

StoRIR: Stochastic Room Impulse Response Generation for Audio Data Augmentation

Piotr Masztalski, Mateusz Matuszewski, Karol Piaskowski, Michał Romaniuk

Samsung R&D Institute Poland

{p.masztalski, m.matuszewski, k.piaskowski, m.romaniuk2}@samsung.com

Abstract

In this paper we introduce StoRIR - a stochastic room impulse response generation method dedicated to audio data augmentation in machine learning applications. This technique, in contrary to geometrical methods like image-source or ray tracing, does not require prior definition of room geometry, absorption coefficients or microphone and source placement and is dependent solely on the acoustic parameters of the room. The method is intuitive, easy to implement and allows to generate RIRs of very complicated enclosures. We show that StoRIR, when used for audio data augmentation in a speech enhancement task, allows deep learning models to achieve better results on a wide range of metrics than when using the conventional image-source method, effectively improving many of them by more than 5%. We publish a Python implementation of StoRIR online¹.

Index Terms: room impulse response, data augmentation, speech enhancement, deep learning

1. Introduction

Deep learning tasks, like speech enhancement or acoustic event classification and localization, can significantly benefit from augmenting training datasets with room impulse responses (RIR) [1, 2, 3]. RIRs are a set of functions that describe the influence of a given acoustic environment when a sound wave propagates from a source to a receiver. Dataset augmentation with RIRs often leads to better generalization in real-life scenarios, where the acoustic environment has a large impact on the recorded audio signal. However, capturing real-life impulse responses requires a lot of time and resources and the openly available databases containing recorded RIRs can be insufficient for proper data augmentation. To address that issue, researchers have used computational methods that simulate RIRs with a lot of success [4, 5].

2. Background and related work

There are several methods to calculate a RIR. The approach that is able to provide the most accurate results is based on numerically solving the wave equation (e.g. finite element method, boundary element method). However, these techniques are computationally very expensive and require a detailed mesh of the acoustic environment. Faster, but less accurate techniques are those based on the assumptions of geometrical acoustics (e.g. ray tracing, image-source) [6]. In geometrical acoustics, sound is assumed to propagate as rays, and all of its wave properties are neglected. This assumption is valid at high frequencies, where the wavelength of sound is short compared to surface dimensions and the overall dimensions of the space, but

at low frequencies the approximation errors increase as wave phenomena play a larger role [6]. It is important to note, that for audio data augmentation in machine learning applications, the method of choice while simulating RIRs is almost exclusively image-source. We can attribute its popularity in this field to low computation complexity, that allows for online generation during training, and availability of open source software designed to model room acoustics using this method [7, 8]

Accuracy similar to that achieved by employing numerical methods in acoustics can be achievable in a much simpler way. Statistical Room Acoustics (SRA), under certain assumptions, can provide a statistical description of a RIR between a source and a receiver. Sabine [9] presented the earliest attempt at room statistical methods. He introduced, in the late 19th century, a method for reverberation time calculation of a space without considering the details of its geometry. More than 50 years later, Schroeder [10] extended Sabine's fundamental work and derived a set of statistical properties describing the frequency spectrum of a random impulse response. Moorer et al. [11] noted, that impulse responses in the finest concert halls around the world sounded remarkably similar to white noise with an exponential amplitude envelope. To test this observation, they generated synthetic impulse responses by shaping unit-variance Gaussian pseudo-random sequences with an exponential of the desired length. The direct sound was added by including an impulse at the beginning. Later, Polack [12] developed a time-domain model excluding the contribution of the direct path and describing a RIR as a realization of a non-stationary stochastic process.

In this paper, we present a method of calculating RIRs based on SRA dedicated to machine learning data augmentation. To the knowledge of the authors this is the first application of SRA to machine learning tasks. The presented method is particularly useful for data augmentation, because it does not require defining room geometry and acoustic properties of walls, furniture, etc. which is necessary in other methods. Arguably, the fact that we do not define room geometry is a drawback for other purposes like auralization where the RIR of a specific room is needed. However, the purpose of data augmentation is to help machine learning models to generalize, so that they work in diverse acoustic environments. Therefore, using StoRIR, we generate RIRs based on several acoustic parameters that correspond to a class of rooms, without explicitly defining geometry or absorption coefficients.

3. Proposed method

3.1. Room Acoustic Parameters

There exist multiple parameters that describe acoustic properties of a room [13]. For the purpose of implementing StoRIR we use the following:

¹<https://github.com/SRPOL-AUI/storir>

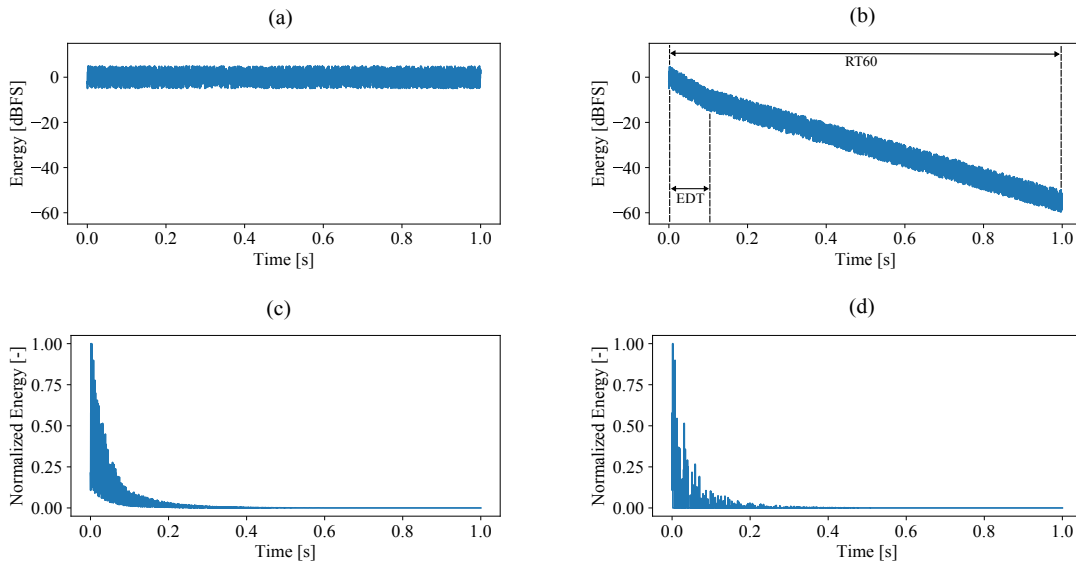


Figure 1: Four steps of stochastic RIR generation. Noise generation - (a). Energy decay curve shaping - (b). Conversion from logarithmic to linear scale - (c). Adjusting energy distribution within the RIR - (d).

- Reverberation time (RT60) - a well-known, and probably the most common, parameter describing the acoustic behaviour of an acoustic cavity. RT60 quantifies the time it takes for the sound to decay by 60 dB after the sound source is removed. This term was introduced in Sabine's pioneering research, in which he noticed that reverberation time was proportional to the volume of the room and inversely proportional to the amount of absorption. Because the absorptive properties of materials vary as a function of frequency, the reverberation time is also frequency dependent,
- Early Decay Time (EDT) - a very similar parameter to RT60. The difference relies in the amount of energy decay. EDT denotes how quickly sound energy drops by 10 dB, not 60 dB like in RT60. This is due to the fact that early reflections may decay with a different rate than the late part of the RIR,
- Direct to Reverberant Ratio (DRR) - describes the energy ratio of the direct sound to the reverberant part of a signal,
- Initial Time Delay Gap (ITD) - specifies the time between the arrival of the direct sound and the first reflection to the receiver.

3.2. Stochastic impulse response generation

The result of a geometric RIR generation method is an energetic impulse response (also called a reflectogram). We aim to simulate a similar outcome, but without providing the algorithm with room geometry, absorption coefficients and microphone/source placement. Hence, the method will not model a specific room but rather the impulse response itself. Normally, room acoustic parameters are obtained from a RIR. In our approach we reverse this property and generate the RIR from scratch by shaping it based on the acoustic parameter values. For the input to our algorithm we choose to consider RT60, EDT, DRR and ITDG described in detail in section 3.1. We argue, that using these parameters can be more intuitive when generating a large amount of

RIRs that should not necessarily correspond to a specific room geometry. This is often the case when augmenting data for machine learning purposes where we want models to generalize in a diverse range of acoustic environments.

The proposed stochastic RIR generation algorithm can be split into four steps: noise generation, energy decay curve shaping, conversion from logarithmic to linear scale and adjusting the energy distribution within the RIR (Figure 1).

3.2.1. Noise generation

We start constructing the RIR with generating an uniformly distributed noise vector v of length l equal to the RT60 parameter in samples. Both EDT and RT60 parameters describe the sound energy decay in decibels so, for convenience, we use the logarithmic scale in the beginning of RIR generation. Due to that, the range of the v 's uniform distribution will become the deviation range of reflection energies (in dB) in the resulting energetic impulse response. This range can be arbitrarily chosen or sampled from a distribution adding to the randomness of the generated RIR. We assume that the mean value of v is 0 [dB].

3.2.2. Energy decay curve shaping

In order to obtain a simulated energy decay curve (EDC), we introduce a negative slope to v so that its mean value drops by 10 dB in time specified by the EDT parameter using

$$\forall i \in 1, 2 \dots k : EDC_i = v_i - \frac{10i}{k} \quad (1)$$

where k is the EDT parameter value in samples.

After that, we concatenate the remainder of v (decreased in level by 10 dB) with the EDC vector and decrease it gradually by further 50 dB so that the overall level of EDC drops by 60 dB in time specified by the RT60 parameter. We do it using the following equation:

$$\forall i \in k + 1, k + 2, \dots, l : \\ EDC_i = v_i - \left[10 + \frac{50(i - k)}{l} \right] \quad (2)$$

3.2.3. Conversion from logarithmic to linear scale

The third step is converting EDC to linear scale using

$$EDC^{lin} = 10^{\frac{EDC}{10}}. \quad (3)$$

After that, we normalize the resulting signal to its peak value. EDC^{lin} can be interpreted as a maximally dense energetic impulse response, meaning energy is detected in every time slot (sample). We will refer to these energy peaks as reflections or rays.

3.2.4. Adjusting energy distribution

In order to account for the fact that, in real life, sound reflections reach the receiver with time gaps in between them, we make the obtained energetic impulse response EDC^{lin} more sparse. We simulate delays between reflections by first eliminating a chunk of reflections of length equal to the ITDG parameter in samples after the initial sound ray. Then, we iteratively delete rays at random locations from the remaining part of EDC^{lin} until we reach a desired DRR, hence the stochastic nature of the generated RIR. Analysis of sound propagation in enclosures shows, that in the early part of the RIR reflections are less densely distributed than in the reverberation tail [13], therefore we assign a higher probability of deleting rays from the early reflection part of the RIR (EDC^{lin}) than from its remainder.

A RIR generated in this manner does not represent any existing or modeled room but rather some imaginary room with desired acoustic parameters. Due to the randomness introduced in generating the energy decay curve from random noise, and the process of adjusting the energy ratios, the resulting RIR will differ with every generation even if we do not change the values of parameters it is based on. The more variability we allow in the acoustic parameters, the more the generated impulse responses will differ from each other. We can e.g. make a whole set of impulse responses where $RT60 = 1$ s, but with other parameter values sampled from chosen ranges, therefore creating an augmentation dataset of RIRs representing a class of virtual rooms with 1 second reverberation time.

3.3. Other representations of the stochastic impulse response

After executing all of the steps described above, we obtain an energetic impulse response that can be directly convolved with an arbitrary audio signal resulting in a reverberation effect. However, the reverberation time of the space represented by this RIR is equal across the whole frequency spectrum which, most of the time, is not the case with real impulse responses due to stronger air and surface attenuation of sound in higher frequencies [13]. To account for that, the RIR can be generated separately for each 1/1 or 1/3 octave band with decreasing reverberation time. Another modification can be to generate the RIR with lower time resolution (e.g. 5 ms) and then adapt it to the sampling frequency using the Poisson-distributed noise method described in [14] (however, it should be noted that this method requires an estimation of the room volume). In our experiments we decided to use the most basic version of the energetic impulse response without any of the above described modifica-

tions, as we have experimentally established that it produces the best results in our evaluation task.

4. Experimental evaluation

4.1. Model description

Although our method can be employed in virtually any task that can benefit from RIR augmentation, we decided to evaluate it on a speech enhancement task. A large amount of research has been conducted in the field of noise suppression in anechoic conditions [15, 16, 17, 18]. However, these experimental scenarios are not applicable to real-life use cases, as the acoustic characteristics of the space where the sound sources are placed cause a significant difference to the spectral structure of resulting audio signals. In order to account for this difference, it is common to convolve noisy audio with various RIRs so that the model can suppress noise and reverberation at once. To evaluate the proposed augmentation method we use one deep learning model designed for denoising [16], and one designed for dereverberation [19] which we call DenoiseNet and DereverbNet. Both are based on the U-Net [20] architecture with DenoiseNet using complex-valued operations, and DereverbNet employing DenseNet-like blocks [21] and two LSTM layers in the bottleneck. Both models are used for the exact same task of simultaneous speech denoising and dereverberation.

4.2. Dataset

4.2.1. Training

For all experiments, we use the clean and noisy parallel speech database by Valentini et al. [22] which consists of around 400 utterances from each of 14 male and 14 female speakers. The authors obtained the training dataset by mixing speech files from the Voice Bank corpus [23] with 10 different types of noise (2 artificial, and 8 from the DEMAND database [24]) at four different SNRs (15 dB, 10 dB, 5 dB and 0 dB). To evaluate the performance of our RIR generation method, we compare it with the image-source method. We generate 50,000 impulse responses with the image method using Pyroomacoustics [8], an open source Python package for room acoustics simulation. The details of simulated room geometries and placement of sound sources and microphones are the same as described in [19]. The one difference is, we change the desired reverberation time range to 0.2 s - 0.7 s (it has to be converted to absorption coefficients using the Sabine's formula for the purpose of using the image-source method). For comparison, we also generate 50,000 equivalent RIRs with the proposed method setting the RT60 range to 0.2 s - 0.7 s, the EDT range to 50 ms - 100 ms, the ITDG range to 3 ms - 10 ms and the DRR range to -7 dB - 0 dB. During training the generated RIRs are randomly convolved with noisy speech utterances to obtain noisy and reverberant signals. It is important to note, that during training, only simulated RIRs were used without adding any real-life recorded ones.

4.2.2. Testing

The performance of each model is evaluated on two test sets. The first one (Testset 1) is an unchanged version of the noisy and reverberant test set proposed by Valentini [25]. Because of the fact that this dataset covers only 3 room configurations (two of which are recorded in a somewhat artificial setting using a room with configurable reverberation time), we create a second test set (Testset 2) utilizing the same underlying clean and

Table 1: Results on Testset 1 (higher is better, bold text indicates best score per model within a given metric)

Model	Aug. method	CSIG	CBAK	COVL	fwSegSNR	STOI	PESQ
Noisy reverberant speech	-	3.00	1.80	2.25	7.1	0.62	1.91
DenoiseNet	No augmentation	2.87	1.87	2.20	7.1	0.63	2.09
	Image-Source	3.63	2.07	2.73	9.2	0.67	2.40
	StoRIR	3.66	2.08	2.75	9.3	0.69	2.41
DereverbNet	No augmentation	3.69	2.26	2.98	10.2	0.64	2.07
	Image-Source	3.97	2.45	3.26	11.3	0.76	2.34
	StoRIR	3.95	2.35	3.20	11.2	0.70	2.45

Table 2: Results on Testset 2 (higher is better, bold text indicates best score per model within a given metric)

Model	Aug. method	CSIG	CBAK	COVL	fwSegSNR	STOI	PESQ
Noisy reverberant speech	-	3.25	1.87	2.44	8.1	0.77	2.01
DenoiseNet	No augmentation	3.17	1.96	2.42	7.9	0.74	2.08
	Image-Source	3.87	2.08	2.89	9.9	0.83	2.53
	StoRIR	4.02	2.23	3.05	10.6	0.84	2.59
DereverbNet	No augmentation	3.69	2.29	2.98	10.0	0.77	2.03
	Image-Source	4.20	2.46	3.40	11.7	0.81	2.55
	StoRIR	4.34	2.51	3.60	12.6	0.86	2.62

noisy files, but convolved randomly with 12 RIRs from 6 different rooms (two offices, two meeting rooms, a lecture room and a building lobby) in order to further evaluate generalization in real-life settings and more complex room geometries. The RIRs used in Testset 2 are obtained from the ACE challenge corpus [26] and their reverberation time ranges from 0.3 s to 0.75 s. All RIRs in both test sets are real-life recordings.

4.3. Data preprocessing

The original audio signals were first downsampled from 48kHz to 16 kHz (which is a common practice in speech enhancement tasks) and then, for the final model input, complex-valued spectrograms were obtained from raw waveforms. The STFT was computed with a 64 ms window and a 16 ms hop size for DenoiseNet, and a 32 ms window and 8 ms hop size for DereverbNet as stated in the original papers. The training target for all experiments was clean anechoic speech.

4.4. Evaluation metrics

To evaluate the performance of the speech enhancement models, and hence the viability of data augmentation using the proposed method, we choose six objective speech quality and intelligibility metrics: CSIG - mean opinion score (MOS) predictor of signal distortion, CBAK - MOS predictor of background-noise intrusiveness, COVL - MOS predictor of overall signal quality, fwSegSNR - frequency-weighted segmental signal to noise ratio, STOI - short time objective intelligibility and PESQ - perceptual evaluation of speech quality. For calculating metrics we used code repositories available online ^{2 3 4}.

²<https://github.com/IMLHF/Speech-Enhancement-Measures>

³<https://github.com/mpariente/pystoi>

⁴<https://github.com/vBaiCai/python-pesq>

5. Results and discussion

Tables 1 and 2 summarize the metric scores on both test sets. It is worth to mention the substantial boost of performance caused by both RIR augmentation techniques when comparing with no augmentation. In some cases, models trained with no augmentation achieve even worse scores than the original noisy and reverberant mixture. Both RIR generation methods lead to improved speech quality metrics, however, when comparing the two, using StoRIR produces better results in more cases overall. The scores are improved on Testset 2 across all metrics, which we see resulting from better representation of unconventional room geometries (like building lobby) by our RIR generation method. The 0.9 dB improvement in fwSegSNR on DereverbNet seems to be the most impressive result on this testset. More comparable results on Testset 1 may result from the fact that two out of three RIRs used in this test set are recorded in a configurable reverberation room with a perfect ShoeBox geometry, identical to the ones modeled with the image-source method. Still, even though the image-source method has a somewhat unfair advantage in this test scenario, StoRIR manages to provide similar or better results (especially on DereverbNet where the PESQ improvement reached 0.11). The presented results also show that the performance of our method is model agnostic.

6. Conclusions

In this paper we proposed StoRIR, a method for room impulse response generation that does not require any information about room geometry, absorption coefficients or microphone and sound source placement. We show that when used for data augmentation, our method substantially improves the performance of deep learning models employed for a speech enhancement task in reverberant conditions. When compared with the image-source method for RIR generation, StoRIR achieves superior results when dealing with real reverberation in a wide range of acoustic environments.

7. References

- [1] K. Kinoshita, M. Delcroix, S. Gannot, E. Habets, R. Haeb-Umbach, W. Kellermann, V. Leutnant, R. Maas, T. Nakatani, B. Raj, A. Sehr, and T. Yoshioka, "A summary of the reverb challenge: State-of-the-art and remaining challenges in reverberant speech processing research," *Journal on Advances in Signal Processing*, vol. 2016, 12 2016.
- [2] S. Adavanne, A. Politis, J. Nikunen, and T. Virtanen, "Sound event localization and detection of overlapping sources using convolutional recurrent neural networks," *IEEE Journal of Selected Topics in Signal Processing*, vol. 13, no. 1, p. 34–48, Mar 2019. [Online]. Available: <http://dx.doi.org/10.1109/JSTSP.2018.2885636>
- [3] D. Emmanouilidou and H. Gamper, "The effect of room acoustics on audio event classification," 09 2019.
- [4] T. Ko, V. Peddinti, D. Povey, M. L. Seltzer, and S. Khudanpur, "A study on data augmentation of reverberant speech for robust speech recognition," in *2017 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2017, pp. 5220–5224.
- [5] R. Hsiao, J. Ma, W. Hartmann, M. Karafiát, F. Grézl, L. Burget, I. Szöke, J. H. Černocký, S. Watanabe, Z. Chen, S. H. Mallidi, H. Hermanský, S. Tsakalidis, and R. Schwartz, "Robust speech recognition in unknown reverberant and noisy conditions," in *2015 IEEE Workshop on Automatic Speech Recognition and Understanding (ASRU)*, 2015, pp. 533–538.
- [6] L. Savioja and U. P. Svensson, "Overview of geometrical room acoustic modeling techniques," *The Journal of the Acoustical Society of America*, vol. 138, 2015.
- [7] E. Habets, "RIR-Generator." [Online]. Available: <https://github.com/ehabets/RIR-Generator>
- [8] R. Scheibler, E. Bezzam, and I. Dokmanic, "Pyroomacoustics: A python package for audio room simulation and array processing algorithms," *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, Apr 2018. [Online]. Available: <http://dx.doi.org/10.1109/ICASSP.2018.8461310>
- [9] W. C. Sabine, *Collected Papers On Acoustics*. Cambridge: Harvard University Press, 1923.
- [10] M. R. Schroeder, "Frequency-correlation functions of frequency responses in rooms," *The Journal of the Acoustical Society of America*, vol. 34, no. 12, pp. 1819–1823, 1962.
- [11] J. A. Moorer, "About this reverberation business," *Computer music journal*, pp. 13–28, 1979.
- [12] J.-D. Polack, "La transmission de l'énergie sonore dans les salles," Ph.D. dissertation, Le Mans, 1988.
- [13] H. Kuttruff, *Room Acoustics, Fourth Edition*, ser. E-Libro. Taylor & Francis, 2000.
- [14] D. Schröder, "Physically based real-time auralization of interactive virtual environments," Ph.D. dissertation, RWTH Aachen University, 01 2011.
- [15] S. Pascual, A. Bonafonte, and J. Serrà, "SEGAN: Speech enhancement generative adversarial network," 2017. [Online]. Available: <https://arxiv.org/abs/1703.09452>
- [16] H.-S. Choi, J. Kim, J. Huh, A. Kim, J.-W. Ha, and K. Lee, "Phase-aware speech enhancement with deep complex U-Net," in *International Conference on Learning Representations*, 2019. [Online]. Available: <https://openreview.net/forum?id=SkeRTsAcYm>
- [17] C. Macartney and T. Weyde, "Improved speech enhancement with the Wave-U-Net," 2018. [Online]. Available: <https://arxiv.org/abs/1811.11307>
- [18] S. R. Park and J. Lee, "A fully convolutional neural network for speech enhancement," 2016. [Online]. Available: <https://arxiv.org/abs/1609.07132>
- [19] Z. Wang and D. Wang, "Deep learning based target cancellation for speech dereverberation," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 28, pp. 941–950, 2020.
- [20] O. Ronneberger, P. Fischer, and T. Brox, "U-net: Convolutional networks for biomedical image segmentation," in *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015*, N. Navab, J. Hornegger, W. M. Wells, and A. F. Frangi, Eds. Cham: Springer International Publishing, 2015, pp. 234–241.
- [21] G. Huang, Z. Liu, L. Van Der Maaten, and K. Q. Weinberger, "Densely connected convolutional networks," in *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017, pp. 2261–2269.
- [22] C. Valentini Botinhao, X. Wang, S. Takaki, and J. Yamagishi, "Speech enhancement for a noise-robust text-to-speech synthesis system using deep recurrent neural networks," in *Proceedings of Interspeech 2016*, 9 2016, pp. 352–356.
- [23] C. Veaux, J. Yamagishi, and S. King, "The voice bank corpus: Design, collection and data analysis of a large regional accent speech database," in *2013 International Conference COCODA held jointly with 2013 Conference on Asian Spoken Language Research and Evaluation (O-COCODA/CASLRE)*, 2013, pp. 1–4.
- [24] J. Thiemann, N. Ito, and E. Vincent, "The diverse environments multi-channel acoustic noise database (demand): A database of multichannel environmental noise recordings," *The Journal of the Acoustical Society of America*, vol. 133, p. 3591, 05 2013.
- [25] C. Valentini-Botinhao, "Noisy reverberant speech database for training speech enhancement algorithms and TTS models," 2017. [Online]. Available: <https://doi.org/10.7488/ds/2139>
- [26] J. Eaton, N. D. Gaubitch, A. H. Moore, and P. A. Naylor, "The ACE challenge — corpus description and performance evaluation," in *2015 IEEE Workshop on Applications of Signal Processing to Audio and Acoustics (WASPAA)*, 2015, pp. 1–5.