



HAL
open science

Brouhaha: multi-task training for voice activity detection, speech-to-noise ratio, and C50 room acoustics estimation

Marvin Lavechin, Marianne Métais, Hadrien Titeux, Alodie Boissonnet, Jade Copet, Morgane Rivière, Elika Bergelson, Alejandrina Cristia, Emmanuel Dupoux, Hervé Bredin

► To cite this version:

Marvin Lavechin, Marianne Métais, Hadrien Titeux, Alodie Boissonnet, Jade Copet, et al.. Brouhaha: multi-task training for voice activity detection, speech-to-noise ratio, and C50 room acoustics estimation. IEEE Automatic Speech Recognition and Understanding (ASRU 2023), IEEE, Dec 2023, Taipei, Taiwan. pp.1–7. hal-04247647

HAL Id: hal-04247647

<https://hal.science/hal-04247647v1>

Submitted on 20 Oct 2023

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.

BROUHAHA: MULTI-TASK TRAINING FOR VOICE ACTIVITY DETECTION, SPEECH-TO-NOISE RATIO, AND C50 ROOM ACOUSTICS ESTIMATION

Marvin Lavechin^{*,1,2}, Marianne Métails^{*,1}, Hadrien Titeux¹, Alodie Boissonnet², Jade Copet², Morgane Rivière², Elika Bergelson³, Alejandrina Cristia¹, Emmanuel Dupoux^{1,2}, Hervé Bredin⁴

¹ LSCP, DEC, ENS, EHESS, CNRS, PSL University, Paris, France

² Meta AI Research, France ³ Duke University, North Carolina, USA

⁴ IRIT, Université de Toulouse, CNRS, Toulouse, France

ABSTRACT

Most automatic speech processing systems register degraded performance when applied to noisy or reverberant speech. But how can one tell whether speech is noisy or reverberant? We propose Brouhaha, a neural network jointly trained to extract speech/non-speech segments, speech-to-noise ratios, and C50 room acoustics from single-channel recordings. Brouhaha is trained using a data-driven approach in which noisy and reverberant audio segments are synthesized. We first evaluate its performance and demonstrate that the proposed multi-task regime is beneficial. We then present two scenarios illustrating how Brouhaha can be used on naturally noisy and reverberant data: 1) to investigate the errors made by a speaker diarization model (pyannote.audio); and 2) to assess the reliability of an automatic speech recognition model (Whisper from OpenAI). Both our pipeline and a pretrained model are open source and shared with the speech community.

Index Terms— voice activity detection, speech-to-noise ratio, speech clarity, acoustic environment, reverberation

1. INTRODUCTION

Robustness to degraded acoustic environments is a critical factor limiting the impact and adoption of speech technologies. Numerous sources of variations in the audio can degrade or hide the signal of interest and impact the performance of automatic speech processing systems. Be it automatic speech recognition (ASR) [1, 2, 3], speaker identification/diarization [4, 5], or speaker localization [6], most systems exhibit a loss of performance when applied in noisy or reverberant conditions.

While speech processing systems are being improved to handle degraded acoustic environments [7, 8, 9], little work

has been devoted to automatically predict the properties of the acoustic environment. A proposed approach involves using synthetic audio generated by applying an audio transformation of interest (e.g., reverberation). A neural network is then trained to extract the ‘strength’ of this audio transformation. This approach is most commonly used to develop systems that predict room acoustic measures like speech clarity (C_{50}), reverberation time (T_{60}) or direct-to-reverberant ratio (DRR) [10, 11, 12, 13, 14]. In practice, these values can be estimated directly from the room impulse response (RIR, the recording of a high-energy and bursty sound, such as a pistol shot or a balloon popping). However, in most cases, RIRs are not available, and we need to estimate the values of interest from the observed single channel audio recording. A similar approach has been adopted in [15] to automatically estimate the frame-level speech-to-noise ratio (SNR). The authors evaluate the performance of their system on synthetic data, but not on real data. In practice, real SNRs are not available making it impossible to compare the estimated values to the real ones. Thus, it remains unclear if such a system can generalize to real data.

Given the high interplay between noise and reverberation (the SNR may be influenced by how noise and speech sources reverberate, and it is harder to obtain reliable estimates of reverberation parameters in low SNR conditions [16, 17]), can we design a system that tackles both tasks simultaneously? This is one of the questions we address in this work. Our approach is closest to [18] who proposes to train a neural network for jointly estimating room acoustic parameters and the utterance-level SNR. However, the authors use a restrained set of noise segments which cast doubts on the ability of their model to generalize to unseen noises. More importantly, they do not evaluate their system with respect to the SNR, and they do not address the question of whether the proposed multi-task regime is beneficial for the estimation performance.

We propose *Brouhaha*, a model jointly trained on the speech/non-speech classification task and the SNR and C_{50} regression tasks. Our model is trained on 1,250 hours of synthetic audio generated from clean speech segments con-

* M. Lavechin and M. Métails equally contributed to this work. This work was granted access to the HPC resources of GENCI-IDRIS under the allocation 2022-AD011012554. It also benefited from the support of ANR-16-DATA-0004 ACLEW, ANR-17-EURE-0017, ANR-19-P3IA-0001; the J. S. Mc-Donnell Foundation; and ERC ExELang grant no 101001095.

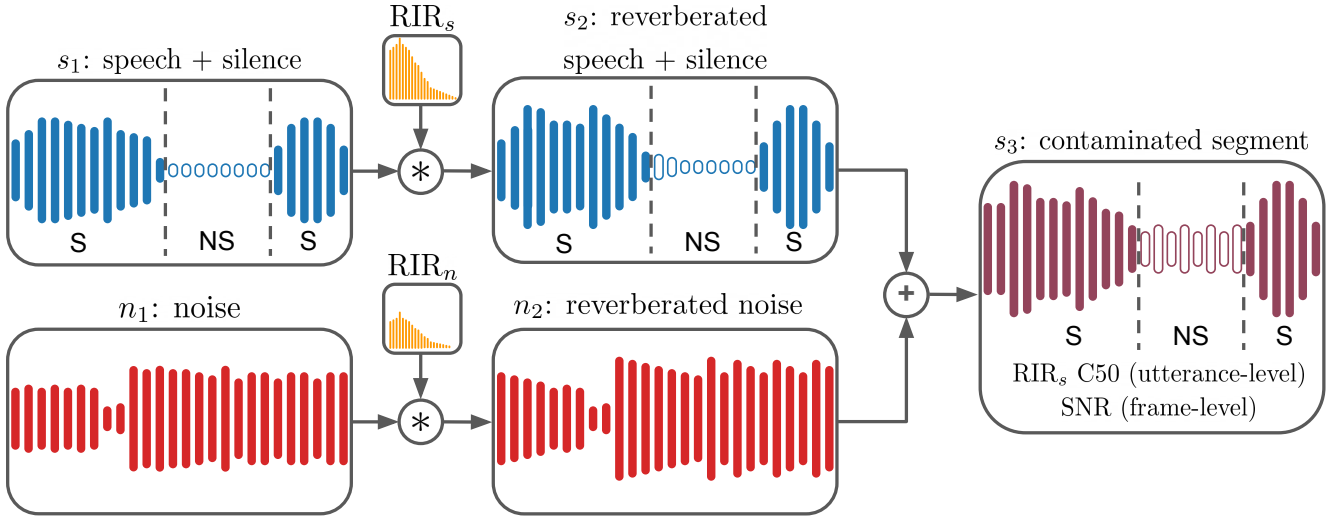


Fig. 1: Audio contamination pipeline. $s_1 \rightarrow s_2$: With probability $p_{RIR} = 0.9$, the clean speech segment (marked as S) contaminated with silence (marked as NS) s_1 is convolved with a randomly drawn impulse response RIR_s . $n_1 \rightarrow n_2$: With probability p_{RIR} , the randomly drawn noise segment n_1 is convolved with a randomly drawn impulse response RIR_n . $s_2 + n_2 \rightarrow s_3$: The reverberated speech segment s_2 and the reverberated noise segment n_2 are added together to obtain a Speech-to-Noise Ratio (SNR) randomly drawn between 0 and 30 dB. C_{50} is computed as the ratio of early (0 to 50 ms) and late (50 ms to the end of the response) energies of the room impulse response RIR_s . As noises can have a wide dynamic range and the utterance-level SNR captures only global information about the noise level, we recompute SNRs using a 2-second long sliding window shifted every 10 ms over s_2 and n_2 . To simulate different locations for the speech and noise sources, we used two randomly RIRs (as opposed to only one). This design choice aimed at forcing the model to rely solely on the reverberation characteristics of speech, disregarding the reverberation of noise, in order to predict the C_{50} measure. Labels obtained via this pipeline include: speech/non-speech (frame-level), C_{50} measure of RIR_s (utterance-level), and SNR (frame-level).

taminated with silence, noise and reverberation. We first demonstrate that the proposed multi-task regime is beneficial and compare the performance of *Brouhaha* against state-of-the-art systems. We then apply *Brouhaha* on real data (under naturally noisy and reverberant conditions) to: 1) analyze the error patterns of a speaker diarization system (*pyannote.audio* [19]); and 2) assess the reliability of an ASR system (Whisper [20]). In addition to showing how *Brouhaha* can be used, these experiments constitute evidence that our system is applicable to real data.

Beyond the scientific interest of exploring the effectiveness of the proposed multi-task training regime and assessing the applicability of the method on real data after training on synthetic ones, we believe our work has a strong practical interest. Unlike previous work [15, 18], *Brouhaha* can be applied to any audio regardless of whether it contains speech, non-speech or both. By using our system, there is no requirement to implement a preliminary voice activity detection system prior to obtaining SNR and C_{50} values. We believe such advancement, in addition to a simple user interface (one python command!), significantly aids empowering researchers who may not possess expertise in speech processing or machine learning to make the most out of speech technology.

2. AUDIO CONTAMINATION PIPELINE

We start from: 1) a set of clean speech segments that will be contaminated; 2) a set of noise segments used to simulate noisy conditions; and 3) a set of RIRs to simulate reverberation. The clean speech segments are contaminated following the steps presented in Figure 1, which we will not repeat here.

3. MULTI-TASK TRAINING

We tackled the voice activity detection problem as a classification problem where, for each 16-ms frame, the expected output is 1 if there is speech, 0 otherwise. C_{50} and SNR estimations were tackled as regression problems where, for each 16-ms frame, the expected output is the actual C_{50} or SNR in dB. We tackled the C_{50} estimation at the frame level during training – despite the label being at the utterance level – to allow the model to return smoother transitions when a change in C_{50} is detected at inference time.

At training time, short fixed length sub-sequences are drawn randomly from the training set and gradient-descent is used to minimize the multi-task loss function $\mathcal{L} = \mathcal{L}_{VAD} + \mathcal{L}_{C_{50}} + \mathcal{L}_{SNR}$, where \mathcal{L}_{VAD} is the binary cross-entropy loss, and $\mathcal{L}_{C_{50}}$ and \mathcal{L}_{SNR} are mean squared error (MSE) losses. Before

training, $\mathcal{L}_{C_{50}}$ and \mathcal{L}_{SNR} are normalized by their maximum value (computed over 10 batches) to ensure all three losses lie between 0 and 1. We computed \mathcal{L}_{SNR} only over speech frames as the SNR is not defined on non-speech frames.

4. EXPERIMENTS

4.1. Datasets

Our audio contamination pipeline requires three types of audio data: 1) clean speech segments; 2) noise segments; and 3) RIRs. A pretrained VAD model [19] was applied to find non-speech segments in 1000 hours of clean read-speech, retrieved from the LibriSpeech [21]. Predicted non-speech segments were extended with silence to obtain a ratio of approximately 30 % of non-speech. We used noise segments from AudioSet [22] and discarded human vocalizations. We also downsampled music segments from 38 % to 5 %, leading to a total of 1500 hours of noise segments. Finally, 385 impulse responses were obtained from EchoThief [23] and the MIT Acoustical Reverberation Scene [24] datasets. We used the same train/dev/test split originally proposed in LibriSpeech. Noise segments and impulse responses were randomly split into 80 %, 10 % and 10 % for the training, development and test set, respectively. All files used in this paper consist of 16-kHz single-channel recordings.

4.2. Evaluation metrics

We evaluated *Brouhaha* performance on the VAD task using the F-score between precision and recall, such as implemented in *pyannote.metrics* [25]. SNR and C_{50} predictions were evaluated using the mean absolute error (MAE) at the frame level. Since SNR is not defined on non-speech frames, the SNR was only evaluated across speech frames.

4.3. Architecture, optimization and training procedure

The model consists of SincNet (using the configuration in [26]), followed by a stack of bidirectional long short-term memory (LSTM) and feed-forward layers. Finally, we have three parallel layers: one classification layer (with *softmax* activation) that returns the predicted probability of speech, and two regression layers that return the predicted SNR and C_{50} (with *sigmoid* activation parametrized between -15 and 80 dB for the SNR, and -10 and 60 dB for the C_{50}).

We trained 144 different architectures across different sets of hyperparameters, varying the duration of the input sequences: 4, 6, 8, or 10 seconds; the batch size: 32, 64, or 128 sequences; the size of the hidden LSTM layers: 128 or 256; the number of LSTM layers: 2 or 3; and the dropout proportion: 0, 30 or 50 %. The best architecture was trained with 6-s segments, a batch size of 64 sequences, 3 LSTM layers of size 256, and a dropout proportion of 50 %. The best architecture was selected on the validation metric: an

average of the VAD F-score, SNR and C_{50} MAEs, with the latter two normalized by the maximum error to balance the contribution of each term.

5. RESULTS

5.1. The effect of multi-task training

Table 1: Performance on unseen synthetic data (our test set) in terms of F-score (VAD) and mean absolute errors (SNR and C_{50}). A checkmark below a given training task indicates that the associated loss is activated during training.

Training tasks:			VAD	SNR	C_{50}
VAD	SNR	C_{50}	F-score (%)	MAE (dB)	MAE (dB)
✓	✓	✓	93.7	4.1	3.5
✓	✓	—	93.7	4.2	—
✓	—	✓	93.6	—	3.8
—	✓	✓	—	4.3	3.7
✓	—	—	93.5	—	—
—	✓	—	—	4.3	—
—	—	✓	—	—	4.2

Table 1 shows performance obtained by models trained to solve either one, two or three of the proposed tasks (VAD, SNR, C_{50}). All models shared the same set of hyperparameters, only the dimension of the output layer differed. Results indicate that the multi-task training regime is beneficial: the model trained simultaneously on the three tasks obtained better performance than models trained on two tasks which themselves obtained better performance than models trained on a single task. The largest performance gain is observed for the C_{50} estimation, with a decrease of 0.7 dB in terms of MAE between the single-task and the three-tasks training regime. These results seem to show that sharing weights during training helps better solve the proposed three tasks. Not only does using a single model provide a performance gain, but it is also more convenient and computationally efficient.

5.2. Voice activity detection

Table 2 shows voice activity detection performance obtained by *Brouhaha* and a state-of-the-art system (*pyannote.audio* [19]). We consider two evaluation sets: 1) our test set made of unseen synthetic audio data (referred as ‘synthetic’ in the table); and 2) BabyTrain [27], a corpus of highly naturalistic child-centered recordings (referred as ‘real’ in the table). Specifically, BabyTrain recordings are acquired via child-worn microphones as they go about their everyday activities and are widely used in language acquisition research [28]. Child-centered recordings are notoriously challenging for speech processing systems as they contain spontaneous

Table 2: Voice activity detection F-score obtained by *Brouhaha* and *pyannote.audio* pretrained system [19]. Numbers are reported on synthetic data (our test set) and on real data (*BabyTrain* [27]).

Data type	System	VAD F-score (%)
synthetic	<i>Brouhaha</i> (ours)	93.7
	<i>pyannote.audio</i> [19]	89.0
real	<i>Brouhaha</i> (ours)	77.2
	<i>pyannote.audio</i> [19]	80.8

and overlapping speech, and a wide variety of noisy and reverberant conditions.

Results show a strong advantage for *Brouhaha* over *pyannote.audio* on unseen synthetic data (4.7% absolute difference in terms of F-score). This indicates that, on highly noisy and reverberant synthetic audio, our system is competitive on the VAD task. Admittedly, *Brouhaha* has an advantage over *pyannote.audio* as the latter has not been trained on synthetically noisy and reverberant audio. Turning to a performance comparison on real data, numbers reveal that *pyannote.audio* outperforms *Brouhaha* by a 3.6% absolute difference in terms of F-score. This result suggests that training a VAD system on LibriSpeech [21] contaminated with reverberation and additive noise might not be optimal, and this is despite the precautions taken in simulating challenging noisy and reverberant conditions. Nonetheless, LibriSpeech is currently the only source of clean speech available in sufficiently large quantities to run our audio contamination pipeline and obtain SNR and C_{50} labels.

5.3. Speech-to-noise ratio estimation

Table 3: Mean absolute error on the SNR estimation task computed on unseen synthetic data (our test set). All predicted and gold SNRs are brought back to the $[-15, 30]$ dB range as done in [15]. For a given speech utterance, the heuristic estimates the noise (resp. speech) power as the mean power of non-speech (resp. speech) frames within a 6-s window centered around each annotated speech frame (defaulting to the average SNR when no non-speech frames were found within the 6-s window).

System	SNR MAE (dB)
<i>Brouhaha</i> (ours)	2.3
Heuristic	8.4
Li et al. [15]	12.5

Table 3 shows MAE performance on the SNR estimation task computed on our test set made of unseen synthetic audio data for: 1) *Brouhaha*; 2) a heuristic using the oracle VAD that es-

timates the noise (resp. speech) as the mean power of neighboring non-speech (resp. speech) frames; and 3) the system proposed in [15] (a 4-layer LSTM trained from mel frequency cepstral coefficients).

Results indicate that *Brouhaha* is better at estimating the frame-level SNR than our heuristic, with an absolute difference of 6.1 dB in terms of MAE (note that both systems use a 6-s window as input, and that our heuristic requires oracle VAD boundaries). Surprisingly, our heuristic performs better than the system proposed in [15] with a 4.1 dB absolute difference in terms of MAE. This indicates that [15] struggles generalizing to unseen noise or to reverberant environments. Unfortunately, we could not compare systems on the test used in [15] as the latter has not been publicly released.

5.4. C_{50} estimation

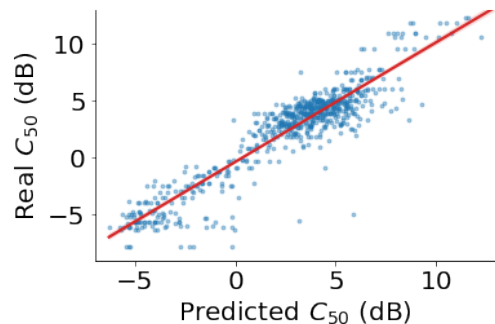


Fig. 2: C_{50} estimation. Real C_{50} against C_{50} predicted by *Brouhaha* on 1000 utterances from the BUT Speech@FIT Reverb dataset [29].

We ran *Brouhaha* on the BUT Speech@FIT Reverb dataset [29]. This dataset consists of LibriSpeech test-clean utterances retransmitted by a loudspeaker in 5 different rooms. For each room, the speaker was placed on 5 positions on average and retransmitted utterances were recorded with 31 microphones. RIRs were measured multiple times for each speaker position. Here, we compare the real C_{50} (averaged over between 1 and 9 duplicated RIR measures) to the C_{50} predicted by *Brouhaha* on 1000 randomly drawn utterances.

Figure 2 shows a strong correlation between the real and the predicted C_{50} , with a R^2 of .85 and a mean average error of 1.1 dB. We would have liked to compare the performance of our system on the C_{50} estimation task with other systems, but we could not find any open-source pre-trained C_{50} estimators despite extensive research in this area [11, 12, 14].

5.5. Investigating speaker diarization errors

We ran a pretrained *pyannote.audio* speaker diarization pipeline [19] on the VoxConverse dataset [30] and evaluated its performance at *Brouhaha* frame resolution (16 ms). Each frame can either be classified as: 1) missed detection

(when the speaker diarization pipeline incorrectly classifies a speech frame as non-speech): 2) false alarm (the other way around); 3) speaker confusion (when a speech frame is assigned to the wrong speaker); or 4) correct. Figure 3 focuses on speaker confusion (but the same pattern holds for missed detections) and shows the distribution of predicted SNR (left) and C_{50} (right) depending on whether the speech frame was assigned to the correct speaker. There is a clear trend as far as SNR is concerned: *pyannote.audio* is much more likely to confuse speakers in low (predicted) SNR regions. Similarly, the accuracy degrades significantly as we get closer to the lowest predicted C_{50} values.

Exploring the errors made by a pretrained system can provide valuable insights for developing effective strategies. In our case, one might devise strategies to address the issue of high speaker confusion in low SNR conditions: increasing the weight of low-SNR sequences in the training loss, or running speech enhancement algorithms on low SNR areas for instance.

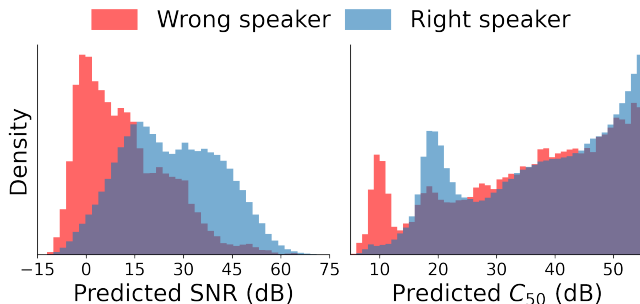


Fig. 3: Investigating speaker diarization errors. Distribution of SNR (left) and C_{50} (right) predicted by *Brouhaha* as a function of whether a pretrained speaker diarization system [19] assigns a speech frame to a wrong (red) or to the right speaker (blue).

5.6. Assessing the reliability of an ASR system

We ran Whisper large ASR system [20] on highly naturalistic speech utterances from the American English Bergelson corpus [31, 32] (child-centered recordings, similar to the ones used in Section 5.2). We evaluate the performance of Whisper using the percentage hits (i.e., percentage of words correctly transcribed). We include a total of 804 utterances at least 5-words long (as short sequences most often led to a score 0% or 100%).

Figure 4 shows the average percentage of hits obtained by Whisper for utterances binned according to their predicted SNR (top panel) or C_{50} (bottom panel) decile. On average, Whisper correctly transcribes 83% of the words on utterances whose SNR belongs in the [12, 24] dB (last SNR decile, top panel). This number decreases as the SNR decreases until Whisper successfully transcribes only 38% of the words on

utterances whose SNR is in the $[-9, -4]$ dB range (first SNR decile). Although utterances whose predicted C_{50} is high tend to be better transcribed by Whisper, the trend with respect to the C_{50} is less clear (bottom panel). In conclusion, this experiment shows that *Brouhaha* can be used to investigate the robustness of ASR systems to noisy conditions.

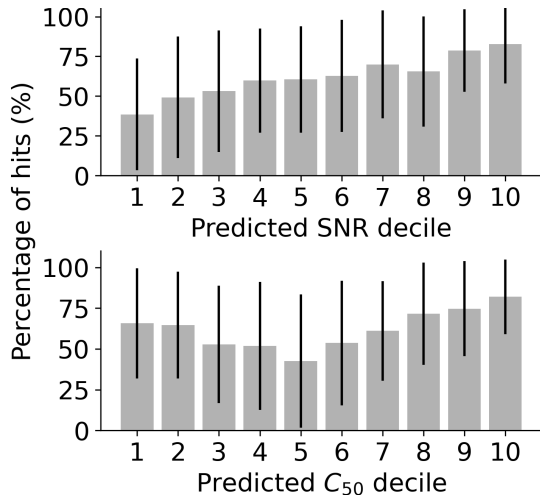


Fig. 4: Assessing the reliability of an ASR system. Percentage of hits obtained by Whisper large as a function of predicted SNR decile (top panel) and predicted C_{50} decile (bottom panel). Bars represent the percentage of hits averaged across utterances. Thin black lines represent standard errors.

6. CONCLUSION AND FUTURE WORK

We proposed *Brouhaha*, a model jointly trained on the voice activity detection, SNR, and C_{50} estimation tasks. After evaluating the performance of our system and demonstrating that the multi-task training regime is beneficial, we illustrated two use cases showing how our model can be used on real data. Beyond investigating errors made by speech processing systems or assessing their reliability in noisy and reverberant conditions, we foresee other potential downstream tasks, e.g., SNR- or C_{50} -based microphone selection [33] or SNR-aware speech enhancement [34]. Future work could explore these downstream tasks, the use of spontaneous clean speech to improve VAD performance, or the estimation of other room acoustic parameters, such as T_{60} or DRR. Both a pre-trained model and our audio contamination pipeline are shared with the community¹.

¹<https://github.com/marianne-m/brouhaha-vad>

7. REFERENCES

- [1] 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 *International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2015, pp. 5014–5018.
- [2] Keisuke Kinoshita, Marc Delcroix, Sharon Gannot, Emanuël A P 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*, pp. 1–19, 2016.
- [3] Hannes Gamper, Dimitra Emmanouilidou, Sebastian Braun, and Ivan J Tashev, “Predicting word error rate for reverberant speech,” in *International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2020, pp. 491–495.
- [4] Xiaojia Zhao, Yuxuan Wang, and DeLiang Wang, “Robust speaker identification in noisy and reverberant conditions,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 22, no. 4, pp. 836–845, 2014.
- [5] Neville Ryant, Prachi Singh, Venkat Krishnamohan, Rajat Varma, Kenneth Ward Church, Christopher Cieri, Jun Du, Sriram Ganapathy, and Mark Y. Liberman, “The third dihard diarization challenge,” in *Interspeech*, 2021.
- [6] Soumitro Chakrabarty and Emanuël AP Habets, “Broadband doa estimation using convolutional neural networks trained with noise signals,” in *Workshop on Applications of Signal Processing to Audio and Acoustics (WASPAA)*. IEEE, 2017, pp. 136–140.
- [7] 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 *International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2017, pp. 5220–5224.
- [8] Chris Donahue, Bo Li, and Rohit Prabhavalkar, “Exploring speech enhancement with generative adversarial networks for robust speech recognition,” in *International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2018, pp. 5024–5028.
- [9] Arun Narayanan, James Walker, Sankaran Panchapagesan, Nathan Howard, and Yuma Koizumi, “Learning mask scalars for improved robust automatic speech recognition,” in *2022 IEEE Spoken Language Technology Workshop (SLT)*. IEEE, 2023, pp. 317–323.
- [10] James Eaton, Nikolay D Gaubitch, Alastair H Moore, and Patrick A Naylor, “The ace challenge—corpus description and performance evaluation,” in *Workshop on Applications of Signal Processing to Audio and Acoustics (WASPAA)*. IEEE, 2015, pp. 1–5.
- [11] Pablo Peso Parada, Dushyant Sharma, Jose Lainez, Daniel Barreda, Toon van Waterschoot, and Patrick A Naylor, “A single-channel non-intrusive c50 estimator correlated with speech recognition performance,” *Transactions on Audio, Speech, and Language Processing*, vol. 24, no. 4, pp. 719–732, 2016.
- [12] Feifei Xiong, Stefan Goetze, Birger Kollmeier, and Bernd T Meyer, “Exploring auditory-inspired acoustic features for room acoustic parameter estimation from monaural speech,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 26, no. 10, pp. 1809–1820, 2018.
- [13] Nicholas J Bryan, “Impulse response data augmentation and deep neural networks for blind room acoustic parameter estimation,” in *International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2020, pp. 1–5.
- [14] Hannes Gamper, “Blind c50 estimation from single-channel speech using a convolutional neural network,” in *International Workshop on Multimedia Signal Processing (MMSP)*. IEEE, 2020, pp. 1–6.
- [15] Hao Li, DeLiang Wang, Xueliang Zhang, and Guanglai Gao, “Frame-level signal-to-noise ratio estimation using deep learning,” in *Interspeech*, 2020, pp. 4626–4630.
- [16] Heinrich Löllmann, Andreas Brendel, and Walter Kellermann, “Comparative study of single-channel algorithms for blind reverberation time estimation,” in *International Congress on Acoustics (ICA)*, 2019.
- [17] 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.
- [18] David Looney and Nikolay D Gaubitch, “Joint estimation of acoustic parameters from single-microphone speech observations,” in *International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2020, pp. 431–435.
- [19] Hervé Bredin, Ruiqing Yin, Juan Manuel Coria, Gregory Gelly, Pavel Korshunov, Marvin Lavechin, Diego

- Fustes, Hadrien Titeux, Wassim Bouaziz, and Marie-Philippe Gill, “Pyannote.audio: neural building blocks for speaker diarization,” in *International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2020, pp. 7124–7128.
- [20] Alec Radford, Jong Wook Kim, Tao Xu, Greg Brockman, Christine McLeavey, and Ilya Sutskever, “Robust speech recognition via large-scale weak supervision,” Tech. Rep., OpenAI, 2022.
- [21] Vassil Panayotov, Guoguo Chen, Daniel Povey, and Sanjeev Khudanpur, “Librispeech: an asr corpus based on public domain audio books,” in *International conference on acoustics, speech and signal processing (ICASSP)*. IEEE, 2015, pp. 5206–5210.
- [22] Jort F Gemmeke, Daniel PW Ellis, Dylan Freedman, Aren Jansen, Wade Lawrence, R Channing Moore, Manoj Plakal, and Marvin Ritter, “Audio set: An ontology and human-labeled dataset for audio events,” in *International conference on acoustics, speech and signal processing (ICASSP)*. IEEE, 2017, pp. 776–780.
- [23] Chris Warren, “Echothief impulse response library,” .
- [24] James Traer and Josh H McDermott, “Statistics of natural reverberation enable perceptual separation of sound and space,” *Proceedings of the National Academy of Sciences*, vol. 113, no. 48, pp. E7856–E7865, 2016.
- [25] Hervé Bredin, “pyannote.metrics: a toolkit for reproducible evaluation, diagnostic, and error analysis of speaker diarization systems,” in *Interspeech*, 2017.
- [26] Mirco Ravanelli and Yoshua Bengio, “Speaker recognition from raw waveform with sincnet,” in *Spoken Language Technology Workshop (SLT)*, 2018, pp. 1021–1028.
- [27] Marvin Lavechin, Ruben Bousbib, Hervé Bredin, Emmanuel Dupoux, and Alejandrina Cristia, “An open-source voice type classifier for child-centered daylong recordings,” in *Interspeech*, 2020.
- [28] Marvin Lavechin, Maureen de Seyssel, Lucas Gautheron, Emmanuel Dupoux, and Alejandrina Cristia, “Reverse engineering language acquisition with child-centered long-form recordings,” *Annual Review of Linguistics*, vol. 8, pp. 389–407, 2022.
- [29] Igor Szöke, Miroslav Skácel, Ladislav Mošner, Jakub Paliesek, and Jan Černocký, “Building and evaluation of a real room impulse response dataset,” *Journal of Selected Topics in Signal Processing*, vol. 13, no. 4, pp. 863–876, 2019.
- [30] Joon Son Chung, Jaesung Huh, Arsha Nagrani, Triantafyllos Afouras, and Andrew Zisserman, “Spot the Conversation: Speaker Diarisation in the Wild,” in *Interspeech*, 2020, pp. 299–303.
- [31] Elika Bergelson, Marisa Casillas, Melanie Soderstrom, Amanda Seidl, Anne S. Warlaumont, and Andrei Amatuni, “What do North American babies hear? A large-scale cross-corpus analysis.,” *Developmental science*, vol. 22 1, pp. e12724, 2019.
- [32] Elika Bergelson, “SEEDLingS HomeBank corpus,” 2017.
- [33] Martin Wolf and Climent Nadeu, “Towards microphone selection based on room impulse response energy-related measures,” in *Workshop on Speech and Language Technologies for Iberian Languages, Porto Salvo, Portugal*, 2009, pp. 61–64.
- [34] Szu-Wei Fu, Yu Tsao, and Xugang Lu, “SNR-aware convolutional neural network modeling for speech enhancement.,” in *Interspeech*, 2016, pp. 3768–3772.