



Unsupervised Voice Activity Detection by Modeling Source and System Information using Zero Frequency Filtering

Eklavya Sarkar^{1,2}, RaviShankar Prasad¹, Mathew Magimai.-Doss¹

¹Idiap Research Institute, Martigny, Switzerland

²Ecole polytechnique fédérale de Lausanne, Switzerland

{eklavya.sarkar, ravi.prasad, mathew}@idiap.ch

Abstract

Voice activity detection (VAD) is an important pre-processing step for speech technology applications. The task consists of deriving segment boundaries of audio signals which contain voicing information. In recent years, it has been shown that voice source and vocal tract system information can be extracted using zero-frequency filtering (ZFF) without making any explicit model assumptions about the speech signal. This paper investigates the potential of zero-frequency filtering for jointly modeling voice source and vocal tract system information, and proposes two approaches for VAD. The first approach demarcates voiced regions using a composite signal composed of different zero-frequency filtered signals. The second approach feeds the composite signal as input to the rVAD algorithm. These approaches are compared with other supervised and unsupervised VAD methods in the literature, and are evaluated on the Aurora-2 database, across a range of SNRs (20 to -5 dB). Our studies show that the proposed ZFF-based methods perform comparable to state-of-art VAD methods and are more invariant to added degradation and different channel characteristics.

Index Terms: Voice activity detection, zero-frequency filtering, speech analysis, signal processing.

1. Introduction

Voice activity detection (VAD) refers to the task of identifying segment boundaries in audio signals which essentially contain voicing information, and typically is one of the first steps to be carried out in any speech technology. Computational efficiency and robustness to noisy data are thus essential pre-requisites for any state-of-the-art voice activity detector. VAD methods can be broadly categorised as unsupervised and supervised methods.

Early unsupervised methods made use of simple energy-based features and temporal parameters such as zero-crossing rate (ZCR) [1, 2], before applying a discriminator model to compute the speech/non-speech decision boundary. Spectral features based on autocorrelation [3–5], mel-frequency cepstral coefficients (MFCCs) [4], skewness and kurtosis of linear prediction (LP) residual [6], spectral shape [7], harmonic structure [8], voicing [9], cepstral features [10, 11], perceptual spectral flux [12], spectral flatness (SF) and short-term energy [13], and speech enhancement and denoising through pitch indicators [14] were proposed to improve the performance and robustness of these systems in the presence of noise.

Supervised VAD models mostly rely on a likelihood ratio test (LRT) over the estimated parameters in a maximum likelihood (ML) framework [15]. Recent deep learning models have also shown success. Deep belief networks [16] combine various acoustic features through multiple non-linear hidden layers to discover the manifold of the features and observe regularity

among them to predict the frame class. Hybrid deep architectures, incorporating convolutional neural networks (CNNs) and long short-term memory (LSTM) models, based on raw waveform [17] and spectrograms [18], have also been proposed. Although these methods can yield good performance, they come at a high computational cost, requiring training or fine-tuning a pre-trained model, ground truth labels, and do not rely as much on prior knowledge as unsupervised methods.

The focus of this paper lies on unsupervised VAD methods, which tend to incorporate prior knowledge about voice source and vocal tract system, typically through source-system decomposition methods, such as linear prediction analysis and cepstral analysis. Unlike such methods, which make a mathematical model assumption, it has been shown in recent years that voice source and vocal tract system information can be effectively extracted using zero-frequency filtering [19, 20] without making any such explicit model assumptions about the speech signal. This paper investigates the potential of zero-frequency filtering for jointly modeling voice source and vocal tract system system information to perform VAD. To that end, we demonstrate that voice activity detection can be effectively achieved by combining the outputs of a bank of zero-frequency filters that carry information related to the fundamental frequency (f_0), and the first (F_1) and second (F_2) formants.

The remainder of the paper is organized as follows. Section 2 provides an overview on the background of zero-frequency filtered signals. Section 3 presents the two approaches for VAD using the ZFF signals. Section 4 gives the experimental setup used to validate our method, and Section 5 summarizes the results. Finally, Section 6 concludes the paper.

2. Background

Zero-frequency filtering (ZFF) was originally proposed in the context of extracting information related to voice source [19]. In this method, a speech signal s is first passed through cascaded digital resonators, implemented as an integrator, centered at 0 Hz, i.e. a zero-frequency filter. The resulting impulse response is given in eq. (1) and its equivalent transfer function in eq. (2).

$$x[n] = s[n] - 2x[n-1] + x[n-2] \quad (1)$$

$$H[z] = \frac{1}{1 - 2z^{-1} + z^{-2}} \quad (2)$$

A trend removal (i.e. local mean subtraction) step, based on an estimate of the periodicity of the speech signal, is then applied to the output of the cascaded resonators to obtain glottal closure instance (GCI) locations and strength of excitation information. The trend removal operation is described by eq. (3).

$$y[n] = x[n] - \frac{1}{2N+1} \sum_{k=n-N}^{n+N} x[k]; N+1 \leq n \leq L-N. \quad (3)$$

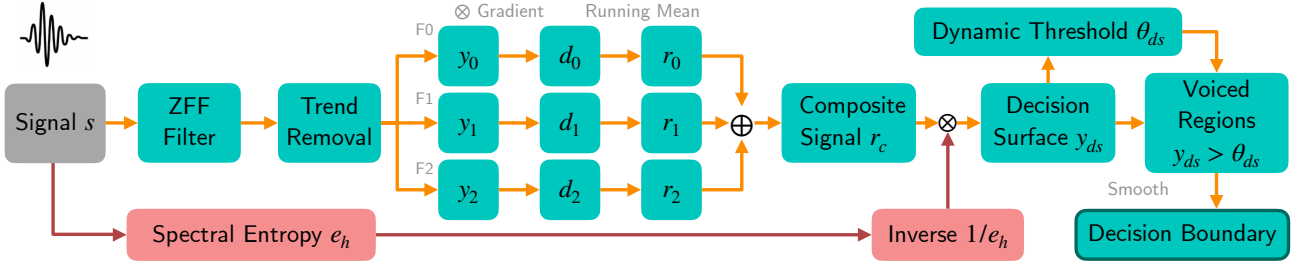


Figure 1: Complete pipeline of proposed method to derive a decision boundary for voice activity detection.

L corresponds to the length of the signal and $2N + 1 \sim T_0$ is the trend removal window duration. T_0 is estimated through autocorrelation.

In a recent work, it was shown that, although the zero-frequency filter heavily damps the high frequency regions in the speech signal, F_1 and F_2 information can still be estimated by moderating the trend removal window duration [20].

Figure 2 illustrates the extraction of f_0 , F_1 , and F_2 evidences. A voiced speech signal $s(n)$ and the location of glottal closure instants (GCIs) (---) are presented in (a1). The latter is derived from the negative-to-positive zero-crossing locations using the method in [19]. The corresponding discrete Fourier transform (DFT) spectrum $S(\omega)$ (—) and the inverse filter response (—) obtained through LP analysis are given in (b1). The fundamental frequency value is obtained at the first peak (●) of the DFT spectrum $S(\omega)$. The formants can also be located at global and relative peak locations (●) in the spectral envelope. The signals $y_0(n)$, $y_1(n)$, $y_2(n)$ and the corresponding GCI locations, shown in (a2–a4), are obtained by passing the speech signal through the ZFF filter, and computing a trend removal step with three different trend removal windows, in this case T_0 , $T_0/5$, $T_0/10$, where the estimate of fundamental period T_0 is calculated through autocorrelation. Their respective DFT response $Y_0(\omega)$, $Y_1(\omega)$, $Y_2(\omega)$, and corresponding peaks (●), as well an overlay of the envelope of $S(\omega)$, are given in (b2–b4).

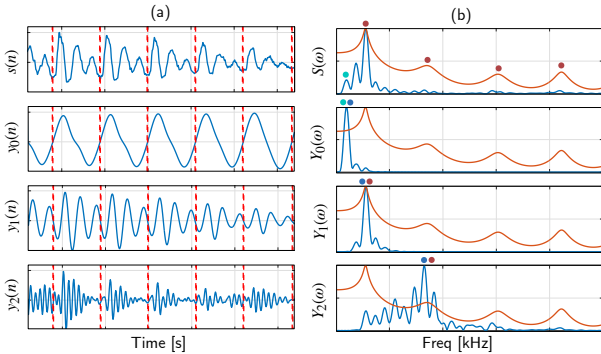


Figure 2: (a1) Speech signal. (a2–a4) ZFF signals $y_0(n)$, $y_1(n)$, $y_2(n)$. GCI locations (---). (b1) $S(\omega)$ (—) and its envelope (—). Formant peaks (●). Fundamental frequency peak (●). (b2–b4) $Y_0(\omega)$, $Y_1(\omega)$, $Y_2(\omega)$ (—), and respective peaks (●).

3. Proposed Method

As motivated in the previous section, zero-frequency filtered signals effectively encode salient source and system speech information, such as the fundamental frequency f_0 , and for-

nants F_1 and F_2 . Furthermore, these evidences are also robust to noise, as the SNR is high in the regions where the source and system related information manifests in the time-frequency domain. More precisely, GCIs in the time domain are high SNR regions, and similarly, F_1 and F_2 appear as peaks in the short-time spectrum. Furthermore, while selectively focusing on the source and system evidences, ZFF heavily damps rest of the spectral information, potentially suppressing other interferences. Considering these aspects, we propose the following two approaches for VAD based on zero-frequency filtering:

1. In the first approach, illustrated in Figure 1, the outputs of the different zero-frequency filters are combined to obtain a composite signal carrying f_0 , F_1 , and F_2 related information. Voiced regions are demarcated by applying a spectral-based weighing.
2. In the second approach, the composite signal is given as the input to another VAD algorithm.

In the remainder of this section, we present the details of the first approach.

Algorithm 1 Proposed VAD method using ZFF.

1. Compute $x[n]$: $x[n] = s[n] \otimes H_Z[n]$, using eq. (2).
 2. Estimate T_0 (i.e., $1/f_0$) for $s[n]$ using autocorrelation.
 3. Obtain $y_0[n]$, $y_1[n]$, $y_2[n]$ from $x[n]$ using eq. (3), with windows sizes: $2N + 1 \sim [T_0, T_0/5, T_0/10]$.
 4. Determine $d_0[n]$, $d_1[n]$, $d_2[n]$ by weighing $y_0[n]$, $y_2[n]$, $y_3[n]$ with their gradients: $d_i = y_i[n] \cdot (y_i[n] - y_i[n - 1])$ to highlight the regions of interest [21, 22].
 5. Obtain the running mean signal $r_0[n]$, $r_1[n]$, $r_2[n]$ from $d_0[n]$, $d_1[n]$, $d_2[n]$ over a duration of 40 ms.
 6. Calculate the accumulated signal $r_c[n] = r_0[n] + r_1[n] + r_2[n]$, and normalize it between [0–1].
 7. Obtain the spectral entropy $e_h[n]$ from $x[n]$ by computing the spectrum with FFT over a window of 20 ms.
 8. Obtain the decision surface $y_{ds}[n] = r_c[n] \cdot 1/e_h[n]$.
 9. Derive a dynamic threshold every 300 ms: $\theta_{ds} = ds_{\min} + (ds_{\text{med}}/3)$, where $ds_{\min} = \min\{y_{ds}[n]\}$ and $ds_{\text{med}} = \text{median}\{y_{ds}[n]\}$.
 10. Demarcate voiced regions as $y_{ds}[n] \geq \theta_{ds}$.
 11. Smooth decision boundary by eliminating short duration outlier segments and merging those in close proximity.
-

Algorithm 1 presents the proposed VAD method, with successive steps represented in Figure 1. Figure 3 illustrates the principal components of this technique. Figure 3 (a) shows a naturally corrupted speech signal along with the boundary demarcations for voiced segments, obtained using the proposed method. Figure 3 (b) shows the composite signal r_c obtained

after applying the zero-frequency filtering stage, and Figure 3 (c) the inverse spectral entropy weight that is applied to it. Finally, Figure 3 (d) shows the resulting decision surface based on dynamic threshold estimation every 300 ms. The duration is chosen on the heuristic assumption that the baseline for the decision surface will not significantly change within this period.




Figure 3: a) Naturally corrupted speech signal s and final decision boundary. b) Accumulated ZFF signals r_c c) Inverse spectral entropy $1/e_h$ d) Decision surface y_{ds} and dynamic threshold θ_{ds} .

4. Experimental Setup

This section presents the dataset, baseline methods, and the evaluation metrics used for our experimental setup.

4.1. Dataset

We demonstrate the performance of our VAD method on the **Aurora-2** dataset [23], which contains clean and synthetically degraded speech utterances. Clean speech utterances are borrowed from the TIDigits dataset [24], downsampled at 8 kHz with ideal low-pass filter characteristics, to which noise is added from the NOISEX-92 dataset [25]. The overall data is distributed in 4 directories as train (*Train*), and test sets (*Test A*, *Test B*, and *Test C*). The *Train* set contains 8440 utterances, with a mix of clean signals, and signals degraded with ‘train’, ‘car’, ‘babble’, and ‘exhibition hall’ noises. The *Test A* set contains different utterances subjected to a channel characteristics and degraded with same noise types, at SNR levels of 20, 15, 10, 5, 0, and -5 dB. The *Test B* and *Test C* sets contain signals degraded with different noise types and channel characteristics, to present a disjoint environment from training. Together the three test sets contain 4004 utterances. In total, the train and test sets contain 8440 and 70070 audio files respectively. The labels are obtained from [14], generated using a HTK recognizer [26], trained on 12 MFCC coefficients, $\Delta+\Delta\Delta$ s, and log-energy, computed over the *Train* set, modeled by 16 HMM states, each represented by 3 Gaussian mixtures.

4.2. Baseline Methods

Performance of the proposed method (V_{ZFF}) is compared against several supervised and unsupervised methods in the literature, implemented based on the hyper-parameters values given in their respective papers.

rVAD (V_{RVP}) [14]: implements two passes of denoising as enhancement to high energy speech regions: *i*) an *a-posteriori*

weighted energy difference measure, subjected to a pitch detection routine, used to classify speech from noise; *ii*) a spectral subtraction method for speech enhancement. The VAD stage uses a SNR weighted energy measure along with the pitch information to determine voice segments.

rVAD-Fast (V_{RVS}) [14]: a faster implementation which uses spectral flatness as a measure to identify the presence of pitch in a segment.

VAD-Wavlet (V_{DWT1} and V_{DWT2}) [27]: uses detail coefficient in wavelet based decomposition of speech. Daubechie’s wavelet (*dB3*) is used to derive details at multiple levels. Two different methods are implemented: *i*) V_{DWT1} , which uses RMS energy of details coefficients to discriminate between speech and noise, and *ii*) V_{DWT2} , which uses four energy based parameters.

VAD-Fusion (V_{FUS}) [28]: a MLP based voice/non-voice classifier implemented over a fusion of multiple spectral features derived by exploiting the spectro-temporal modulations, harmonicity, and long term spectral variability of signal.

VAD-LTSD (V_{LTSD}) [29]: compares long term spectral envelope characteristics of a segment against average noise spectrum. An adaptive threshold updates over noise power which is estimated after each non-speech segment is discovered.

GP-VAD (V_{GP}) [30]: a convolutional recurrent neural network (CRNN) trained on noisy log-Mel power spectrograms in a weakly supervised fashion using only clip-level labels.

VAD-TEO (V_{TEO}) [31]: highlights formant information using a band-spectral mass function, derived from a convex spectral energy function, used to compute the spectral entropy to determine voiced/unvoiced regions.

VAD-LSD (V_{LSD}) [29]: uses maximal spectrum information to model the contrast within spectral behavior across voicing and noise segments, and derives average noise spectral characteristic from silence and pause regions to obtain a divergence function.

VAD-LSE (V_{LSE}) [32]: compares the energy content within high and low spectral bands in the DFT spectra, knowing that voicing information is predominantly confined in the latter band.

4.3. Evaluation Metrics

VAD can be treated as a binary classification problem of sorting the input signal frames into voiced/non-voiced classes. To that end, frame-level results such as true positives (TP), false positives (FP), false negative (FN), and false positives (FP) can be used to compute standard classification metrics and thus measure a model’s performance over time. We use precision (P), recall (R), and F1-score, which are calculated as given in eq. (4).

$$P = \frac{TP}{TP + FP}; \quad R = \frac{TP}{TP + FN}; \quad F1 = 2 \cdot \frac{P \cdot R}{P + R} \quad (4)$$

5. Results and Discussion

Figure 4 shows the performance of the different VAD methods. V_{ZFF} refers to the method based on the first approach, presented in Section 3. $V_{ZFF-on-RVP}$ denotes the second approach, where the composite signal r_c is fed as input to rVAD algorithm *without application of the denoising routine*. The results show that V_{ZFF} , which uses minimal spectral information, outperforms V_{TEO} , V_{LSD} and V_{LSE} , and achieves a performance close to most of the other supervised and unsupervised methods, with the exception of V_{RVP} . It is also interesting to observe that $V_{ZFF-on-RVP}$, computed without the denoising routine, yields performance close to V_{RVP} . Furthermore, it can also be observed

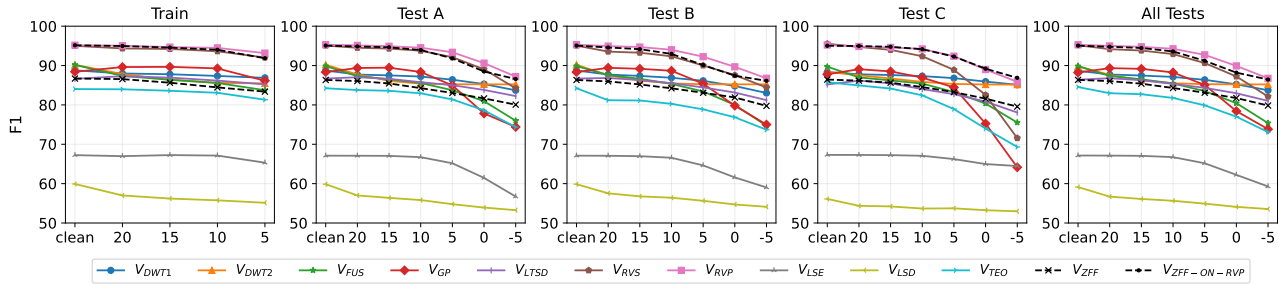


Figure 4: Performance of methods across all SNRs in different sets of the Aurora-2 database.

that at very high and very low SNRs, the V_{ZFF} and $V_{ZFF-ON-RVP}$ methods yield competing performances. Together these observations show that the outputs of zero-frequency filters indeed carry the source and system information in a reliable and robust manner, and can be effectively employed for VAD.

Table 1 shows the standard deviation of the F1-scores of each method, across all SNRs for the entire *Test* set. It can be observed that performance of V_{ZFF} remains invariant to added interferences across a range of SNRs (20 dB to -5 dB). Figure 4 shows that the performance of the method not only suffers very marginally at low SNRs, but in fact gives good results for the very low SNR value of -5 dB. This is particularly noticeable for the *Test C* set, where the channel characteristics are different than in the *Train* set (G.712). Nonetheless, V_{ZFF} outperforms V_{GP} , V_{RVS} , V_{FUS} , and V_{LTSD} .

Table 1: The standard deviation of the F1-scores ([%]) for each method, computed over the entire test set and across all SNRs. A high value indicates a significant variance and degradation with noise.

V_{DWT}	V_{LSD}	V_{LTSD}	V_{ZFF}	V_{LSE}	V_{RVP}	$V_{ZFF-ON-RVP}$	V_{TEO}	V_{RVS}	V_{FUS}	V_{GP}
1.6	1.7	2.0	2.2	2.8	3.0	3.2	3.7	4.3	4.5	5.7

Figure 5 shows the decision boundaries obtained for different methods on a given speech recording. It can be observed that the decision boundary produced by V_{ZFF} is successfully able to segment the speech signal into intervals which are significantly tighter than the other methods, as well as those given in the ground truth. In other words, the performance of our method is not as well reflected in the F1-scores compared to the other methods because of the broader boundaries of the ground truth segments, even though our method is able to provide a much more granular segmentation. The methods yielding higher performance comply closely to the broader VAD boundaries and consequently yield higher F1-scores, as can be noted in Figure 5 for V_{RVP} , V_{RVS} , V_{DWT1} , V_{DWT2} , and V_{GP} methods.

6. Conclusion

In this paper, we investigated modeling source and system information jointly using zero-frequency filtering technique for voice activity detection. In that direction, we proposed and validated two approaches for VAD on the Aurora-2 dataset with different noise, channel, and SNR conditions. Our investigations demonstrated that VAD can be effectively performed by combining the filter outputs together to compose a composite signal carrying f_0 , F_1 , and F_2 related information, and then applying a dynamic threshold after spectral entropy-based weighting (first approach), or else by passing the composite signal to an




Figure 5: Performance of some baseline methods, proposed method, and ground truth, for noisy speech (SNR = 10 dB).

other VAD (second approach). The experiments also illustrate that the proposed method produces more refined boundary demarcations for the VAD task compared to other supervised and unsupervised methods in the literature. It is also robust against degradation as well as channel characteristics, and yields stable performance across a range of SNRs. The first approach operates in the time domain and is relatively less complex to implement. The second approach illustrates that the composite signal, obtained by modulation of trend removal in the zero-frequency filtering, is an effective representation of speech characteristics, and can hence be used in conjunction with other VADs.

One of the main advantages of the proposed zero-frequency filtering based approach is that it does not explicitly assume any mathematical model for the produced speech signal in order to acquire source and system information. It can also thus be extended to other types of audio signals, such as animal and birds vocalizations. Our future work will focus in this direction, along with modeling the composite signal using the raw waveform neural network based modeling approach [33] for supervised voice activity detection [17].

7. Acknowledgement

This work was partially funded by the Swiss National Science Foundation (SNSF) through projects: NCCR Evolving language (grant agreement no. 51NF40.180888) and Towards Integrated processing of Physiological and Speech signals (TIPS) (grant agreement no. 200021.188754).

8. References

- [1] L. R. Rabiner and M. R. Sambur, "An algorithm for determining the endpoints of isolated utterances," *The Bell System Technical Journal*, vol. 54, pp. 297–315, 1975.

- [2] L. Lamel, L. Rabiner, A. Rosenberg, and J. Wilpon, "An improved endpoint detector for isolated word recognition," *IEEE Transactions on Acoustics, Speech, and Signal Processing*, vol. 29, no. 4, pp. 777–785, 1981.
- [3] B. Kingsbury, G. Saon, L. Mangu, M. Padmanabhan, and R. Sarikaya, "Robust speech recognition in Noisy Environments: The 2001 IBM spine evaluation system," in *Proc. of ICASSP*, vol. 1, 02 2002, pp. I-53–I-56.
- [4] T. T. Kristjansson, S. Deligne, and P. A. Olsen, "Voicing features for robust speech detection," in *Proc. of Interspeech*, 2005.
- [5] S. O. Sadjadi and J. H. L. Hansen, "Unsupervised Speech Activity Detection Using Voicing Measures and Perceptual Spectral Flux," *IEEE Signal Processing Letters*, vol. 20, no. 3, pp. 197–200, 2013.
- [6] E. Nemer, R. Goubran, and S. Mahmoud, "Robust voice activity detection using higher-order statistics in the LPC residual domain," *IEEE Transactions on Speech and Audio Processing*, vol. 9, no. 3, pp. 217–231, 2001.
- [7] L. Rabiner and M. Sambur, "Application of an LPC distance measure to the voiced-unvoiced-silence detection problem," *IEEE Transactions on Acoustics, Speech, and Signal Processing*, vol. 25, no. 4, pp. 338–343, 1977.
- [8] R. Tucker, "Voice activity detection using a periodicity measure," *IEE Proceedings I (Communications, Speech and Vision)*, vol. 139, pp. 377–380(3), August 1992.
- [9] I. McCowan, D. Dean, M. McLaren, R. Vogt, and S. Sridharan, "The Delta-Phase Spectrum With Application to Voice Activity Detection and Speaker Recognition," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 19, no. 7, pp. 2026–2038, 2011.
- [10] J. Haigh and J. S. Mason, "A Voice Activity Detector Based On Cepstral Analysis," in *Proc. of European Conference Speech Communication and Technology*, 1993, pp. 1103–1106.
- [11] J. Haigh and J. Mason, "Robust voice activity detection using cepstral features," in *Proceedings of TENCON '93. IEEE Region 10 International Conference on Computers, Communications and Automation*, vol. 3, 1993, pp. 321–324 vol.3.
- [12] S. O. Sadjadi and J. H. Hansen, "Unsupervised Speech Activity Detection Using Voicing Measures and Perceptual Spectral Flux," *IEEE Signal Processing Letters*, vol. 20, no. 3, pp. 197–200, 2013.
- [13] M. H. Moattar and M. M. Homayounpour, "A simple but efficient real-time Voice Activity Detection algorithm," in *Proc. of European Signal Processing Conference (EUSIPCO)*, 2009, pp. 2549–2553.
- [14] Z.-H. Tan, A. kr. Sarkar, and N. Dehak, "rVAD: An unsupervised segment-based robust voice activity detection method," *Computer Speech & Language*, vol. 59, pp. 1–21, 2020.
- [15] J. Sohn and W. Sung, "A voice activity detector employing soft decision based noise spectrum adaptation," in *Proc. of ICASSP*, vol. 1. IEEE, 1998, pp. 365–368.
- [16] X.-L. Zhang and J. Wu, "Deep belief networks based voice activity detection," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 21, no. 4, pp. 697–710, 2012.
- [17] R. Zazo, T. N. Sainath, G. Simko, and C. Parada, "Feature Learning with Raw-Waveform CLDNNs for Voice Activity Detection," in *Proc. of Interspeech*, 2016, pp. 3668–3672.
- [18] N. Wilkinson and T. Niesler, "A hybrid CNN-BiLSTM voice activity detector," in *Proc. of ICASSP*, Toronto, Canada, 2021.
- [19] K. S. R. Murty and B. Yegnanarayana, "Epoch Extraction From Speech Signals," *IEEE Transactions on Audio, Speech, and Language Processing*, vol. 16, no. 8, pp. 1602–1613, 2008.
- [20] R. Prasad and M. Magimai-Doss, "Identification of F1 and F2 in Speech Using Modified Zero Frequency Filtering," in *Proc. of Interspeech*, 2021, pp. 56–60.
- [21] S. R. Kadiri and P. Alku, "Excitation Features of Speech for Speaker-Specific Emotion Detection," *IEEE Access*, vol. 8, pp. 60 382–60 391, 2020.
- [22] R. Prasad, G. Yilmaz, O. Chetelat, and M. Magimai.-Doss, "Detection of S1 And S2 Locations In Phonocardiogram Signals Using Zero Frequency Filter," in *Proc. of ICASSP*, 2020, pp. 1254–1258.
- [23] D. Pearce and H.-G. Hirsch, "The AURORA experimental framework for the performance evaluations of speech recognition systems under noisy condition," in *Proc. of ICSLP*, vol. 4, 01 2000, pp. 29–32.
- [24] R. Leonard, "A database for speaker-independent digit recognition," in *Proc. of ICASSP*, vol. 9. IEEE, 1984, pp. 328–331.
- [25] A. Varga and H. J. M. Steeneken, "Assessment for automatic speech recognition: II. NOISEX-92: A database and an experiment to study the effect of additive noise on speech recognition systems." *Speech Communication*, vol. 12, no. 3, pp. 247–251, 1993.
- [26] S. J. Young, D. Kershaw, J. Odell, D. Ollason, V. Valtchev, and P. Woodland, *The HTK Book*. Cambridge University Press, 2006.
- [27] "Voice activity Detection Using Wavelet Transform," https://github.com/pvenuprasad/VAD_wavelet, 2019, accessed in 29. 03. 2022.
- [28] M. Van Segbroeck, A. Tsiartas, and S. Narayanan, "A robust frontend for VAD: exploiting contextual, discriminative and spectral cues of human voice," in *Proc. of Interspeech*, 2013, pp. 704–708.
- [29] J. Ramirez, J. C. Segura, C. Benitez, A. De La Torre, and A. Rubio, "Efficient voice activity detection algorithms using long-term speech information," *Speech communication*, vol. 42, no. 3-4, pp. 271–287, 2004.
- [30] H. Dinkel, Y. Chen, M. Wu, and K. Yu, "Voice Activity Detection in the Wild via Weakly Supervised Sound Event Detection," in *Proc. of Interspeech*, 2020, pp. 3665–3669.
- [31] R. Hegde and R. Muralishankar, "Voice Activity Detection Using Novel Teager Energy Based Band Spectral Entropy," in *Proc. of International Conference on Communication and Electronics Systems (ICCES)*. IEEE, 2019, pp. 1272–1278.
- [32] J. Pang, "Spectrum energy based voice activity detection," in *Proc. of 7th Annual Computing and Communication Workshop and Conference*. IEEE, 2017, pp. 1–5.
- [33] D. Palaz, R. Collobert, and M. Magimai.-Doss, "Estimating Phoneme Class Conditional Probabilities from Raw Speech Signal using Convolutional Neural Networks," in *Proc. of Interspeech*, 2013.