

# EVALUATION OF DEEP-LEARNING-BASED VOICE ACTIVITY DETECTORS AND ROOM IMPULSE RESPONSE MODELS IN REVERBERANT ENVIRONMENTS

*Amir Ivry    Israel Cohen    Baruch Berdugo*

Andrew and Erna Viterbi Faculty of Electrical Engineering  
Technion – Israel Institute of Technology, Technion City, Haifa 3200003, Israel

## ABSTRACT

State-of-the-art deep-learning-based voice activity detectors (VADs) are often trained with anechoic data. However, real acoustic environments are generally reverberant, which causes the performance to significantly deteriorate. To mitigate this mismatch between training data and real data, we simulate an augmented training set that contains nearly five million utterances. This extension comprises of anechoic utterances and their reverberant modifications, generated by convolutions of the anechoic utterances with a variety of room impulse responses (RIRs). We consider five different models to generate RIRs, and five different VADs that are trained with the augmented training set. We test all trained systems in three different real reverberant environments. Experimental results show 20% increase on average in accuracy, precision and recall for all detectors and response models, compared to anechoic training. Furthermore, one of the RIR models consistently yields better performance than the other models, for all the tested VADs. Additionally, one of the VADs consistently outperformed the other VADs in all experiments.

*Index Terms*— Voice activity detection, reverberation, room impulse response, deep learning

## 1. INTRODUCTION

Voice activity detection (VAD) aims to determine the boundaries in which speech exists in an observed audio signal. State-of-the-art deep-learning-based VADs are often trained with anechoic data. However, real-life acoustic environments are reverberant, which deteriorates VAD performance in practical scenarios. In this study, we mitigate the mismatch between training data and real data by generating an augmented training set that integrates anechoic and reverberant audio signals. The reverberant training corpus is generated by convolving anechoic utterances with a simulated room impulse responses (RIRs). Enhanced VAD in reverberant environments may benefit a variety of audio-based applications such as speech enhancement [1–3], dereverberation [4,5] and speech and speaker recognition [6,7].

Deep-learning-based VADs have attained leading performances during recent years, due to the ability of neural networks to learn non-linear relations and complex patterns of audio signals. To detect voice activity, Ariav and Cohen [8] encoded spectral audio features via an auto-encoder that fed a recurrent neural network. Wagner et al. [9] introduced automatic feature engineering through the convolutional layers of a deep neural network. Leading performance was obtained by Kim and Hahn [10] that integrated an attention model to weight context information into existing deep learning architectures. Combined end-to-end VAD system was introduced by Ariav et al. [11], that comprised of WaveNet for feature extraction and a deep residual network for speech detection. Ivry et al. [12] applied ensemble learning with two deep encoder-decoder structures to learn the unique temporal and spatial patterns of speech through the diffusion maps method.

In latest decades, several RIR models were proposed to produce reverberant utterances via simulations. An extension of the known image method [13] to arbitrary polyhedra was first introduced by Borish [14]. Vorländer [15] suggested a combined modeling that considers both the image method and ray-tracing techniques. Rindel [16] employed reflection coefficients that are incidence angle-dependent in the frequency domain, to offer a more accurate characterization of a room response. A similar model was implemented by Lam [17], but it focused on low frequencies for more realistic boundary conditions. Valeau et al. [18] applied the diffusion equation to predict room acoustics.

We consider the aforementioned five deep-learning-based VADs [8–12] and five RIR models [14–18]. First, we show that training these detectors with solely anechoic corpus and testing them in real reverberant rooms and spaces leads to a significantly impeded detection capability. To include unique acoustic patterns of reverberant data during training, we generated an augmented training set of nearly five million utterances. This extended corpus comprises of anechoic and reverberant signals, where the latter is generated by convolving the anechoic signals with a variety of RIRs, generated using a fixed RIR model. Then, all five VADs are independently trained with this augmented training set. This experiment is repeated for each of the five RIR models. All trained detection systems are tested in three real reverberant spaces of a

---

This work was supported by the Israel Science Foundation (grant no. 576/16) and the ISF-NSFC joint research program (grant No. 2514/17).

classroom, a large concert hall, and an octagon shaped library. Experimental results demonstrate that the performance of all detectors is enhanced in each of the tested reverberant environments, regardless of the RIR model employed during training. Evaluation measures such as accuracy, precision and recall increase by 20% on average, compared to non-reverberant training. An interesting outcome shows that the leading accuracy of each detector was consistently achieved by the Valeau RIR model [18]. In a similar manner, the detector introduced in Ivry [12] prevailed competing VADs across all experiments.

The remainder of this paper is organized as follows. In Section 2, we describe the database. In Section 3, experimental results are given. We conclude in Section 4.

## 2. DATABASE GENERATION

In this section, we detail the construction of two disjoint datasets: An augmented training set and a test set. The training set contains both anechoic and reverberant utterances, that are generated by simulating a fixed RIR model and convolving the anechoic data with it. In contrast, the test set is constructed with real reverberant conditions, not simulations.

For the training stage, we employ the TIMIT [19] training dataset that contains 4620 anechoic utterances, sampled at 16 kHz. Since this corpus is imbalanced and does not comprise of noises, we perform several preprocessing steps. Initially, since in TIMIT there are more speech frames than silence frames, we manually add 2 s of silence for each existing recording in the corpus. Next, we acquire recordings of stationary noises such as white and colored Gaussian noise, musical instruments and babble. These noises are randomly added to both speech and silence frames in signal-to-noise-ratios (SNRs) that are distributed uniformly between 10-20 dB relative to clean anechoic speech.

We perform augmentation of this anechoic training set, so it holds both anechoic and simulated reverberant data. To simulate varied reverberant environments, 50 rectangular spaces are considered, such that the length, width and height are uniformly chosen from the range 3 – 20 m. This permits both small, medium and large spaces. To cover various scenarios, each of the 50 spaces is simulated 20 times, with different locations of the speaker and the receiver. To obtain a realistic setting, the speaker and the microphone are limited to height range of 1 – 2 m, and a distance of at least 0.5 m from each other. Each room is simulated with a reverberation time (RT60) that is chosen uniformly from the interval 0.1 – 1 s, such that both low and high reflective surfaces are accounted for.

Given an RIR model, we simulate  $50 \times 20$  RIR signals. Each of these responses is convolved with the anechoic utterances in [19], which results in a reverberant training set. The augmented training set is simply a composition of the original anechoic signals with their aforementioned reverberant mod-

ifications. Ultimately, for a given RIR model, the training set comprises of  $4620 \times 1001$  utterances.

In the test stage, we use 100 anechoic utterances from the TIMIT test dataset. To obtain the reverberant test set, convolution is applied between this corpus and real recordings of room responses. These RIRs are taken from three reverberant environments [20] of a classroom, a large concert hall and an octagon shaped library. For each environment, 130 recordings are available, from various locations in the room. Thus, three test sets are formed, each comprises of  $100 \times 130$  reverberant utterances.

## 3. EXPERIMENTAL RESULTS

In the following experiments, voice activity detection performance is evaluated by several measures. The receiver operating characteristic (ROC) curve is used to present a trade-off between speech detection and false-alarm rates in various operation points. The robustness of the VAD and the sensitivity of its classifier to noises is derived by the area under curve (AUC) measure. Accuracy, precision, recall and F1-score [21] are also employed in this study. When combined, all measures strongly indicate on the accuracy, generalization and robustness abilities of the detector.

In this study, we consider FIVE VADs [8–12] and address them as Ariav-R, Wagner, Kim, Ariav-W and Ivry, respectively. Also, we employ five RIR models [14–18], and refer them as Borish, Vorlnder, Randel, Lam and Valeaua, correspondingly. We perform the following experiment, comprises of two-stages; training stage and test stage. In the first part, a fixed RIR model is simulated. Then, the steps described in Section 2 are implemented with respect to the chosen RIR model. As a result, an augmented training set is obtained. Next, a VAD system is chosen and trained with the derived training set. We repeat this experiment for each VAD system and for each RIR model. Ultimately, this stage yields  $5 \times 5$  trained VAD systems. In the second stage, we test each trained detector on three test sets, generated in three reverberant environments of a classroom, a large concert hall and an octagon shaped library, as detailed in Section 2. An experiment conducted by the authors of this study showed that these three acoustic spaces are characterized by long (1 s), medium (0.8 s) and short (0.6 s) reverberation time, in correspondence.

By observing Fig. 1, several conclusions can be derived. First and foremost, if the training set contains merely anechoic data, then the performance of all VADs is significantly degraded when tested in real reverberant conditions. Respectively, employing the suggested augmented training set that comprises of reverberant utterances consistently enhances VAD performance in practical scenarios. The reason is that acoustic patterns and features highly differ between reverberant and anechoic environments, and this mismatch between the training data and real data is mitigated by the reverberant augmented training set. Another interesting derivation is that

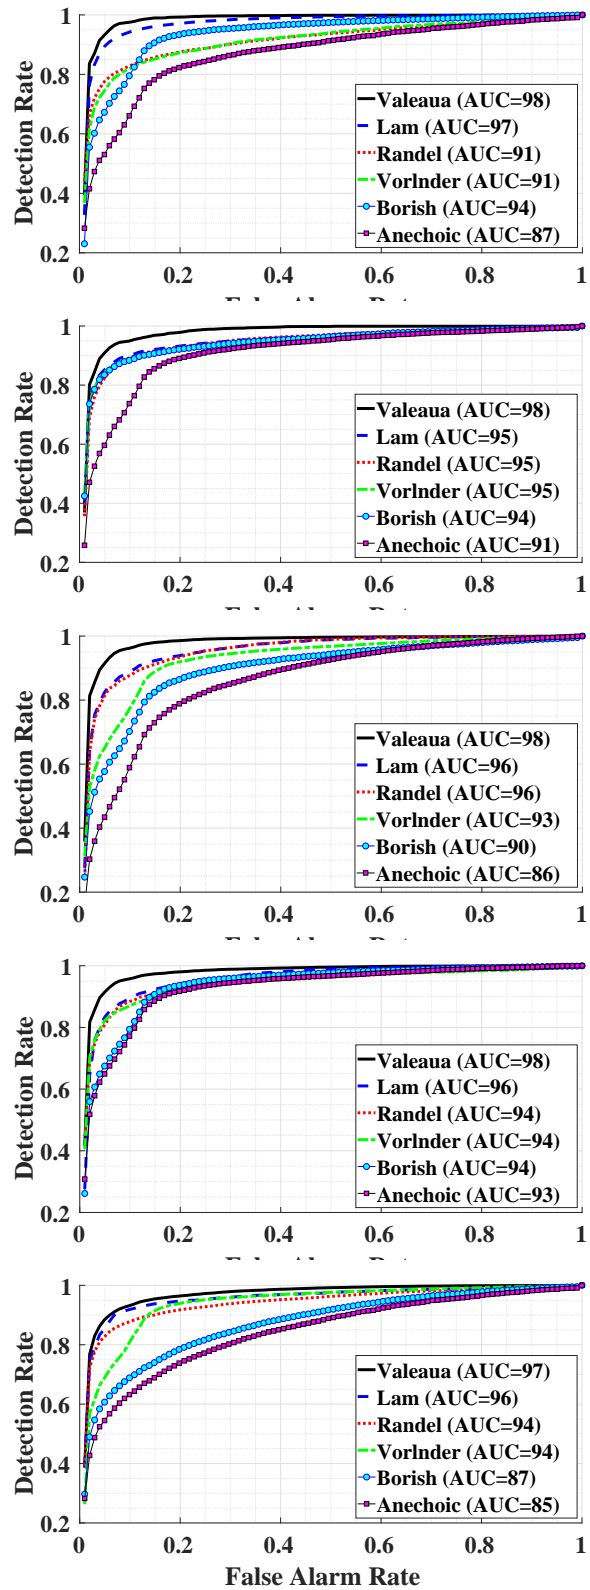
the RIR model introduced by Valeaua [18] consistently leads to the highest performance, relative to competing RIR models, for all VADs and in all tested acoustics. One explanation is that the model proposed in [18] predicts room acoustics better than the remaining models. It should be noticed that training with Valeaua RIR model leads to rapid convergence of the ROC curves and leading AUC values. These results indicate that detectors trained with Valeaua impulse response achieve wide margins of separation between speech and silence. Therefore, these detectors experience high robustness from noises and interferences that might shift the classifier.

Further derivations can be made based on Fig 2. The reported results reaffirm that augmentation of the training set with respect to Valeaua RIR model leads to enhanced VAD performance in reverberant conditions, compared to training that merely considers anechoic data. This enhancement can be quantified by approximately 20% gap across all performance measures of accuracy, precision, recall and F1-score. This conclusion also implies high generalization ability of all VADs that are trained with [18], since they consistently achieve enhanced performance for all measures and in all three acoustic environments. Next, let us focus on the interpretation of the accuracy, precision and recall measures. Since the training and test sets are balanced, these values strongly characterize the capabilities of the detector. The accuracy measure confirms that the Valeaua model leads to accurate detection in frames of both speech and silence. Also, the enhanced precision measure correspondingly lowers the false-positive value, i.e., non-speech frames has lower probability of being classified as speech. This result highly benefits applications such as speech enhancement, in which interferences may lead to severe degradation in practical performance. In a similar manner, the increase in recall decreases the false-negative measure. Thus, loss of information that typically lies in speech frames is obviated with higher probability.


Additional results are depicted in Fig. 3. Here, we focus on Ivory VAD [12] that achieved leading performance in all previous experimental results of this study. It can be deduced that this detector obtains a state-of-the-art performance of 95% in all reported measures when trained with Valeaua RIR model, which prevails competing VAD methods. Also, this detection system obtains leading accuracy, precision and recall measures across all tested reverberant setups. This outcome points on high generalization ability, robustness for noises and interferences, and prime accuracy in correctly distinguishing speech from silence.

#### 4. CONCLUSIONS

In this study, we have considered five different state-of-the-art deep-learning-based voice activity detectors that are trained with merely anechoic data. We have shown that these detectors experience substantial degradation in performance when

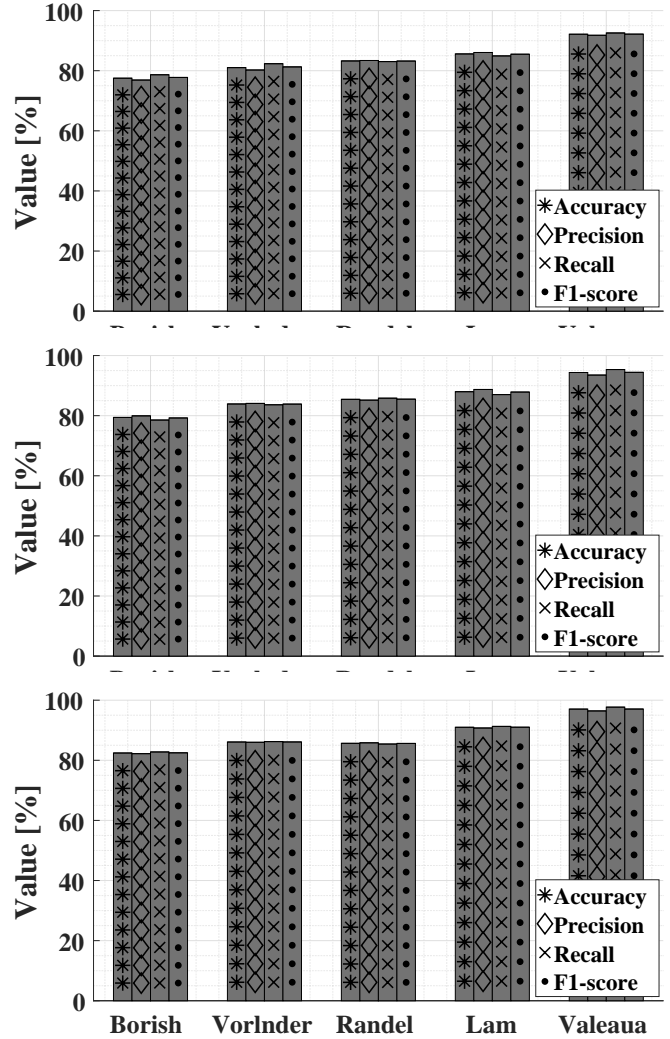


**Fig. 1.** Detection rate versus false alarm rate in a reverberant setup of a classroom. Comparison is made between the five different training RIR models. VADs (top to bottom): Ivory, Ariav-W, Kim, Wagner and Ariav-R.



**Fig. 2.** Performance of the five VADs in real reverberant conditions of (top to bottom): classroom, large concert hall, octagon library. Comparison is made between employing anechoic training (dark) and augmented training with Valeaua RIR model (light).

tested in reverberant conditions. To mitigate this mismatch, we simulated an augmented training set that contains both an anechoic corpus and its reverberant transformation, where the latter was calculated by convolution of the anechoic corpus with a fixed room impulse response model. This extension permitted detectors to learn unique patterns and audio-based features that represent reverberant settings. All five detection systems were trained with this augmented training set. This experiment was performed independently with five different room impulse response models. The training augmentation led to enhanced performance of all VAD systems when tested in three different real-life reverberant spaces. Improvement was obtained in terms of both accuracy, generalization



**Fig. 3.** Performance of Ivry VAD [12] in real reverberant conditions of (top to bottom): classroom, large concert hall, octagon library. Comparison is made between the five RIR training models.

and robustness abilities. Also, an average increase of 20% was held in accuracy, precision and recall measures with respect to non-reverberant training corpus. This study has also shown that the response model introduced by Valeaua [18] consistently leads to the best performance, regardless of the detector and the tested acoustic environment. That and more, the VAD introduced by Ivry [12] has achieved leading performance across all experiments. In future work, additional aspects such as feature engineering and dedicated architecture will be addressed in order to further enhance Ivry detector and adjust it for practical and reverberant acoustic scenarios.

## 5. REFERENCES

- [1] Keisuke Kinoshita, Marc Delcroix, Sharon Gannot, Emanuël AP Habets, Reinhold Haeb-Umbach, Walter Kellermann, Volker Leutnant, Roland Maas, Tomohiro Nakatani, Bhiksha Raj, et al., “A summary of the REVERB challenge: state-of-the-art and remaining challenges in reverberant speech processing research,” *EURASIP Journal on Advances in Signal Processing*, vol. 2016, no. 1, pp. 7, 2016.
- [2] Adam Kuklasiński, Simon Doclo, Søren Holdt Jensen, and Jesper Jensen, “Maximum likelihood PSD estimation for speech enhancement in reverberation and noise,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 24, no. 9, pp. 1599–1612, 2016.
- [3] Yan Zhao, DeLiang Wang, Ivo Merks, and Tao Zhang, “DNN-based enhancement of noisy and reverberant speech,” in *2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2016, pp. 6525–6529.
- [4] Kun Han, Yuxuan Wang, DeLiang Wang, William S Woods, Ivo Merks, and Tao Zhang, “Learning spectral mapping for speech dereverberation and denoising,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 23, no. 6, pp. 982–992, 2015.
- [5] Andreas Schwarz and Walter Kellermann, “Coherent-to-diffuse power ratio estimation for dereverberation,” *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 23, no. 6, pp. 1006–1018, 2015.
- [6] Tom Ko, Vijayaditya Peddinti, Daniel Povey, Michael L Seltzer, and Sanjeev 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)*. IEEE, 2017, pp. 5220–5224.
- [7] Ritwik Giri, Michael L Seltzer, Jasha Droppo, and Dong Yu, “Improving speech recognition in reverberation using a room-aware deep neural network and multi-task learning,” in *2015 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2015, pp. 5014–5018.
- [8] Ido Ariav, David Dov, and Israel Cohen, “A deep architecture for audio-visual voice activity detection in the presence of transients,” *Signal Processing*, vol. 142, pp. 69–74, 2018.
- [9] Johannes Wagner, Dominik Schiller, Andreas Seiderer, and Elisabeth André, “Deep learning in paralinguistic recognition tasks: Are hand-crafted features still relevant?,” in *Interspeech*, 2018, pp. 147–151.
- [10] Juntae Kim and Minsoo Hahn, “Voice activity detection using an adaptive context attention model,” *IEEE Signal Processing Letters*, vol. 25, no. 8, pp. 1181–1185, 2018.
- [11] Ido Ariav and Israel Cohen, “An end-to-end multi-modal voice activity detection using WaveNet encoder and residual networks,” *IEEE Journal of Selected Topics in Signal Processing*, vol. 13, no. 2, pp. 265–274, 2019.
- [12] Amir Ivry, Baruch Berdugo, and Israel Cohen, “Voice activity detection for transient noisy environment based on diffusion nets,” *IEEE Journal of Selected Topics in Signal Processing*, vol. 13, no. 2, pp. 254–264, 2019.
- [13] Jont B Allen and David A Berkley, “Image method for efficiently simulating small-room acoustics,” *The Journal of the Acoustical Society of America*, vol. 65, no. 4, pp. 943–950, 1979.
- [14] Jeffrey Borish, “Extension of the image model to arbitrary polyhedra,” *The Journal of the Acoustical Society of America*, vol. 75, no. 6, pp. 1827–1836, 1984.
- [15] Michael Vorländer, “Simulation of the transient and steady-state sound propagation in rooms using a new combined ray-tracing/image-source algorithm,” *The Journal of the Acoustical Society of America*, vol. 86, no. 1, pp. 172–178, 1989.
- [16] JH Rindel, “Modelling the angle-dependent pressure reflection factor,” *Applied Acoustics*, vol. 38, no. 2-4, pp. 223–234, 1993.
- [17] Yiu Wai Lam, “Issues for computer modelling of room acoustics in non-concert hall settings,” *Acoustical science and technology*, vol. 26, no. 2, pp. 145–155, 2005.
- [18] V Valeau, J Picaut, and M Hodgson, “On the use of a diffusion equation for room-acoustic prediction,” *The Journal of the Acoustical Society of America*, vol. 119, no. 3, pp. 1504–1513, 2006.
- [19] John S Garofolo, Lori F Lamel, William M Fisher, Jonathan G Fiscus, and David S Pallett, “Getting started with the DARPA TIMIT CD-ROM: An acoustic phonetic continuous speech database,” *National Institute of Standards and Technology (NIST), Gaithersburgh, MD*, vol. 107, pp. 16, 1988.
- [20] Rebecca Stewart and Mark Sandler, “Database of omnidirectional and B-format room impulse responses,” in *2010 IEEE International Conference on Acoustics, Speech and Signal Processing*. IEEE, 2010, pp. 165–168.
- [21] David Martin Powers, “Evaluation: from precision, recall and F-measure to ROC, informedness, markedness and correlation,” 2011.