

# BLIND ACOUSTIC ROOM PARAMETER ESTIMATION USING PHASE FEATURES

Christopher Ick<sup>1</sup>, Adib Mehrabi<sup>2</sup>, Wenyu Jin<sup>3</sup>.

<sup>1</sup> New York University, New York, NY 10003, United States

<sup>2</sup> Sonos Experience Limited, United Kingdom <sup>3</sup> Sonos, Inc., Boston, MA 02111, United States

## ABSTRACT

Modeling room acoustics in a real-world settings involves some degree of blind parameter estimation from noisy and reverberant audio. Modern approaches leverage convolutional neural networks (CNNs) in tandem with time-frequency representations. Using short-time Fourier transforms to develop these spectrogram-like features has shown promising results, but this method implicitly discards a significant amount of audio information in the phase domain. Inspired by recent works in speech enhancement, we propose utilizing phase-related features to extend recent approaches to blindly estimate the so-called “reverberation fingerprint” parameters, namely, volume and  $RT_{60}$ . The addition of these features is shown to outperform existing methods that rely solely on magnitude-based spectral features across a wide range of acoustics spaces. We evaluate the effectiveness of the deployment of these phase features in both single-parameter and multi-parameter estimation strategies, using a task-specific dataset that consists of publicly available room impulse responses (RIRs), synthesized RIRs, and in-house measurements of real acoustic spaces.

## 1. INTRODUCTION

In recent years, the dynamic parameterization of listeners’ local acoustic space has become increasingly relevant in the area of audio processing. Knowledge of the local room acoustics can be used to enhance speech signals or aid dereverberation algorithms, for the purpose of word recognition or voice communication clarity [1, 2, 3]. Spatial sound reproduction systems [4, 5, 6] may also benefit from this information for acoustic room equalization tasks. To this end, one of the major challenges is to estimate the best acoustical parameters of the room to improve the plausibility of immersive audio. In [7], the concept of “reverberation fingerprint”, which consists of the room’s volume and its frequency-dependent diffuse reverberation decay time, was introduced to characterize rooms for realistic binaural rendering on audio augmented reality headphones. This fingerprint is limited to the position-independent part of the reverberation as it characterizes a room in isolation from the orientation and directivity of sources and receivers.

Conventionally, room parameters such as reverberation time ( $RT_{60}$ ) and direct-to-reverberant ratio (DRR) can be derived directly from measured RIRs. Room volume is highly related to the estimation of the so-called critical distance, which is defined as the distance at which the direct and reverberant power portions of a sound source become equivalent, i.e., the point at which the DRR is 0 dB. Under the ideal diffuse soundfield assumption, the relation between these parameters is given by the well known Sabine’s equation [8]. In practice, in-situ RIR measurements of users’ local acoustic environments are typically difficult to achieve. An attractive alternative is to estimate room acoustic parameters blindly from audio recordings of unknown sound sources using microphones in the presence of noise. The 2015 ACE challenge [9] set the bar for blindly estimating  $RT_{60}$  and DRR from noisy speech signals in the presence of various noise

types, with signal-modelling based approaches for  $RT_{60}$ [10, 11] being the top-performing systems. On the other hand, room volume estimation has long been formulated as a classification problem. The systems in [12, 13] leverage MFCC-based features to identify the room that corresponds to a environmental sound or speech recording from a closed set for the purpose of audio-forensics. In [14], a blind estimation algorithm for the reverberation fingerprint was proposed based on the decay envelope of a single-channel clean speech signal using simulated data.

Recent advancements in deep neural networks also push the boundary of room acoustic parameters estimation. An approach of blind room volume estimation from single-channel noisy speech deploying a CNN was proposed in [15], estimating volume within a factor of 2 from the ACE challenge [9]. A number of neural-network-based systems were proposed to blindly estimate  $RT_{60}$  and DRR from single-channel speech signals [16, 17, 18, 19] and demonstrated promising results in terms of both estimation accuracy and the ability to follow temporal variations in dynamically evolving acoustic environments. Srivastava et al. [20] devised a system that jointly estimates multiple room parameters including the total surface area, the volume, the frequency-dependent  $RT_{60}$  and mean surface absorption coefficients. However, the system relies on two-channel noisy speech recordings to explore inter-channel features, which implies it is more hardware-demanding.

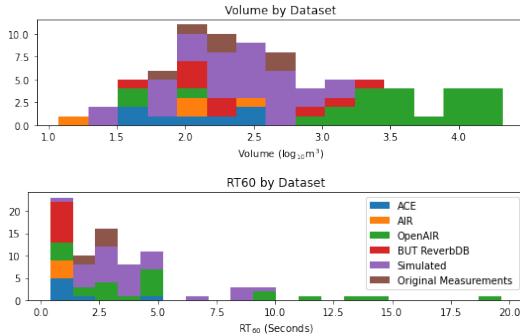
In this paper, we propose using a CNN model to estimate room acoustic parameters blindly from single-channel noisy speech signals with the special focus on the reverberation fingerprint (i.e. volume and  $RT_{60}$ ). Unlike prior works that generally rely on the log-energy calculation of spectro-temporal features, we introduce phase-related features. To the authors’ best knowledge, this is the first DNN-based room parameter estimation system that leverages phase-related features. Specifically, the Gammatone phase spectrogram and its numerical derivatives that denote relative phase variations across the time and frequency axes are considered. Results on both simulated and real-world RIRs show that the proposed model outperforms state-of-the-art single channel blind room acoustic parameter estimation systems on similarly-structured datasets.

## 2. DATA GENERATION PIPELINE

Approaching the task of room parameter estimation with neural-network-based approaches is inherently challenging due to the scale of data required. Because the task relies on having audio samples from a wide variety of rooms, building an appropriately diverse dataset by hand would be impractically expensive and time consuming. For our work, we instead develop a multi-stage audio generation pipeline that makes use of a wide range of RIRs across various room configurations with labeled volume and broad-band  $RT_{60}$ .

### 2.1. Real-world RIR Dataset

To maximize the range of room acoustic parameters, we sampled audio simulations from four different publicly available datasets



**Fig. 1:** A histogram showing the distribution of RIR over various volumes and  $RT_{60}$  ranges from various datasets

to provide 43 distinct room profiles, with volumes ranging from single-person phone booths, to nuclear reactor halls. Meanwhile,  $RT_{60}$  values ranging from less than half a second, to several over 10 seconds, which were computed using the Schroeder method [21]. To ensure consistency across datasets, all RIRs were downsampled to 16 kHz. Because many of these datasets were captured in academic settings, most rooms are rectangular in shape, and fall generally into the category of offices, classrooms, and auditoriums/lecture halls. This includes the ACE Challenge dataset [9], the Aachen impulse response (AIR) dataset [22], and the Brno University of Technology Reverb Database (BUT ReverbDB) [23]. Lastly, the OpenAIR dataset [24] covers primarily large acoustic spaces (cathedrals, nuclear reactors, and other large structures). To account for the natural gap of real-world acoustic spaces falling in the  $100 \text{ m}^3$  to  $1000 \text{ m}^3$  range of volumes, we made several in-house RIR measurements within this range to augment our available data, by recording sine-sweeps in shoebox-shaped rooms of measured volume.

## 2.2. Synthetic RIRs

To increase the scale of data, particularly in the range of less frequently available room volumes, we added in 30 simulated RIRs for rooms of various geometries. Each room was simulated with a singular source near the center of the room, and five point receivers uniformly sampled over the volume of each room. To generate this dataset, we made use of *pyroomacoustics* [25], a software package that uses the image-source model for RIR simulation for specific room geometries. While this geometric method doesn't account for effects such as diffraction and scattering, experiments show that the use of simulated data benefits the performance of the model, which generalizes well to natural data [15]. Overall, Fig. 1 illustrates the distribution of room volumes and  $RT_{60}$  on logarithmic and linear scales, respectively. To ensure that our data was not overfitting to the well-behaved simulated data, our validation and test data splits did not include any data from the simulated datasets (except in the case of experiments that were explicitly measuring the performance of the model on simulated data).

## 3. AUDIO GENERATION

From this RIR dataset with known volume and  $RT_{60}$  values, we could then generate reverberant audio data that would be eventually fed into our CNNs. To do this, we convolved speech data recorded in anechoic chambers with RIRs, effectively mapping the acoustic response of the room to the audio. This audio was combined with white noise convolved with the same RIR, the same method for

Data Split	# of	Real Rooms	Simulated Rooms
Train	12000	29	18
Validation	4000	16	12
Test	4000	16	0

**Table 1:** Summary of data splits

noisy data generation as done in the ACE challenge [9]. To enforce robustness across various noisiness conditions, we combined these two signals to simulate the recording at 5 distinct SNR levels, including no noise (+Inf) and [+30,+20,+10,+0]dB.

Speech samples were drawn from the ACE dataset [9] of anechoic voice recordings in randomly selected 4 second increments, with both male and female speakers. Within the RIR dataset, certain rooms had more RIR measurements than others, so to ensure uniformity of representation, each room was equally sampled, so that the distribution of audio samples in our dataset matched the distribution of volumes in our RIR dataset, as seen in Fig. 1.

The data was split into a 6-2-2 train, validation and test split by rooms, as shown in Table 1. Specific rooms were selected for the validation and test splits to ensure a wide range of specific volumes that remained unseen by the models during the training.

## 4. FEATURIZATION

As the primary input to a CNN, the featurization of audio into a two-dimensional time-frequency representation has a non-trivial amount of influence on the performance of the model. Multiple featurization schemes can be combined to train a single model, but increasing the dimensionality of the input space can lead to highly complex models that are prohibitively data hungry and costly to train. As such, choosing an audio representation that maximally captures information about the acoustic space must be balanced by discarding features that provide no benefit to the model. This ensures that we are not confined by the relatively small data size we have access to, as well as hopefully providing more generalizable results.

### 4.1. Spectral Representations

Prior literature [15, 17, 26] suggests the importance of low-frequency effects for room acoustic parameter estimation, constraining their feature representations to relatively low frequency bands ( $< 2000$  kHz). To keep model complexity low while preserving signal information that is relevant for the parameter estimation problem, we used Gammatone ERB filterbank to generate our time-frequency representation. The Gammatone filterbank has 20 bands, ranging from 50Hz to 2000Hz. In computing the STFT of our audio, we used a Hann window size of 64 samples, with a hop size of 32. The resulting spectrogram is convolved with our filterbank, resulting in a  $20 \times 1997$  complex Gammatone spectrogram. Additionally, as detailed in [15], a number of separate energy-based features that aim to capture the low-frequency behaviour were considered specifically for improving the room volume estimation performance.

### 4.2. Phase Representation

During the calculation of the log-energy of the Gammatone spectrogram, phase information is typically discarded. However, recent literature in the area of speech enhancement has demonstrated the effectiveness of including various forms of phase information, either explicitly or implicitly [27]. Being inspired by this, we propose retaining the phase information captured in the STFT and using it as an input for similar networks. By computing the phase angle at every time-frequency bin of Gammatone features, we can generate a

Gammatone phase spectrogram equivalently sized to our log-energy spectrogram.

In addition to the direct phase calculations, recent work in [27] has shown value in numerical derivatives of the phase spectrograms. This provides us with a new feature that could potentially better characterize room reflection patterns at low frequencies, compensating for the nature of low spectral resolution due to the ERB filterbank. Additionally, we also compute the *phase continuity*  $\phi$ , which is a second degree partial derivative of the phase information along the frequency axis  $k$ , and the time axis  $n$ .

$$\phi(\theta_n^k) = \lim_{i \rightarrow 0} \lim_{j \rightarrow 0} \frac{f(\theta_n^k) - f(\theta_{n+j}^{k+i})}{\theta_n^k - \theta_{n+j}^{k+i}} \quad (1)$$

where  $f(\theta) = \cos(\theta) + \sin(\theta)$  is a phase-wrapping function as defined in [27]. By concatenating various features together, we can compare the performance benefits of these features, while observing the trade-offs of training time and model complexity.

## 5. EXPERIMENT DESIGN

### 5.1. Network Architecture

The model architecture was selected to be a feed-forward 2D CNN due to its suitability for capturing 2-dimensional time-frequency signal patterns. Because the focus of this study is comparison of audio featurizations, the choice of architecture was not critical, and such, this relatively simple architecture was employed, for it's success in the same task in [15]. The model makes use of 4 time-domain convolutional layers, followed by 2 spectral convolutions, resulting in a total of 6 convolutional layers, each one containing an average pooling layer. The convolutional block was then followed by a 50% dropout layer and a final, single fully-connected layer that mapped down to output dimensionality.

Fig. 2 illustrates the full architecture employed.

### 5.2. Evaluation Metrics

The models were evaluated on their ability to estimate the room parameters within a feasible accuracy. Perceptual evaluation of the dataset had shown that there were many more variations in the lower volume range than across the wide range of larger volumes, so to account for that, we evaluated our data on the base-10 logarithm of the volume. This also mitigated the effect that larger acoustic spaces would be over-weighted in training due to the relative higher contribution to error estimations.

While this effect is less pronounced for  $RT_{60}$ , we also computed the log of  $RT_{60}$  for evaluation purposes, as it allows smaller values closer to 0 to be more widely distributed.

We evaluated the model on several standard regression metrics, including mean squared error (MSE) and the Pearson correlation coefficient ( $\rho$ ) between our target and predicted values. One additional metric we used is the exponential of the mean absolute logarithm of the ratio between estimated and target room parameters. Given target  $y$ , and estimate  $\hat{y}$ , we can write this as:

$$\text{MeanMult} = \exp \left( \frac{1}{N} \sum_{n=1}^N \left| \ln \left( \frac{\hat{y}_n}{y_n} \right) \right| \right) \quad (2)$$

This metric summarizes mean error in terms of the ratio between estimated and target parameters in a method that is less sensitive to overall scale, allowing more equal weighting across all scales.

Finally, we used one additional metric comparing the variance between all estimates and all targets, which is mainly used to determine if our model is overfitting to a singular estimation, in which

Feature Set	Volume Estimation			
	MSE	$\rho$ -corr.	MM	VR
Baseline [15]	0.4184	0.5736	1.6941	0.6624
	<b>0.3630</b>	<b>0.6479</b>	<b>1.6084</b>	<b>0.7461</b>
	0.3843	0.6299	1.6365	0.7415
Feature Set	RT <sub>60</sub> Estimation			
	MSE	$\rho$ -corr.	MM	VR
	0.1754	0.9159	2.0424	<b>0.7928</b>
Baseline [17]	0.1569	0.9330	1.9687	0.7842
	<b>0.1408</b>	<b>0.9463</b>	<b>1.9210</b>	0.7840

**Table 2:** Performance of volume and  $RT_{60}$  estimation across baseline (log-magnitude), phase augmented, and continuity augmented models

case it approaches 0, or is capturing the range of our dataset, in which case it approaches 1. This can be written as  $VR = \frac{\text{var}(\hat{y}_n)}{\text{var}(y_n)}$ .

### 5.3. Training

Each model was trained using MSE on the log of the target parameter using Pytorch's Adam optimizer. Each training run ran for 1000 epochs, used L2 regularization, and implemented an adaptive learning rate to ensure efficient convergence. We used grid search to tune hyperparameters including batch size, initial learning rate, regularization strength. We varied the feature combinations on each model as inputs, but left the optimized hyperparameters the same to ensure consistency across model configurations, so that variations in performance would be due to the variations in feature representation.

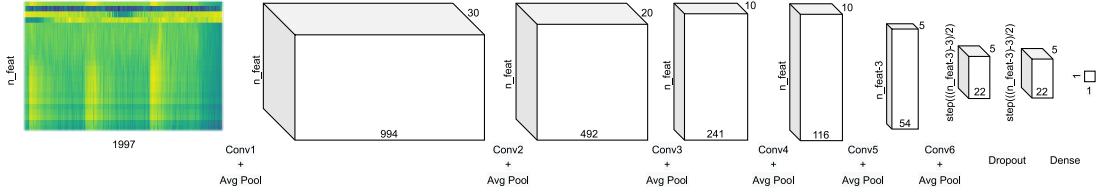
## 6. RESULTS

### 6.1. Room parameter estimation using phase-related features

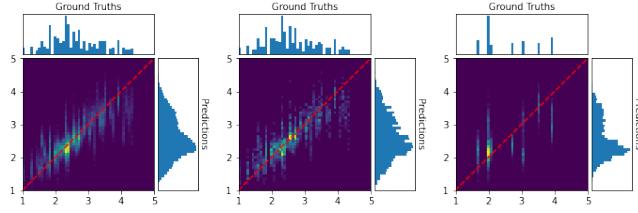
To assess the effectiveness of phase-related features, we trained several models containing a variety of feature combinations. As our baseline, we used a model trained on log-scaled magnitude spectral features for  $RT_{60}$  estimation [17], as well as a set of additional hand-picked features that were proposed in [15] for volume estimation.

Based on our early exploratory results, we built a phase-augmented model (+Phase), taking the same log-scaled magnitude spectral coefficients, and concatenating them with the phase spectrogram features and the first order derivative of the phase, as described in Sec. 4.2. To limit the dimensionality of our input features, these phase features were truncated to only include the bands corresponding to frequencies at, or below 500Hz, as these lower frequencies typically carry more information corresponding to the acoustic parameters of interest. We found that this model best leveraged the additional phase information provided by phase spectrograms and it's first-order derivatives while also balancing model complexity for volume estimation. Similarly, we also built an additional model (+Continuity) by adding the continuity features described in [27] to our +Phase model limited to sub-500Hz frequencies. Each of these models were trained and evaluated on the same data splits, and contained the same model architecture other than the input layer.

Table 2 shows the performance comparison between investigated models. As it can be seen, in general, the two phase-relevant models feature clear performance benefits over the baseline, despite being slightly more complex due to a higher input dimensionality. In terms of volume estimation, +Phase model demonstrates the best performance, achieving the lowest MSE as well as the highest Pearson correlation coefficient. The inclusion of additional phase-continuity information in the +Continuity does not lead to improved



**Fig. 2:** A visualization of our CNN architecture from our featurization. Note that the height of each layer is dependent on the dimensionality of the input feature, which varies from experiment to experiment



**Fig. 3:** Confusion Matrices of our volume estimation model using *+Phase* features on our train, validation, and testing dataset splits (from left to right). The dashed red line indicates a perfect prediction.

performance for the volume estimation task, perhaps due to the additional parameterization requiring more data and processing power to fully take advantage of. However, the additional phase continuity features do provide improved performance over the baseline model as well as the basic phase-augmented model in the estimation of  $RT_{60}$ , which confirms that the tracking of phase variations across time better characterizes  $RT_{60}$  behaviors.

Fig. 3 visualizes the confusion matrices of the proposed volume estimation model for the training, the validation and the test set with x and y axis indicating index numbers of base-10 logarithm values of volume sizes. The volume estimates are well distributed around the ground truth across the tested range, indicating that the CNN successfully learned a representation of the underlying regression problem. Note that the variance is generally smallest in the  $100 m^3$  to  $1000 m^3$  range of estimations, where the bulk of our data lies (including both real and simulated RIRs as shown in Fig. 1); estimation variance is more challenging outside this range due to both a lack of data, as well as the diverging effects of reverb in small acoustic spaces (which are relatively minimal) and large ones (which can be exceptionally hard to process due to long  $RT_{60}$  values).

It's important to note that the model architecture is tuned to support the baseline featurization as is done in [15]; this additionally complex featurization could be better leveraged by a deeper model with wider filters, or a multi-channel featurization as shown in sound event detection literature [28].

## 6.2. Joint estimation of room parameters

The purpose of this section is to investigate the feasibility of jointly estimating multiple room parameters with similar performance to training separate parallel models. The assumption is that parameters such as  $RT_{60}$  and volume draw reverberation information from similar frequency bands and acoustic characteristics, so that we can optimize the simultaneous estimation of all parameters jointly by training a single model. Note that we limited to  $RT_{60}$  and volume due to the lack of ground truth for other room parameters (e.g. total surface area and mean surface absorption coefficients) in the considered real-world RIR datasets.

The architecture remained largely the same as the *+Phase* and *+Continuity* in the prior section, except we modified the final output

Pearson correlation coefficient ( $\rho$ )		
Method	$RT_{60}$	Volume
Baseline [16]	-	0.5736
Baseline [18]	0.9159	-
<i>+Phase</i> w/joint	0.9304	0.6090
<i>+Continuity</i> w/joint	<b>0.9445</b>	<b>0.6195</b>

**Table 3:** Pearson correlation coefficients of parameter estimates using various systems.

layers in both models to fully-connected layers that output two parameters (instead of only one). To avoid issues due to significant scale differences between room volume in  $m^3$  and  $RT_{60}$  in seconds, both (log-scaled) parameters were normalized such that the maximum value in each dataset was 1; we then trained the models on the sum of the normalized MSE for each parameter to ensure equal weighting. Table 3 demonstrates correlation coefficient scores for different systems. Note that other evaluation metrics are not selected for the purpose of consistency due to that parameters are scaled differently in the normalization process. Compared to the results in Table 2, we can observe very similar performance between the joint estimation systems and the same architectures that estimate each of these quantities individually with minimum performance loss. In particular, the joint system with the feature set of *+Continuity* provides the most stable prediction when estimating volume and  $RT_{60}$  simultaneously.

## 7. CONCLUSION

In this paper, we propose improving the task of room parameter estimation by utilizing phase-related features, including phase spectrograms, numerical derivatives of phase, and phase continuity. We demonstrate the benefit of these features in the context of blind parameter estimation from single-channel noisy speech recordings using a CNN architecture. We assemble a dataset from a variety of sources, including publicly available RIR datasets, acoustic room simulations, and in-house RIR measurements. Using the proposed phase-related features, we're able to demonstrate clear improvements over previous state-of-the-art methods in terms of reverberation fingerprint estimation on unseen real-world rooms. Furthermore, we demonstrate the effectiveness of the proposed method in a joint-estimation configuration, showing that sharing parameters improves stability over single-parameter models.

The use of phase-related features for parameter estimation should potentially be extended to multi-band  $RT_{60}$ , absorption coefficients, or other acoustic parameters, and is an area for future work. Further opportunity lies in additional, more-complex phase-related featurizations; recent work in sound event detection demonstrate benefits from multichannel approaches to spectral featurizations [28] allowing for cross-feature representations through 3-D convolution, which could prove beneficial in this task.

## 8. REFERENCES

- [1] Zixing Zhang, Jürgen Geiger, Jouni Pohjalainen, Amr El-Desoky Mousa, Wenyu Jin, and Björn Schuller, “Deep learning for environmentally robust speech recognition: An overview of recent developments,” *ACM Trans. Intell. Syst. Technol.*, vol. 9, no. 5, apr 2018.
- [2] Bo Wu, Kehuang Li, Fengpei Ge, Zhen Huang, Minglei Yang, Sabato Marco Siniscalchi, and Chin-Hui Lee, “An end-to-end deep learning approach to simultaneous speech dereverberation and acoustic modeling for robust speech recognition,” *IEEE Journal of Selected Topics in Signal Processing*, vol. 11, no. 8, pp. 1289–1300, 2017.
- [3] Nasser Mohammadiha and Simon Doclo, “Speech dereverberation using non-negative convolutive transfer function and spectro-temporal modeling,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 24, no. 2, pp. 276–289, 2016.
- [4] Stefania Cecchi, Alberto Carini, and Sascha Spors, “Room response equalization – a review,” *Applied Sciences*, vol. 8, no. 1, 2018.
- [5] Wenyu Jin, “Adaptive reverberation cancelation for multizone soundfield reproduction using sparse methods,” in *2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2016, pp. 509–513.
- [6] Wenyu Jin, Patrick McPherson, Chris Pike, and Adib Mehrabi, “Acoustic room compensation using local PCA-based room average power response estimation,” in *2022 17th International Workshop on Acoustic Signal Enhancement (IWAENC)*, 2022, pp. 166–170.
- [7] Jean-Marc Jot and Keun Sup Lee, “Augmented reality head-phone environment rendering,” in *Audio Engineering Society Conference: 2016 AES International Conference on Audio for Virtual and Augmented Reality*. Audio Engineering Society, 2016.
- [8] Heinric Kuttruff, “Room acoustics,” in *CRC Press*, 2016.
- [9] James Eaton, Nikolay D. Gaubitch, Alastair H. Moore, and Patrick 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.
- [10] Thiago de M. Prego, Amaro A. de Lima, Rafael Zambrano-López, and Sergio L. Netto, “Blind estimators for reverberation time and direct-to-reverberant energy ratio using subband speech decomposition,” in *2015 IEEE Workshop on Applications of Signal Processing to Audio and Acoustics (WASPAA)*, 2015, pp. 1–5.
- [11] Heinrich Loellmann, Andreas Brendel, Peter Vary, and Walter Kellermann, “Single-channel maximum-likelihood T60 estimation exploiting subband information,” *arXiv preprint arXiv:1511.04063*, 2015.
- [12] A. Moore, M. Brookes, and P. Naylor, “Room identification using roomprints,” *Journal of the Audio Engineering Society*, june 2014.
- [13] Nils Peters, Howard Lei, and Gerald Friedland, “Name that room: Room identification using acoustic features in a recording,” in *Proceedings of the 20th ACM International Conference on Multimedia*. 2012, p. 841–844, Association for Computing Machinery.
- [14] Prateek Murgai, Mark Rau, and Jean-marc Jot, “Blind estimation of the reverberation fingerprint of unknown acoustic environments,” *Journal of the Audio Engineering Society*, October 2017.
- [15] Andrea F. Genovese, Hannes Gamper, Ville Pulkki, Nikunj Raghuvanshi, and Ivan J. Tashev, “Blind room volume estimation from single-channel noisy speech,” in *ICASSP 2019 - 2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2019, pp. 231–235.
- [16] Feifei Xiong, Stefan Goetze, and Bernd T Meyer, “Joint estimation of reverberation time and direct-to-reverberation ratio from speech using auditory-inspired features,” *arXiv preprint arXiv:1510.04620*, 2015.
- [17] Hannes Gamper and Ivan J. Tashev, “Blind reverberation time estimation using a convolutional neural network,” in *2018 16th International Workshop on Acoustic Signal Enhancement (IWAENC)*, 2018, pp. 136–140.
- [18] Nicholas J. Bryan, “Impulse response data augmentation and deep neural networks for blind room acoustic parameter estimation,” in *ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2020, pp. 1–5.
- [19] Philipp Götz, Cagdas Tuna, Andreas Walther, and Emanuël A. P. Habets, “Blind reverberation time estimation in dynamic acoustic conditions,” in *ICASSP 2022 - 2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2022, pp. 581–585.
- [20] Prerak Srivastava, Antoine Deleforge, and Emmanuel Vincent, “Blind room parameter estimation using multiple multichannel speech recordings,” in *2021 IEEE Workshop on Applications of Signal Processing to Audio and Acoustics (WASPAA)*, 2021, pp. 226–230.
- [21] M. R. Schroeder, “New method of measuring reverberation time,” *The Journal of the Acoustical Society of America*, vol. 37, no. 3, pp. 409–412, 1965.
- [22] Marco Jeub, Magnus Schafer, and Peter Vary, “A binaural room impulse response database for the evaluation of dereverberation algorithms,” in *2009 16th International Conference on Digital Signal Processing*, 2009, pp. 1–5.
- [23] Igor Szöke, Miroslav Skácel, Ladislav Mošner, Jakub Paliesek, and Jan Černocký, “Building and evaluation of a real room impulse response dataset,” *IEEE Journal of Selected Topics in Signal Processing*, vol. 13, no. 4, pp. 863–876, 2019.
- [24] D. Murphy and S. Shelley, “Openair: an interactive auralization web resource and database,” *Journal of the Audio Engineering Society*, November 2010.
- [25] Robin Scheibler, Eric Bezzam, and Ivan Dokmanic, “Pyroomacoustics: A python package for audio room simulations and array processing algorithms,” *CoRR*, vol. abs/1710.04196, 2017.
- [26] Prerak Srivastava, Antoine Deleforge, and Emmanuel Vincent, “Realistic sources, receivers and walls improve the generalisability of virtually-supervised blind acoustic parameter estimators,” *arXiv preprint arXiv:2207.09133*, 2022.
- [27] Doyeon Kim, Hyewon Han, Hyeon-Kyeong Shin, Soo-Whan Chung, and Hong-Goo Kang, “Phase continuity: Learning derivatives of phase spectrum for speech enhancement,” 2022.
- [28] Christopher Ick and Brian McFee, “Sound event detection in urban audio with single and multi-rate PCEN,” 2021.