	2	$s\text{IR}$	P , R
Kemar BRIRs [19]	1	50	B	$r\text{IR}$	P , G
Motus [20]	1	3,320	32, 4^{th} OA	$r\text{IR}$	P , R , G
Aachen IR database [21]	4	17	B	$r\text{IR}$	P , R , G , O
ACE Challenge [9]	7	10	2 - 32	$r\text{IR}$	P , R
DIHRA [22]	2	62	3	$r\text{A}$, $s\text{A}$	P , G
Voice-Home [23]	12	24	8	$r\text{IR}$, $r\text{A}$	P , G
Sweet-Home [24]	4	7	1	$r\text{A}$	-
Microsoft DNS 2001 [4]	14,576	1	1	$s\text{IR}$, $r\text{IR}$	R
TUT Sound Events 2018 [8]	5	487-4,366	8, 1^{st} OA	$s\text{A}$, $r\text{A}$	P , G
Spatial LibriSpeech	8,952	20	19, 1^{st} OA	$s\text{A}$	P , SD , R , G

Example tasks include estimation of DRR [28], material absorption and scattering [29], and room volume [30]. Finally, spatial processing tasks involve the transformation of audio signals with information extracted from the acoustic environment. Example tasks include de-reverberation [31], beamforming [32], and audio source separation [33].

Table 1 contains a summary of the main differences between 18 published spatial audio datasets in terms of the number of environments, number of physical configurations per environment, recorded/simulated channels, and data-types and labels included in the dataset. We found only two datasets that include over 50k unique configurations across all environments: BIRD [18] and DCASE2021 Task 3 [15]. However, BIRD features only two microphone channels, and DCASE2021 Task 3 includes only position labels. Spatial LibriSpeech is the only dataset to feature over 200k unique configurations, labels for a large number of acoustic tasks, and both mic-domain audio and ambisonics. This lack of diversity and labels limits the applicability of existing datasets to modern machine learning techniques [34] such as multi-task learning or contrastive representation learning [35].

3. The Spatial LibriSpeech Dataset

In this section, we describe the generation of Spatial LibriSpeech and its defining characteristics.

3.1. Generation Methodology

Our pipeline to generate Spatial LibriSpeech consists of three steps: *parametric room generation*, *room impulse response simulations*, and *mixing*.

Parametric room generation. We start by defining a realistic set of conditions for commonly encountered living spaces following [36]. These consist of a predefined distribution of reverberation times and room shapes and sizes. The materials associated with the room surfaces are chosen from a database of typical construction materials, for which the absorption and scattering coefficients as a function of frequency are available.

Room impulse response simulations. We use a geometrical acoustic solver, which includes a high fidelity model of the

room that accounts for the directivity of the microphones (including diffraction effects from the array body). The room is populated with several acoustic objects including the recording device placed at a randomized but bounded position, several sources surrounding the device with a randomized looking direction and directivity function. When sources are intended to represent speech, their directivity is computed using a boundary element code using a set of artificially generated models. The direct path from each source to the microphone array, the early reflections, and the late reverberations are used to assemble each impulse response.

Mixing. Once we have computed RIRs for each configuration of the simulated rooms, we mix them with LibriSpeech samples by randomly assigning them to sources in every simulated room. We then remove leading and trailing silence from each sample², convolve the sample with the selected RIR and scale the output to a random active speech level (ASL) between 85dB-ASL and 100dB-ASL at mouth reference point³. Finally, we save the resulting audio samples and store the relevant labels from both the room simulation and LibriSpeech. For the distractor noise we follow the same process, except we assign a random source, different to the main signal source, and set the signal strength to a random signal-to-noise ratio (SNR) between 10dB and 40dB.

3.2. Dataset Characteristics

The main characteristics of Spatial LibriSpeech are summarized in Table 2. The dataset consists of two splits: TRAIN which is derived from LibriSpeech’s `train-clean-100` and `train-clean-360` subset, and TEST which is derived from `test-clean`. The audio is sampled at 16kHz.

Spatial LibriSpeech includes the following labels: source localization (azimuth and elevation), speaking direction, room volume, surface and floor area, voice directivity identifier, narrow-band C50, DRR, EDT, T20, and T30 [38]; and all

²Silence is defined as any leading or trailing sound before a 100ms segment at more than 1% volume.

³Mouth reference point is the point on the reference axis 25 mm in front of the lip plane ITU-T P.581 [37]

Table 2: Main characteristics of Spatial LibriSpeech.

	TRAIN	TEST
Total duration	573h 13m 12s	85h 29m 20s
Speech samples	171,951	49,505
Simulated rooms	8,952	4,970
Source azimuth	[-180.0°, +180.0°]	[-180.0°, +180.0°]
Source elevation	[-48.1°, +48.7°]	[-49.5°, +42.4°]
Source distance	[0.5m, 4.0m]	[0.9m, 4.0m]
Speaking azimuth	[-180.0°, +180.0°]	[-180.0°, +180.0°]
Speaking elevation	[-89.3° +87.7°]	[-74.5° +68.2°]
Voice directivities	16 variations	1 variation
Room floor area	[13.3m ² , 277.4m ²]	[14.3m ² , 277.4m ²]
Full-band T30	[144.5ms, 2846.0ms]	[157.8ms, 1267.7ms]

original metadata from LibriSpeech, including ids for speaker, books, and chapters, and the reference transcriptions.

We verify that our models do not overfit to the virtual rooms in the training set by checking that the performance of models trained with Spatial LibriSpeech transfer to two existing evaluation datasets (see Section 4.1).

4. Training with Spatial LibriSpeech

To demonstrate the utility of Spatial LibriSpeech in real-world situations, we trained several neural networks for: i) 3D source localization, ii) source distance, iii) DRR, and iv) T30. We chose these tasks since they are representative of the source and environment parameter estimation spatial audio tasks and because public third party evaluation datasets were available to analyse transferability.

For 3D source localization, we report the median absolute error of the 3D angle (cf. [39]):

$$\alpha = \cos^{-1} \left(\sin(\phi) \sin(\hat{\phi}) + \cos(\phi) \cos(\hat{\phi}) \cos(\theta - \hat{\theta}) \right), \quad (1)$$

where ϕ and θ are the ground-truth azimuth and elevation respectively, and $\hat{\phi}$ and $\hat{\theta}$ are the predicted azimuth and elevation. For distance regression, we report the median absolute error in meters. As DRR and T30 are 20-dimensional vectors representing the third-octave frequencies between 100Hz and 8kHz, we report both the median absolute error as well as the Pearson correlation across frequencies.

Our objective is to show that models trained with Spatial LibriSpeech transfer to two established evaluation datasets: ACE Challenge [9] and TUT Sound Events 2018 [8]. All models share the same architecture. We first transform and normalize segments⁴ of spatial audio encoded as 4-channel first-order ambisonics into active and reactive components following [40]. The active and reactive components are fed through two independent branches of four 3D-convolutional layers each (of 2, 4, 8 and 16 channels respectively), with max-pool, batch-norm and exponential linear units between the convolutional layers. Next, the output of both branches is flattened and concatenated, and fed to a 3-layer multi-layer perceptron (MLP). The number of outputs of the last layer depends on the task: two for source localization (one for azimuth, and another for elevation), one for distance regression, 20 for DRR, and 20 for T30. Our model has 101,933 parameters (for 3D source localization and distance) and 747,148 parameters (for T30 and DRR). Models are trained for 20 epochs for 3D source localization and distance, and 50 epochs for T30 and DRR⁵. We use the Adam optimizer, a weight decay of 0.01, dropout of 50% on the output

⁴The duration of the segments varies by task, for 3D source localization and distance we use 0.5s, for T30 and DRR we use 4.0s.

⁵This represents 80.8M gradient updates for 3D source localization and distance, and 21.2M gradient updates for T30 and DRR estimation.

Table 3: Performance of models trained with the Spatial LibriSpeech on 3D Sound Localization, Distance, T30 and DRR regression. For each task, 5 models are trained with the train set of Spatial LibriSpeech and evaluated against the test set of Spatial LibriSpeech, TUT Sound Events 2018 [8], and ACE Challenge [9]. We report the median absolute error ($\bar{\epsilon}_*$) and IQR across predictions (Δ_*) for the best performing model on the test set of Spatial LibriSpeech. We also report the IQR across the median absolute error of each model (Δ_m). For DRR and T30, $\bar{\epsilon}_*$, Δ_* , and Δ_m are taken over all frequency bins. Additionally, for DRR and T30, we report the median Pearson correlation of the best performing model across frequency bin for each prediction ($\bar{\rho}_*$). Purple violin plots show the density of the absolute error, while pink violin plots show the density of the Pearson correlations, both for the best performing models.

	3D SOURCE LOCAL.	DISTANCE ESTIMATION
SPATIAL LIBRISPEECH (TEST)	$\bar{\epsilon}_*: 6.60^\circ$ $\Delta_*: 7.75^\circ$ $\Delta_m: 0.52^\circ$	$\bar{\epsilon}_*: 0.43\text{m}$ $\Delta_*: 0.56\text{m}$ $\Delta_m: 0.01\text{m}$
TUT SOUNDS (ANSIM)	$\bar{\epsilon}_*: 8.19^\circ$ $\Delta_*: 9.49^\circ$ $\Delta_m: 1.82^\circ$	$\bar{\epsilon}_*: 4.76\text{m}$ $\Delta_*: 4.39\text{m}$ $\Delta_m: 0.04\text{m}$
TUT SOUNDS (RESIM)	$\bar{\epsilon}_*: 5.80^\circ$ $\Delta_*: 5.57^\circ$ $\Delta_m: 0.51^\circ$	$\bar{\epsilon}_*: 0.25\text{m}$ $\Delta_*: 0.30\text{m}$ $\Delta_m: 0.02\text{m}$
TUT SOUNDS (REAL)	$\bar{\epsilon}_*: 12.43^\circ$ $\Delta_*: 15.96^\circ$ $\Delta_m: 1.01^\circ$	$\bar{\epsilon}_*: 0.60\text{m}$ $\Delta_*: 0.85\text{m}$ $\Delta_m: 0.03\text{m}$
	DRR	T30
SPATIAL LIBRISPEECH (TEST)	$\bar{\epsilon}_*: 2.74\text{dB}$ $\Delta_*: 3.49\text{dB}$ $\bar{\rho}_*: 0.98$ $\Delta_m: 0.35\text{dB}$	$\bar{\epsilon}_*: 90.66\text{ms}$ $\Delta_*: 116.83\text{ms}$ $\bar{\rho}_*: 0.46$ $\Delta_m: 35.12\text{ms}$
ACE CHALLENGE	$\bar{\epsilon}_*: 14.50\text{dB}$ $\Delta_*: 14.75\text{dB}$ $\bar{\rho}_*: 0.74$ $\Delta_m: 1.40\text{dB}$	$\bar{\epsilon}_*: 157.32\text{ms}$ $\Delta_*: 230.13\text{ms}$ $\bar{\rho}_*: 0.32$ $\Delta_m: 70.94\text{ms}$
ACE CHALLENGE (FINETUNED)	$\bar{\epsilon}_*: 4.81\text{dB}$ $\Delta_*: 6.61\text{dB}$ $\bar{\rho}_*: 0.40$ $\Delta_m: 3.79\text{dB}$	$\bar{\epsilon}_*: 105.57\text{ms}$ $\Delta_*: 187.28\text{ms}$ $\bar{\rho}_*: 0.43$ $\Delta_m: 23.16\text{ms}$

of the convolutional block, Kaiming uniform initialization, and a learning rate of 10^{-5} . Our models completed training in a median 8h15min for 3D Source Localization and Distance Estimation, and 20h for T30 and DRR on a 80-core Intel Xeon CPU with 8 NVIDIA Tesla V100s GPUs, and 375GiB of RAM.

Table 4: *Performance contextualization for different tasks.* RANDOM refers to the performance of the model described in Section 4 without any training. S.L.S indicates the performance of our baseline (see Table 3). EXT. refers to external benchmarks. The comparison with external benchmark may not be fair, as they were trained with different architectures, training regimes, and datasets closer to the evaluation dataset. *Model was fine-tuned on ACE Dev.

TASK	Eval.	Metric	RANDOM	S.L.S.	EXT.
3D Src	ANSIM [8]	median abs.	88.49°	8.19°	~4.3° [39]
Local.	RESIM [8]	3D angle	88.55°	5.80°	~7.4° [39]
	REAL [8]	error (eq. 1)	88.53°	12.43°	~4.3° [39]
T30			-712.43ms	156.44ms*	22.1ms [28]
DRR	ACE [9]	mean error	9.84dB	4.04dB*	0.81dB [28]

4.1. Performance and transferability

Table 3 shows the performance of our models on all four tasks when evaluated against the test set of Spatial LibriSpeech, TUT Sounds Events 2018 [8], and ACE Challenge Eval set⁶ [9]. All models were trained using Spatial LibriSpeech only.

Our results show that the difference in performance between the test set of Spatial LibriSpeech and the performance on TUT Sound Events 2018 is small, +5.83° and +0.17m worse on REAL than Spatial LibriSpeech, though better on RESIM than on Spatial LibriSpeech (-0.90° and -0.18m). We find that our models do not exploit spurious correlations between the speech level and the distance labels, instead incorporating information from reverberations to make accurate predictions. This is evidenced by the median absolute error of ANSIM (an anechoic subset of TUT Sound Events 2018) being 4.76m.

Looking at the performance differences between Spatial LibriSpeech and ACE for DRR and T30 regression we find that the transfer penalty is higher (+11.76dB for DRR, and +66.66ms for T30). We also investigate fine-tuning the last three layers (the MLP) with the ACE Challenge Dev set for 20 epochs, and find that the median absolute error improves significantly for both T30 and DRR (resulting in 2.07dB and 14.91ms gaps with the Spatial LibriSpeech test set).

For additional context into these results, Table 4 compares the performance of our models against the performance of a randomly initialized version of our model as well as the reported performance of recent benchmarks [28, 39]. In general, benchmarks obtain better results, as they use more complex architectures and train on data specifically modeled after the evaluation sets. We hypothesize that the performance of our model is close to the benchmarks due to the diversity of Spatial LibriSpeech, since we did not carry out an architecture or hyper-parameter search. Furthermore, note that none of these benchmarks published their training sets. In contrast, LibriSpeech was primarily designed for training spatial audio models.

4.2. Visualization of Dataset Representations

Plotting the representations from our models trained on Spatial LibriSpeech further illustrates the transferability to real-world test data seen in Section 4.1. Figure 1 shows UMAP [42] plots of embeddings extracted from representations before the MLP block of the network⁷ We see that the model trained for 3D

⁶For ACE Challenge, we obtained first-order ambisonics from EM32 audio samples [41].

⁷UMAP embeddings are constructed by first partitioning the train and test sets into 50 batches of data, then sampling 10,000 samples from




Figure 1: *UMAP projection of different datasets from representations before the MLP block of the model.* Representations are collected from two networks trained for 3D source localization (a) and DRR (b) using 10,000 samples across all datasets.

source localization (Figure 1a) yields overlapping representations for Spatial LibriSpeech and TUT Sounds Events 2018, while the model trained on DRR (Figure 1b) yields representations on ACE at the tail of the Spatial LibriSpeech representations indicating the need for a small amount of fine-tuning of the MLP block on the target data.

4.3. Training with a smaller subset of Spatial LibriSpeech

We explore the performance of models trained with just 10% of Spatial LibriSpeech⁸, which we sample uniformly, maintaining the diversity of the dataset. We find that these models tend to perform worse on the test set of Spatial LibriSpeech than models trained with the full dataset, for instance the best median absolute error on 3D sound localization is 0.39° higher, and 1.36ms higher for T30 regression on the 10% models. However, when looking at performance on external baselines, models trained with 10% of the training data tend to perform better, for example the best median absolute error on 3D sound localization on TUT Sound Events 2018 was 0.88° lower, and the median Pearson correlation on T30 regression on ACE Challenge was 0.12 higher. Still, all performance differences were smaller than the IQR between models. Based on these results, we recommend researchers prototype with the smaller version of the dataset, which will be available as a separate download. Additionally, we release the full Spatial LibriSpeech dataset to enable researchers to explore other tasks, such as ablations of acoustic conditions or representation learning.

5. Conclusion & Further Work

We have presented Spatial LibriSpeech, a spatial audio dataset for multiple spatial audio tasks and representation learning. We have shown the utility of our dataset with a simple convolutional network trained on Spatial LibriSpeech, the performance of which transfers to established baselines with minimal intervention, and that the results are close to the state-of-the-art, despite a less sophisticated architecture.

We intend to use Spatial LibriSpeech for a number of other tasks, such as denoising, or room identification. Another interesting line of research is whether we can further improve performance using representation learning to train a single model to regress to many of the spatial audio tasks with the same embedding. We look forward to the community using Spatial LibriSpeech to accelerate research in spatial audio.

all train/test partitions to fit UMAP. Figure 1 depicts a random subset of 2048 samples from train Spatial LibriSpeech, and 128 samples from all other evaluation sets.

⁸The number of training epochs is increased by 10x to keep the number of model updates constant.

6. References

- [1] M. Ohuchi, Y. Iwaya, Y. Suzuki, and T. Munekata, "A comparative study of sound localization acuity of congenital blind and sighted people," *Acoust. Sci. and Tech.*, vol. 27, no. 5, pp. 290–293, 2006.
- [2] S. Hameed, H. Pakarinen, K. Valde, and V. Pulkki, "Psychoacoustic cues in room size perception," *J. of the Audio Eng. Soc.*, 2004.
- [3] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, "Librispeech: An ASR corpus based on public domain audio books," in *ICASSP*. IEEE, 2015.
- [4] C. K. Reddy, H. Dubey, K. Koishida, A. Nair, V. Gopal, R. Cutler, S. Braun, H. Gamper, R. Aichner, and S. Srinivasan, "INTERSPEECH 2021 Deep Noise Suppression Challenge," in *Interspeech*, 2021.
- [5] D. Pinardi and A. Farina, "Metrics for evaluating the spatial accuracy of microphone arrays," in *IBDA*. IEEE, 2021.
- [6] M. J. Gerzon, "Periphone (with height sound reproduction)," in *Audio Eng. Soc. Convention*. AES, 1972.
- [7] H. Gamper, M. R. Thomas, L. Corbin, and I. Tashev, "Synthesis of Device-Independent Noise Corpora for Realistic ASR Evaluation," in *Interspeech*, 2016.
- [8] S. Adavanne, A. Politis, J. Nikunen, and T. Virtanen, "Sound event localization and detection of overlapping sources using convolutional recurrent neural networks," *J. of Sel. Top. in Signal Process.*, vol. 13, no. 1, 2019.
- [9] J. Eaton, N. D. Gaubitch, A. H. Moore, and P. A. Naylor, "Estimation of Room Acoustic Parameters: The ACE Challenge," *Trans. on Audio, Speech, and Lang. Process.*, vol. 24, 2016.
- [10] D. Di Carlo, P. Tandeitnik, C. Foy, N. Bertin, A. Deleforge, and S. Gannot, "dEchorate: a calibrated room impulse response dataset for echo-aware signal processing," *J. of Audio Speech Music Proc.*, vol. 39, 2021.
- [11] T. McKenzie, L. McCormack, and C. Hold, "Dataset of spatial room impulse responses in a variable acoustics room for six degrees-of-freedom rendering and analysis." arXiv, 2021.
- [12] R. Rust, A. Xydis, K. Heutschi, N. Perraudin, G. Casas, C. Du, J. Strauss, K. Eggenschwiler, F. Perez-Cruz, F. Gramazio, and M. Kohler, "A data acquisition setup for data driven acoustic design," *Build. Acoust.*, vol. 28, no. 4, pp. 345–360, 2021.
- [13] J. Donley, V. Tourbabin, J.-S. Lee, M. Broyles, H. Jiang, J. Shen, M. Pantic, V. K. Ithapu, and R. Mehra, "Easycom: An augmented reality dataset to support algorithms for easy communication in noisy environments." arXiv, 2021.
- [14] T. McKenzie, S. J. Schlecht, and V. Pulkki, "Acoustic analysis and dataset of transitions between coupled rooms," in *ICASSP*. IEEE, 2021, pp. 481–485.
- [15] A. Politis, S. Adavanne, D. Krause, A. Deleforge, P. Srivastava, and T. Virtanen, "A dataset of dynamic reverberant sound scenes with directional interferers for sound event localization and detection," in *Detect. and Classif. of Acoust. Scenes and Events Workshop*, 2021.
- [16] I. Szöke, M. Skácel, L. Mošner, J. Paliesek, and J. Černocký, "Building and evaluation of a real room impulse response dataset," *J. of Sel. Top. in Signal Process.*, vol. 13, no. 4, 2019.
- [17] D. Satongar, Y. W. Lam, and C. Pike, "Measurement and analysis of a spatially sampled binaural room impulse response dataset," in *Int. Congr. on Sound and Vib.*, 2014.
- [18] F. Grondin, J.-S. Lauzon, S. Michaud, M. Ravanelli, and F. Michaud, "BIRD: Big Impulse Response Dataset." arXiv, 2020.
- [19] C. Mittag, S. Werner, M. Böhme, and F. Klein, "Dataset of Binaural Room Impulse Responses at Multiple Recording Positions, Source Positions and Orientations in a Real Room," in *DAGA*. DEGA Akustik, 2017.
- [20] G. Götz, S. J. Schlecht, and V. Pulkki, "A dataset of higher-order ambisonic room impulse responses and 3d models measured in a room with varying furniture," in *IBDA*. IEEE, 2021.
- [21] M. Jeub, M. Schafer, and P. Vary, "A binaural room impulse response database for the evaluation of dereverberation algorithms," in *Int. Conf. on Digit. Signal Process.* IEEE, 2009.
- [22] M. Ravanelli, L. Cristoforetti, R. Gretter, M. Pellin, A. Sosi, and M. Omologo, "The DIRHA-ENGLISH corpus and related tasks for distant-speech recognition in domestic environments," in *ASRU*. IEEE, 2015.
- [23] N. Bertin, E. Camberlein, E. Vincent, R. Lebarbenchon, S. Peillon, É. Lamandé, S. Sivasankaran, F. Bimbot, I. Illina, A. Tom, S. Fleury, and E. Jamet, "A French corpus for distant-microphone speech processing in real homes," in *Interspeech*, 2016.
- [24] M. Vacher, B. Lecouteux, P. Chahua, F. Portet, B. Meillon, and N. Bonnefond, "The Sweet-Home speech and multimodal corpus for home automation interaction," in *LREC*, 2014.
- [25] P.-A. Grumiaux, S. Kitić, L. Girin, and A. Guérin, "A survey of sound source localization with deep learning methods," *J. of the Acoust. Soc. Am.*, vol. 152, no. 1, 2022.
- [26] T. Gburrek, J. Schmalenstroer, A. Brendel, W. Kellermann, and R. Haeb-Umbach, "Deep neural network based distance estimation for geometry calibration in acoustic sensor networks," in *Eusipco*, 2021.
- [27] K. Ahuja, A. Kong, M. Goel, and C. Harrison, "Direction-of-Voice (DoV) Estimation for Intuitive Speech Interaction with Smart Devices Ecosystems," in *UIST*. ACM, 2020.
- [28] N. J. Bryan, "Impulse response data augmentation and deep neural networks for blind room acoustic parameter estimation," in *ICASSP*, 2020.
- [29] S. Pelzer and M. Vorländer, "Inversion of a room acoustics model for the determination of acoustical surface properties in enclosed spaces," *Proc. of Meet. on Acoust.*, vol. 19, no. 1, 2013.
- [30] P. Srivastava, A. Deleforge, and E. Vincent, "Blind room parameter estimation using multiple multichannel speech recordings," in *Workshop on Appl. of Signal Process. to Audio and Acoust.*, 2021.
- [31] K. Kinoshita, M. Delcroix, H. Kwon, T. Mori, and T. Nakatani, "Neural network-based spectrum estimation for online wpe dereverberation." in *Interspeech*, 2017.
- [32] J. Heymann, L. Drude, and R. Haeb-Umbach, "Neural network based spectral mask estimation for acoustic beamforming," in *ICASSP*, 2016.
- [33] E. Tzinis, Z. Wang, and P. Smaragdis, "Sudo rm -rf: Efficient networks for universal audio source separation," in *Int. Workshop on Mach. Learn. for Signal Process.* IEEE, 2020.
- [34] Z. Gong, P. Zhong, and W. Hu, "Diversity in machine learning," *IEEE Access*, vol. 7, 2019.
- [35] A. Baevski, W.-N. Hsu, Q. Xu, A. Babu, J. Gu, and M. Auli, "data2vec: A general framework for self-supervised learning in speech, vision and language," in *ICML*, 2022.
- [36] C. Díaz and A. Pedrero, "The reverberation time of furnished rooms in dwellings," *Appl. Acoustics*, vol. 66, no. 8, 2005.
- [37] ITU-T P.581, "Use of head and torso simulator for hands-free and handset terminal testing," International Telecommunication Union, standard, Jul. 2022.
- [38] ISO 3382-1, "Measurement of room acoustic parameters — Part 1: Performance spaces," International Organization for Standardization, standard, Jun. 2009.
- [39] C. Schymura, B. Bönninghoff, T. Ochiai, M. Delcroix, K. Kinoshita, T. Nakatani, S. Araki, and D. Kolossa, "PILOT: Introducing Transformers for Probabilistic Sound Event Localization," in *Interspeech*, 2021.
- [40] F. Jacobsen, "Active and reactive sound intensity in a reverberant sound field," *J of Sound and Vib.*, vol. 143, no. 2, 1990.
- [41] B. Rafaely, *Fundamentals of spherical array processing*. Springer, 2015, vol. 8.
- [42] L. McInnes, J. Healy, and J. Melville, "UMAP: Uniform manifold approximation and projection for dimension reduction." arXiv, 2018.