

Blind room parameter estimation using multiple multichannel speech recordings

Prerak Srivastava, Antoine Deleforge, Emmanuel Vincent

► **To cite this version:**

Prerak Srivastava, Antoine Deleforge, Emmanuel Vincent. Blind room parameter estimation using multiple multichannel speech recordings. WASPAA 2021- IEEE Workshop on Applications of Signal Processing to Audio and Acoustics, Oct 2021, New Paltz, NY, United States. hal-03304656

HAL Id: hal-03304656

<https://hal.archives-ouvertes.fr/hal-03304656>

Submitted on 28 Jul 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

BLIND ROOM PARAMETER ESTIMATION USING MULTIPLE MULTICHANNEL SPEECH RECORDINGS

Prerak Srivastava, Antoine Deleforge, Emmanuel Vincent

Université de Lorraine, CNRS, Inria, Loria, F-54000 Nancy, France
{prerak.srivastava, antoine.deleforge, emmanuel.vincent}@inria.fr

ABSTRACT

Knowing the geometrical and acoustical parameters of a room may benefit applications such as audio augmented reality, speech dereverberation or audio forensics. In this paper, we study the problem of jointly estimating the total surface area, the volume, as well as the frequency-dependent reverberation time and mean surface absorption of a room in a blind fashion, based on two-channel noisy speech recordings from multiple, unknown source-receiver positions. A novel convolutional neural network architecture leveraging both single- and inter-channel cues is proposed and trained on a large, realistic simulated dataset. Results on both simulated and real data show that using multiple observations in one room significantly reduces estimation errors and variances on all target quantities, and that using two channels helps the estimation of surface and volume. The proposed model outperforms a recently proposed blind volume estimation method on the considered datasets.

1. INTRODUCTION

Audio augmented reality (AAR) has received increased interest in the recent years [1]. A key task in AAR is to simulate virtual sound sources that are consistent with the user’s environment via, e.g., a hear-through headset. To achieve this in an immersive way, a major hurdle is to estimate the relevant acoustical parameters of the room the user evolves in. The notion of *reverberation fingerprint* (RF) was introduced in [2] as a compact way to characterize rooms for realistic binaural rendering on AAR headphones. The RF consists of the room’s volume V in m^3 and its reverberation time per octave band $\text{RT}_{60}(b)$ in s , where $b \in \mathcal{B} = \{125, 250, 500, 1\text{k}, 2\text{k}, 4\text{k}\}$ Hz. Under ideal diffuse sound field conditions, these parameters are related via Sabine’s well-known formula [3]

$$\text{RT}_{60}(b) \approx 0.16 \frac{V}{\bar{\alpha}(b) \cdot S} \quad (1)$$

where $S = \sum_{i=1}^K S_i$ is the total area of the room’s K surfaces in m^2 and $\bar{\alpha}(b) = \sum_{i=1}^K \alpha_i(b) S_i / S$ is the area-weighted mean absorption coefficient in octave band $b \in \mathcal{B}$.

An attractive research direction to retrieve such parameters is to estimate them solely from audio recordings of unknown sound sources in the room, using microphones placed on the headset. In [4], an algorithm to blindly estimate the RF based on the decay envelope of a single-channel clean speech signal was proposed and showed encouraging results on simulated data. In the same year, the authors of [5] trained a statistical model on a large dataset generated from simulated room impulse responses (RIRs) to blindly estimate the position of a broadband source and the mean absorption

coefficient of walls above 1 kHz using a binaural receiver and inter-channel cues. Experiments were limited to a fixed room geometry, a fixed receiver position, and simulated data. Later, [6] paved the way for blind room volume estimation from single-channel noisy speech using a convolutional neural network trained using both simulated and real RIRs. The reported results showed that this method can estimate a broad range of volumes within a factor of 2 on real data from the Acoustic Characterization of Environments (ACE) challenge [7]. Following the ACE challenge, a number of neural-network-based methods were proposed to blindly estimate reverberation times and direct-to-reverberant ratios from speech signals [8, 9, 10, 11]. In particular, [11] proposes a method to generate augmented training datasets from real RIRs. These works estimate broadband values rather than frequency-dependent ones, which may limit the rendering realism in the context of AAR. They also all use single-channel signals. As an interesting alternative approach, [12] proposes to learn a latent representation referred to as *room embedding*, which is then used to condition a waveform-to-waveform network that converts a signal recorded in one environment to another. However, this approach cannot be straightforwardly applied to spatial AAR, whereas explicitly knowing room parameters offers more flexibility in binaural rendering. Moreover, explicit room parameters may be useful in themselves for other applications such as speech dereverberation [13] or audio forensics [14].

In this paper, we push these recent research efforts further by proposing a blind method to jointly estimate the total surface area S , the volume V , the reverberation times $\text{RT}_{60}(b)$ and the mean absorption coefficients $\bar{\alpha}(b)$ in all octave bands $b \in \mathcal{B}$. On top of that, we investigate whether using multichannel input and/or fusing estimates across multiple source-receiver positions inside the same room help. To this end, we propose a new convolutional neural network architecture that combines single- and inter-channel features in a joint embedding layer, and uses a likelihood-based loss function that yields adaptive variance estimates, which allow us to fuse multiple independent observations of the same room in a statistically principled way. The model is trained on a carefully-crafted realistic simulated dataset and tested on both simulated and real wet speech recordings¹.

2. SIMULATED DATASET

To train and validate our approach, a large dataset of noisy speech signals annotated with target room acoustic parameters is created. First, synthetic RIRs are generated using Roomsim [15]. Roomsim

¹Our code to reproduce this paper is available at github.com/prerak23/RoomParamEstim.

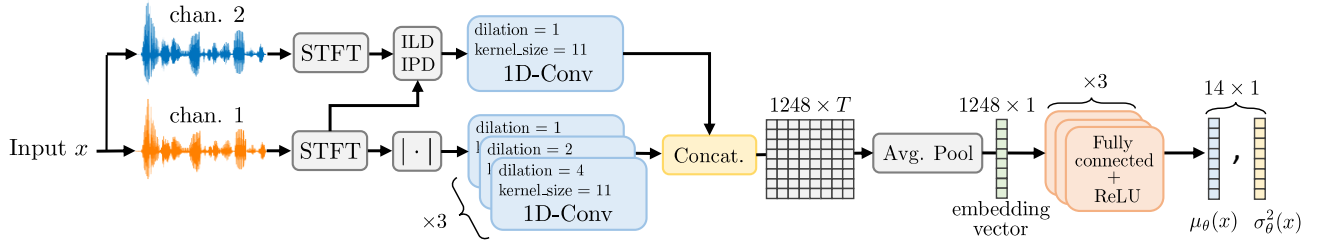


Figure 1: Diagram of the proposed neural network architecture.

is a hybrid shoe-box room acoustic simulator combining the image-source method [16] to simulate specular reflections that dominate the early part of RIRs, and the diffuse-rain method [15] to model scattered reflections using stochastic ray-tracing. Simulations are run using a sampling frequency of 48 kHz, a reflection order of 10 for the image-source method and 2000 rays for the diffuse-rain method. 20,000 rooms are simulated with length, width and height drawn uniformly at random in [3, 10] m, [3, 10] m and [2.5, 4] m, respectively, resulting in $S \in [48, 360] \text{ m}^2$ and $V \in [18, 400] \text{ m}^3$. For each room, 5 RIRs corresponding to different source-receiver positions are generated. Omnidirectional sources and receivers are used and placed uniformly at random at least 30 cm from each surface and from each other. The receiver is a two-microphone array placed parallel to the floor with an aperture of 22.5 cm, that is a typical head or headset width. For each room, a single scattering coefficient drawn uniformly at random in [0.2, 1] is used on all surfaces and octave bands for the diffuse-rain method. The absorption coefficients of the 6 surfaces in 6 octave bands are sampled using the *reflectivity-biased* sampling strategy proposed in [17]. Each $\alpha_i(b)$ is drawn uniformly at random inside ranges based on measured databases of commonly encountered surface materials, while giving each surface a 50% probability of having a frequency-independent reflective profile ($\alpha_i(b) < 0.12$). As showed in [17], this yields more realistic and more diverse distributions of absorption and reverberation times ($\text{RT}_{60}(b) \in [0.2, 3.2] \text{ s}$ and $\bar{\alpha}(b) \in [0.02, 0.6]$ in our dataset) than, e.g., by sampling each $\alpha_i(b)$ uniformly in [0, 1].

The obtained RIRs are downsampled to 16 kHz and convolved with random speech excerpts from the LibriSpeech [18] corpus. The resulting 3 s two-channel reverberated signals are then corrupted with both static microphone noise, i.e., independent additive white Gaussian noise on each channel, and spatially-diffuse babble noise, i.e., speech-shaped noise convolved with the late part (>50 ms) of a random RIR in the room. For the noise levels to be realistic, signals from sources that are placed further away from the receiver should exhibit lower signal-to-noise ratios (SNRs). To achieve this, for each room, we first generate a *reference* signal using a random speech source placed 1 meter in front of a receiver (not used in the final dataset). Static and diffuse noise levels for this signal are set to obtain SNRs drawn uniformly at random in [70, 90] dB and [30, 60] dB, respectively (different levels for each source-receiver position in the room are set). These noise levels are then kept fixed for the final mixtures, irrespective of the distance from the speech source to the receiver. This resulted in an overall SNR range of [-10, 65] dB across the dataset.

The $5 \times 20 \text{ k} = 100 \text{ k}$ two-channel noisy speech signals are divided into training, validation and test sets of respective size 80k,

10k and 10k with no room or speech-signal overlap between them. Each room is annotated with the 14 target parameters, namely, 6 mean absorption coefficients and 6 reverberation times in all octave bands, the total surface area S and the volume V . For each room and each octave band, a unique reverberation time is estimated by taking the median value over the 5 source-receiver positions available. The values are obtained by linear regression over the -5 dB to -25 dB decay of Schroeder curves [19].

3. NEURAL NETWORK MODEL

3.1. Proposed Architecture, Cost Function, and Fusion Method

The proposed neural network architecture is depicted in Fig. 1. Single-channel (SC) and inter-channel (IC) features are extracted from the time-domain two-channel input signal x in the form of spectrograms. We use the short-time Fourier transform (STFT) with 96 ms Hann windows and 50% overlap to obtain a complex spectrogram $\{X_i(f, t)\}_{f, t=1}^{F, T}$ for each channel i , with $F = 769$ positive frequency bins and $T = 63$ time frames for a 3 s input signal (our architecture works on arbitrary input length). This choice performed best among 32, 64, 96 and 128 ms windows in our preliminary experiments. Then, SC features are computed as $|X_1(f, t)|$, which performed better than $|X_1(f, t)|^2$, $|X_1(f, t)|^{1/2}$, or $\log |X_1(f, t)|$. IC features are obtained by concatenating inter-channel level differences (ILD) and phase differences (IPD):

$$\text{ILD}(f, t) = \log |X_1(f, t)| - \log |X_2(f, t)|, \quad (2)$$

$$\text{IPD}(f, t) = \left[\text{Re}, \text{Im} \left(\frac{X_1(f, t)X_2^*(f, t)}{|X_1(f, t)X_2^*(f, t)|} \right) \right]. \quad (3)$$

These features are then processed through 1D convolutional blocks (1D-Conv), which were recently proposed in the Conv-TasNet architecture in the context of speech separation [20]. These blocks consist of separable convolutions (depth-wise and point-wise) intertwined with rectified linear unit (ReLU) activations and followed by layer normalization [21]. The latter proved to be crucial in our experiments, as it creates scale-invariant representations. For SC features, three 1D-Conv blocks with increasing dilation factors along the frequency-axis and a kernel size of 11 are used, while only one block is used on IC features, as this showed to give best results. The obtained representations are concatenated along the frequency axis and average-pooled along the time-axis to yield a time-independent, 1248-dimensional *embedding vector*. The embedding vector is finally passed through 3 fully-connected layers of respective dimensions 96, 48 and 28 to obtain $2 \times D = 14$ outputs consisting of the

estimated room parameters $\mu_{\theta}(\mathbf{x}) \in \mathbb{R}^D$ and the estimated variances $\sigma_{\theta}^2(\mathbf{x}) \in \mathbb{R}^D$ (or uncertainties) on these parameters.

The network parameters θ are optimized by minimizing the following Gaussian negative log-likelihood loss function:

$$\begin{aligned} \mathcal{L}_{\theta}(\mathbf{x}, \mathbf{y}) &= -\log p_{\theta}(\mathbf{y}|\mathbf{x}) = -\log \mathcal{N}(\mathbf{y}; \mu_{\theta}(\mathbf{x}), \sigma_{\theta}^2(\mathbf{x})) \\ &= \frac{c}{2} \sum_{d=1}^D \log \sigma_{d,\theta}^2(\mathbf{x}) + \frac{(y_d - \mu_{d,\theta}(\mathbf{x}))^2}{\sigma_{d,\theta}^2(\mathbf{x})} \end{aligned} \quad (4)$$

where $\mathbf{y} \in \mathbb{R}^D$ denotes the ground truth room parameters. A benefit of this approach is that it adaptively weights errors on individual parameters. The estimated variances can also be used to fuse estimated obtained from J independent observations $\{\mathbf{x}_j\}_{j=1}^J$ of the same room using the following formula derived from Bayes' rule:

$$p_{\theta}(y_d|\bar{\mathbf{x}} = [\mathbf{x}_1, \dots, \mathbf{x}_J]) = \mathcal{N}(y_d; \bar{\mu}_{d,\theta}(\bar{\mathbf{x}}), 1/\bar{\gamma}_{d,\theta}^2(\bar{\mathbf{x}})) \quad (5)$$

$$\text{with } \bar{\mu}_{d,\theta}(\bar{\mathbf{x}}) = \sum_{j=1}^J \frac{\gamma_{d,\theta}^2(\mathbf{x}_j)}{\bar{\gamma}_{d,\theta}^2(\bar{\mathbf{x}})} \mu_{d,\theta}(\mathbf{x}_j), \quad \bar{\gamma}_{d,\theta}^2(\bar{\mathbf{x}}) = \sum_{j=1}^J \gamma_{d,\theta}^2(\mathbf{x}_j)$$

where $\gamma_{d,\theta}^2(\mathbf{x}_j) = 1/\sigma_{d,\theta}^2(\mathbf{x}_j)$ is the estimated precision for observation \mathbf{x}_j and $\bar{\mu}_{d,\theta}(\bar{\mathbf{x}})$ is the fused estimate.

To avoid issues due to scale differences, the ground truth parameters are normalized at training time by dividing them by their standard deviations over the training set, which are saved and multiplied with the network output at test time. The network is trained using the ADAM optimizer [22] with a learning rate of 10^{-4} and a batch size of 120. A dropout rate of 0.2 and 0.4 was used in conv-blocks and in fully connected layers to avoid over-fitting. We used a patience of 15 epochs on the validation set for early stopping. Our models generally converged in 100–150 epochs.

3.2. Other Tried Architectures

In addition to 1D-Conv blocks, Conv-TasNet introduced the idea of learnable input filters [20]. We tried replacing the STFT by such learnable filters to estimate our SC features, but this yielded worse performance. We also tried an architecture inspired by x-vectors, a state-of-the-art model for speaker recognition [23], but this did not perform well on our data. We experimented with replacing the proposed maximum-likelihood loss (4) by a simple mean squared error loss. This yielded significantly higher errors while losing the benefit of easily fusing multiple observations in one room. Estimating $\log(V)$ and $\log(S)$ instead of V and S as suggested in [6] did not improve results on our data, most likely because the ranges we consider for these quantities are smaller. Finally, instead of estimating $\bar{\alpha}$, RT_{60} , V and S jointly in a multi-task fashion, we trained the same architecture to estimate each of these quantities individually. The obtained results were identical, while increasing the number of model parameters and training time by a factor of 4, which suggests that our proposed network achieves efficient weight sharing.

4. EXPERIMENTS AND RESULTS

4.1. Simulated Data

We first evaluate the proposed approach on the 2,000 unseen rooms of the simulated test set (see Section 2), using between 1 and 5 two-channel 3-s noisy speech recordings at different source-receiver positions per room, based on the fusion method in (5). Two variants of

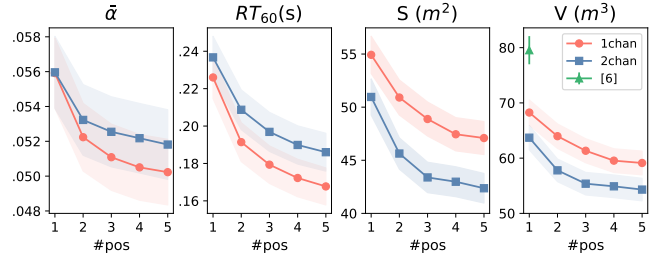


Figure 2: Mean absolute error achieved on simulated data by [6] vs. the proposed model with one- or two-channel inputs, as a function of the number of source-receiver positions fused in each room. Shaded areas indicate 95% confidence intervals.

Input	Feature	$\bar{\alpha}$	RT_{60} (s)	S (m^2)	V (m^3)
1mic, 1sig	SC	0.055	0.225	55.0	68.8
1mic, 2sig	SC	0.058	0.222	55.4	69.7
2mic, 1sig	SC	0.057	0.221	54.0	67.7
2mic, 1sig	SC+IC	0.055	0.236	49.9	63.7

Table 1: Mean absolute error achieved on simulated data for one source-receiver position using different inputs and features. Bold numbers indicate the best statistically significant result per column based on 95% confidence intervals, when there is one.

our model are compared: the full two-channel architecture depicted in Fig. 1, and the same architecture without the upper inter-channel processing part of the network, i.e., using only one channel. These variants are compared to our implementation of the single-channel, single-position blind room volume estimation method of [6] trained on the same data. We could not find competing methods for blind total surface estimation or blind frequency-dependent reverberation time and mean absorption coefficient estimation in the literature. The metric used is the mean absolute error on each room parameter.

As can be seen in Fig. 2, increasing the number of fused observations per room significantly reduces errors on all parameters. Using five positions and two channels, mean absolute errors of 0.052 for $\bar{\alpha}$ (training range [0.02, 0.6]), 0.18 s for RT_{60} (training range [0.2, 3.2]), 42 m^2 for S (training range [48, 360]) and 54 m^3 for V (training range [18, 400]) are obtained. The proposed model significantly outperforms the one in [6] for volume estimation, reducing the error by 13% using one channel and one observation, and by 31% using two channels and five observations. Interestingly, we observe that using two channels instead of one significantly reduces surface and volume estimation error without significantly impacting mean absorption and reverberation time estimation, according to 95% confidence intervals. This may be interpreted by the fact that the latter mostly govern the late, spatially-diffuse regime of RIRs, and hence should have limited correlation with inter-channel cues that mostly capture spatial characteristics. Conversely, S and V are inherently spatial quantities as they relate to the room's geometry and hence early echoes, which do correlate with IC cues [24, 25].

We hence carry an additional experiment to study whether the improvement observed when using two channels is truly explained by concatenating SC and IC representations at the embedding layer (see Fig. 1), or is only due to doubling the number of signal samples in a given source-receiver position, which mitigates noise. Four different embeddings are compared, namely (i) the SC representation of one speech signal (1mic, 1sig), (ii) the average of the SC

Octave bands	1 position		5 positions	
	$\bar{\alpha}$	RT ₆₀ (s)	$\bar{\alpha}$	RT ₆₀ (s)
125 Hz	0.056	0.392	0.051	0.320
250 Hz	0.060	0.305	0.055	0.249
500 Hz	0.057	0.228	0.051	0.170
1 kHz	0.053	0.188	0.050	0.146
2 kHz	0.054	0.165	0.051	0.122
4 kHz	0.052	0.139	0.049	0.106

Table 2: Mean absolute errors on $\bar{\alpha}(b)$ and RT₆₀(b) in the 6 octave bands achieved by the proposed model with two channels and 1 or 5 source-receiver positions per room on simulated data.

representations of two speech signals (1mic, 2sig), (iii) the concatenation of the SC representations of the two individual channels of one speech signal (2mic, 1sig) and (iv) the proposed concatenation of SC and IC representations. Speech signals are always 3-s long. As can be seen in Table 1, the only model which significantly improves surface and volume estimation is the one using IC features, despite the fact that the 2nd and 4th models also benefit from the same number of input signal samples.

So far, the reported errors on $\bar{\alpha}$ and RT₆₀ were averaged over all 6 octave bands. Table 2 reports detailed errors per octave band. As can be seen, the mean absorption is well estimated in all octave bands, while reverberation-time errors steadily decrease by a factor of 3 from 125 Hz to 4 kHz. This could be explained by the fact that less information is available in the narrower, lower octave bands.

4.2. Real Data

To check how well our method generalizes to real data, we use the recently released dEchorate dataset [25], which contains wet speech recordings made inside the acoustic lab at Bar-Ilan University. The lab is a shoe-box room of size $5.7 \times 6 \times 2.4$ m ($S = 125$ m², $V = 82$ m³) where each of the walls, floor and ceiling can be set either to either a reflective mode or an absorbent mode. 5 arrays of 6 omnidirectional microphones and 6 directional loudspeakers are placed inside the room, yielding $5 \times 6 = 30$ multi-channel speech recordings per room configuration. For each room, ground truth $\bar{\alpha}$ and RT₆₀ values are provided for the four octave bands from 500 Hz to 4 kHz. Hence, results in lower octave bands are omitted here. In our experiments, we use $3 \times 30 = 90$ 3-s speech recordings corresponding to the 2-channel sub-arrays with aperture 22.5 cm and to the 3 room configurations involving 3 or more reflective surfaces, as these most closely match the considered scenario. For these rooms, $\bar{\alpha}(b)$ ranges from 0.16 to 0.35 and RT₆₀(b) from 0.25 to 0.66 s.

Table 3 reports mean absolute errors using the proposed approach with or without IC features. We report results using either 1 or 5 source positions and a fixed receiver. For the latter, we exclude one out of the 6 available source positions for each test, so that there are 90 tests in each case. The mean absolute volume estimation error using the single-channel, single-position approach of [6] is reported as well. Encouragingly, errors obtained with our approach are of comparable orders to those obtained on simulated data. In the single-channel, single-position case, volume estimation errors obtained with our model are comparable to [6]. For RT₆₀, S and V , we are able to reproduce the observation that increasing the number of source-receiver positions significantly decreases errors, at least using IC features. We also observe again that using IC

Method	Features	# pos	$\bar{\alpha}$	RT ₆₀	S	V
[6]	SC	1	-	-	-	137.8
Ours	SC	1	0.061	0.134	129.6	154.5
Ours	SC	5	0.060	0.097	125.8	149.1
Ours	SC+IC	1	0.084	0.101	89.4	107.6
Ours	SC+IC	5	0.094	0.062	50.2	68.8

Table 3: Mean absolute error achieved over 3 rooms from the real dEchorate dataset. Bold numbers indicate the best statistically significant result per column, based on 95% confidence intervals.

Method	Features	# pos	$\bar{\alpha}$	RT ₆₀	S	V
[6]	SC	1	-	-	-	10.0
Ours	SC	1	0.030	0.161	27.2	31.8
Ours	SC	5	0.024	0.090	19.6	23.0
Ours	SC+IC	1	0.031	0.100	34.7	39.7
Ours	SC+IC	5	0.015	0.054	16.5	18.9

Table 4: Standard deviation of parameter estimates for room "011100" of the real dEchorate dataset.

features significantly improves the estimation of S and V . The errors obtained on $\bar{\alpha}$ are less consistent than those on simulated data, which may be due to the difficulty of reliably annotating mean absorption coefficients in real rooms.

Finally, Table 4 shows the standard deviations of estimated values by the same methods over the 30 recordings from the room with 3 reflective surfaces. Encouragingly, the relatively low standard deviations reveal the ability of the models to provide parameter estimates that are stable within a room, and do not depend much on the source-receiver position. Moreover, it can be seen that, as expected, using five observations in a room instead of one systematically decreases the standard deviation of estimates.

5. CONCLUSION

This study revealed that using inter-channel cues can significantly improve the blind estimation of a room's volume and surface from noisy speech, while for estimating reverberation and absorption parameters, a single channel is sufficient. It also highlights that fusing multiple measurements reduces estimation errors and variances on all parameters. Finally, we showed that a system trained on a carefully simulated training set offers reasonable generalization capabilities to real data. Future work will include the use of data-augmentation and domain adaptation techniques to improve real data results, extensions to binaural or ambisonic receivers, and the joint estimation of local parameters such as the positions and properties of the source and individual surfaces in the room, with the help of early acoustic echoes.

6. ACKNOWLEDGMENTS

This work was made with the support of the French National Research Agency through project HAIKUS "Artificial Intelligence applied to augmented acoustic scenes" (ANR-19-CE23-0023). Experiments presented in this paper were carried out using the Grid'5000 testbed, supported by a scientific interest group hosted by Inria and including CNRS, RENATER and several Universities as well as other organizations (see <https://www.grid5000.fr>). We thank the authors of [6] for their helpful advice.

7. REFERENCES

- [1] V. Valimaki, A. Franck, J. Ramo, H. Gamper, and L. Savioja, "Assisted listening using a headset: Enhancing audio perception in real, augmented, and virtual environments," *IEEE Signal Processing Magazine*, vol. 32, no. 2, pp. 92–99, 2015.
- [2] J.-M. Jot and K. S. Lee, "Augmented reality headphone environment rendering," in *2016 AES International Conference on Audio for Virtual and Augmented Reality*, 2016, pp. 8–2.
- [3] C. F. Eyring, "Reverberation time in "dead" rooms," *The Journal of the Acoustical Society of America*, vol. 1, no. 2A, pp. 217–241, 1930.
- [4] P. Murgai, M. Rau, and J.-M. Jot, "Blind estimation of the reverberation fingerprint of unknown acoustic environments," in *AES 143rd Convention*, 2017, p. 9905.
- [5] S. Kataria, C. Gaultier, and A. Deleforge, "Hearing in a shoebox: binaural source position and wall absorption estimation using virtually supervised learning," in *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2017, pp. 226–230.
- [6] A. F. Genovese, H. Gamper, V. Pulkki, N. Raghuvanshi, and I. J. Tashev, "Blind room volume estimation from single-channel noisy speech," in *2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2019, pp. 231–235.
- [7] J. Eaton, N. D. Gaubitch, A. H. Moore, and P. A. Naylor, "Estimation of room acoustic parameters: The ACE challenge," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 24, no. 10, pp. 1681–1693, 2016.
- [8] F. Xiong, S. Goetze, and B. T. Meyer, "Joint estimation of reverberation time and direct-to-reverberation ratio from speech using auditory-inspired features," *arXiv preprint arXiv:1510.04620*, 2015.
- [9] H. Gamper and I. J. Tashev, "Blind reverberation time estimation using a convolutional neural network," in *16th International Workshop on Acoustic Signal Enhancement (IWAENC)*, 2018, pp. 136–140.
- [10] M. Chai, T. Li, M. Zhu, T. Wang, and W. Zhang, "Blind estimation of reverberation time using binaural complex ideal ratio mask," in *2019 IEEE International Conference on Multimedia & Expo Workshops (ICMEW)*, 2019, pp. 378–383.
- [11] N. J. Bryan, "Impulse response data augmentation and deep neural networks for blind room acoustic parameter estimation," in *2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2020, pp. 1–5.
- [12] J. Su, Z. Jin, and A. Finkelstein, "Acoustic matching by embedding impulse responses," in *2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2020, pp. 426–430.
- [13] B. Wu, M. Yang, K. Li, Z. Huang, S. M. Siniscalchi, T. Wang, and C.-H. Lee, "A reverberation-time-aware DNN approach leveraging spatial information for microphone array dereverberation," *EURASIP Journal on Advances in Signal Processing*, vol. 2017, no. 1, pp. 1–13, 2017.
- [14] A. H. Moore, M. Brookes, and P. A. Naylor, "Roomprints for forensic audio applications," in *2013 IEEE Workshop on Applications of Signal Processing to Audio and Acoustics (WASPAA)*, 2013, pp. 1–4.
- [15] S. M. Schimmel, M. F. Muller, and N. Dillier, "A fast and accurate "shoebox" room acoustics simulator," in *2009 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2009, pp. 241–244.
- [16] J. Allen and D. Berkley, "Image method for efficiently simulating small-room acoustics," *The Journal of the Acoustical Society of America*, vol. 65, pp. 943–950, 1979.
- [17] C. Bastien, A. Deleforge, and C. Foy, "Mean absorption coefficient estimation from impulse responses: Deep learning vs. Sabine," in *E-FA 2020 - Forum Acusticum 2020*, Dec. 2020.
- [18] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, "Librispeech: An ASR corpus based on public domain audio books," in *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2015, pp. 5206–5210.
- [19] M. R. Schroeder, "New method of measuring reverberation time," *The Journal of the Acoustical Society of America*, vol. 37, no. 6, pp. 1187–1188, 1965.
- [20] Y. Luo and N. Mesgarani, "Conv-TasNet: Surpassing ideal time–frequency magnitude masking for speech separation," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 27, no. 8, p. 1256–1266, Aug 2019.
- [21] J. L. Ba, J. R. Kiros, and G. E. Hinton, "Layer normalization," *stat*, vol. 1050, p. 21, 2016.
- [22] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," in *ICLR (Poster)*, 2015. [Online]. Available: <http://arxiv.org/abs/1412.6980>
- [23] D. Snyder, D. Garcia-Romero, G. Sell, D. Povey, and S. Khudanpur, "X-vectors: Robust DNN embeddings for speaker recognition," in *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2018, pp. 5329–5333.
- [24] D. Di Carlo, A. Deleforge, and N. Bertin, "MIRAGE: 2D source localization using microphone pair augmentation with echoes," in *2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2019, pp. 775–779.
- [25] D. Di Carlo, P. Tandeitnik, C. Foy, N. Bertin, A. Deleforge, and S. Gannot, "dEchorate: a calibrated room impulse response database for echo-aware signal processing," 2021, (under review). [Online]. Available: <https://hal.archives-ouvertes.fr/hal-03207860>