



Auto-Landmark: Acoustic Landmark Dataset and Open-Source Toolkit for Landmark Extraction

Xiangyu Zhang¹, Daijiao Liu¹, Tianyi Xiao³, Cihan Xiao², Tuende Szalay¹, Mostafa Shahin¹, Beena Ahmed¹, Julien Epps¹

¹The University of New South, ²Johns Hopkins University, ³Nanyang Technological University,
xiangyu.zhang2@unsw.edu.au

Abstract

In the speech signal, acoustic landmarks identify times when the acoustic manifestations of the linguistically motivated distinctive features are most salient. Acoustic landmarks have been widely applied in various domains, including speech recognition, speech depression detection, clinical analysis of speech abnormalities, and the detection of disordered speech. However, there is currently no dataset available that provides precise timing information for landmarks, which has been proven to be crucial for downstream applications involving landmarks. In this paper, we selected the most useful acoustic landmarks based on previous research and annotated the TIMIT dataset with them, based on a combination of phoneme boundary information and manual inspection. Moreover, previous landmark extraction tools were not open source or benchmarked, so to address this, we developed an open source Python-based landmark extraction tool and established a series of landmark detection baselines. The first of their kinds, the dataset with landmark precise timing information, landmark extraction tool and baselines are designed to support a wide variety of future research.

Index Terms: Acoustic Landmark, Speech Processing

1. Introduction

In speech processing, frame-based methods are most commonly used to segment the speech waveform. This method treats each frame as the central unit of analysis, with a fixed set of speech attributes measured at each frame, has been used in many domain [1, 2]. However, because this approach relies on fixed-duration processing, it often overlooks important timing factors such as speaking rate and segmental duration. In contrast, acoustic landmark detection focuses on specific, acoustically significant points in the speech signal [3], as illustrated in Figure 1. This method allows for the identification of acoustic landmarks somewhat independently of frames, while also providing valuable timing information for subsequent processing.

Acoustic landmarks have been proven highly effective across various fields. Initially applied in speech recognition [3, 4], they have since been extended to the health domain, including applications such as depression detection [5, 6], clinical analysis of speech abnormalities [7], and the detection of disordered speech [8]. Despite the success of acoustic landmarks in various fields, no dataset currently provides precise ground truth timing information for the occurrence of acoustic landmarks, making it difficult to produce standardized results on a baseline dataset [5, 8]. To address this research gap, we selected the most useful acoustic landmarks based on previous research [6–10] and annotated the TIMIT dataset [11] with them, with some landmarks labeled according to phoneme boundary information and others manually annotated.




Figure 1: Example of acoustic landmarks: The speech signal is discretized into a series of tokens with speech production information

Moreover, previous landmark extraction tools were based on closed-source software, making it difficult to understand their inner workings, and had not been benchmarked. To overcome this limitation, we developed an open source Python-based acoustic landmark extraction tool - Auto-Landmark and established a series of acoustic landmark detection baselines to support future research.

2. Related Work

2.1. Acoustic Landmarks

The idea of acoustic landmarks originates from studies on distinctive features [12]. Distinctive features offer a concise way to describe the sounds of a language at a subsegmental level. They are closely related to both acoustics and articulation, characterized by binary values. This minimal set of features allows for the differentiation of each segment within a language. Landmarks serve as indicators of the underlying segments, which group distinctive features into bundles. Landmarks are categorized into four groups: abrupt-consonantal, abrupt, non-abrupt, and vocalic. While scholars have slightly different definitions of acoustic landmarks, [13] expanded upon [3] by releasing a MATLAB-based landmark detection toolkit (SpeechMark), which has become the most widely used approach for landmark detection.

Table 1: Description of the five landmarks investigated.

Landmark	Description
g	vibration of vocal folds start (+) or end (-)
b	onset (+) or offset (-) of existence of turbulent noise during obstruent regions
s	releases (+) or closures (-) of a nasal
v	voiced frication onset (+) or offset (-)
f	frication onset (+) or offset (-)

2.2. Application of Acoustic Landmarks

Acoustic landmarks were initially used in speech recognition [3,4], with some researchers exploring the combination of acoustic landmarks with CTC in this context [4]. Following their success in speech recognition, acoustic landmarks have been applied to various other domains. Acoustic landmarks have been widely applied in various health-related studies. For example, [5,6] utilized acoustic landmark timing information for depression detection. [9,14,15] further extended the use of acoustic landmarks by tokenizing speech for large language model-based systems in depression detection. [8] used acoustic landmark time duration to describe abnormalities in speech production that affect a speaker’s intelligibility. [7] evaluated the feasibility of differentiating conversational and clear speech produced by individuals with muscle tension dysphonia (MTD) using landmark-based analysis of speech.

3. Acoustic Landmark Dataset

3.1. Dataset and Landmark Selection

As shown in Table 1, we selected five different types of acoustic landmarks based on previous research [5–10], each with the onset and offset states. These landmarks are g(lottis), s(sonorant), f(ricative), v(voiced fricative), and b(ursts), which are pointing in time where different abrupt articulatory events occur.

The **g (glottis)** landmark captures the start or end of the vocal fold vibration, which is crucial in identifying voiced phonemes where the vocal cords are active. The **b (bursts)** landmark marks the onset or offset of turbulent noise during obstruent regions, such as plosives or stops, which are significant for identifying consonant sounds. The **s (sonorant)** landmark indicates the releases or closures of nasal sounds and glides [l], which are produced with a continuous airflow through the vocal tract. Proper identification of these sonorants is important for understanding the nasal quality in speech. The **v (voiced fricative)** landmark captures the onset or offset of voiced frication, with voiced fricatives like [v] and [z] involving continuous vibration of the vocal cords while producing a turbulent airflow. Lastly, the **f (fricative)** landmark indicates the onset or offset of frication, with fricatives like [f] and [s] characterized by the turbulent airflow created by a narrow constriction in the vocal tract [16,17].

3.2. Labeling Convention

As discussed in the previous section, landmarks are closely associated with phonemes, which makes the TIMIT dataset [11] with its manually annotated phoneme boundaries an ideal

Table 2: Rules for Annotating Data

Landmark	Details
v+	Start time of voiced frication phone
v-	End time of voiced frication phone
f+	Start time of frication phone
f-	End time of frication phone
s+	Start time of nasal or [l] phone
s-	End time of nasal or [l] phone
b+	Sudden energy increase in stop and affricate in the spectrum
b-	Sudden energy decrease in stop and affricate in the spectrum
g+	Start time of voiced phones (vowels, voiced fricatives, nasals, b, d, g)
g-	End of glottal phone, or change to voiceless stops, fricatives, or affricates

choice for annotation. We utilized Praat [18] to facilitate the labeling process. We first identified the boundaries of the phonemes associated with landmark, then used Praat to help determine the exact timing of these landmarks.

Table 2 outlines the rules we followed for annotating acoustic landmarks. Each landmark corresponds to specific phonetic events and transitions within the speech signal. The v+ and v- landmarks indicate the start and end times of voiced fricatives, such as [z], [zh], [v], and [dh]. Similarly, f+ and f- mark the start and end times of fricatives like [s], [sh], [f], and [th]. For nasals and the glides [l], the s+ and s- landmarks represent the start and end times, covering phonemes like [m], [n], [ng], [em], [en], [eng], and [nx]. The b+ and b- landmarks capture the sudden energy changes in stops and affricates, including [b], [d], [g], [p], [t], [k], [dx], [q], [jh], and [ch] and manually labeled. Finally, the g+ and g- landmarks denote the start and end times of voiced phones, encompassing a range of vowels (e.g., [iy], [ih], [eh], [ey], [ae], [aa], [aw], [ay], [ah], [ao], [oy], [ow], [uh], [uw], [ux], [er], [ax], [ix], [axr], [ax-h]), voiced fricatives, nasals, and stops such as [b], [d], and [g]. The g- landmark also indicates the transition from voiced to voiceless sounds.

3.3. Statistics of Landmark Dataset

Figure 2 illustrates the proportion of each landmark in the dataset. As shown, the g landmark has the highest proportion in both the training and test sets. However, the proportion of the b landmark in the test set is significantly lower than in the training set. We found that this discrepancy is due to the lower proportion of phonemes associated with the b landmark in the test set compared with the train set.

4. Landmark Detection System

We provide both signal processing and deep learning methods for landmark detection. For the former, we developed an open-source Python tool to address the gap in fully Python based open-source solutions. For the latter, we advanced the field by incorporating landmark detection into cutting-edge research




Figure 2: Landmark distribution in the TIMIT dataset [11] for the training and test sets. Right is Landmark distribution in the training set, left is Landmark distribution in the test set

and provided an ESPnet [19] recipe to facilitate this.

4.1. Filter-based Method

Our implementation of the signal processing-based method is based on [3, 9, 13]. Initially, the spectrogram is divided into six frequency bands. Landmarks are identified through energy changes within these bands using a two-pass strategy. Different landmarks are determined by analyzing either a single band or a combination of multiple bands.

The **g landmark** is identified when both the coarse and fine filters exhibit a peak in band 1. The **b landmark**, in an unvoiced segment (not between +g and the next -g), is identified if at least three out of five frequency bands show simultaneous power increases of no less than 6 dB in both coarse and fine filters. The **s landmark** is identified in a voiced segment (between +g and the next -g) under the same conditions: simultaneous power increases of no less than 6 dB in at least three out of five frequency bands in both coarse and fine filters. The **f+** and **v+** landmarks are detected by identifying a 6 dB power increase in at least three high-frequency bands (4, 5, 6) and a power decrease in low-frequency bands (2, 3). Conversely, for **f-** and **v-**, the criteria involve a 6 dB power decrease in the high-frequency bands and a power increase in the low-frequency bands. The key distinction is that frication landmarks are detected in unvoiced segments (b landmark), while voiced frication landmarks are detected in voiced segments (s landmark). The details of the method can be found in [3, 9, 13].

4.2. Deep Learning Based Method

To date, landmark detection has not been performed with deep learning models. Our first consideration is how to model this problem. In this paper, since both landmark detection and speech recognition tasks involve converting speech signals into symbols, we approach landmark detection as a speech recognition task. We employ a Hybrid CTC/attention end-to-end architecture [20]. We utilized two types of models as the Encoder: the Conformer [21] and ConExtBimamba [22–24]. For the decoder, we employed the Transformer model [25]. Given the success of self-supervised models in various speech tasks [26], we therefore explored how they perform on the landmark detection task. Specifically, we utilized the wav2vec2.0 large model [27]. We followed the Superb setup [28], using the weighted sum feature from different layers of wav2vec2.0.

4.3. Toolkit for Landmark Extraction

Previous landmark extraction tools were not open source, making it difficult to understand their inner workings. To address

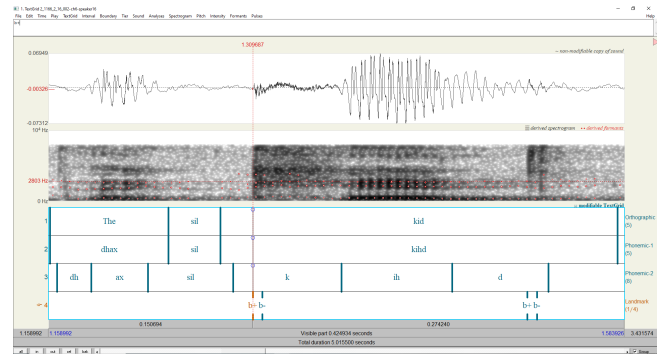


Figure 3: Example of labeling Using Praat: the b landmark is marked when the energy suddenly changes

this, we developed an open-source Python-based tool, **Auto-Landmark**, for landmark extraction.

C.1) Directory Structure

Auto-Landmark adopts a structure similar to the directory organization of ESPnet [19]. There are four main folders: **egs**, which is organized Recipe for each dataset; **methods**, which stores different landmark extraction methods; **utils**, containing commonly used functions; and **visualisation**, which houses functions related to visualization. There are two datasets now supported by Auto-landmark, TIMIT [11] and DAIC-WOZ [29]

C.2) Toolkit workflow

Stage 1-2 Data preparation: The initial data format follows the Kaldi style. After preprocessing, all speech data is converted into a standard format for resampling and segmentation.

Stage 3 Landmark Extraction: At this stage, the primary task is to execute different algorithms for landmark extraction. So far, we have provided two distinct types of landmarks: basic and advanced. The main difference between the two lies in the implementation of the filters and the smoothing methods. The **basic method** applies moving window average smoothing, a simpler approach that focuses on noise reduction in energy signals. This method is effective for peak detection but offers less flexibility and control, making it more suitable for straightforward tasks. On the other hand, the **advanced method** utilizes linear convolution smoothing with customizable kernels, such as Hanning windows, providing greater control and adaptability for handling more complex data. For more details on their differences, please refer to the code.

Optional Visualization: We have provided additional visualization tools to help health researchers in analyzing cases.

5. Results and Discussion

5.1. Experimental Setup

Signal Processing Method: Our two-pass filter utilizes six band energies, specifically in the ranges of 0.0–0.4, 0.8–1.5, 1.2–2.0, 2.0–3.5, 3.5–5.0, and 5.0–8.0 kHz. Following previous work [3, 9, 13], the coarse processing uses 20ms smoothing with an 8 dB threshold for peak detection. In fine processing, 10ms smoothing is applied, and peak detection targets thresholds within the range of 5 dB to 8 dB.

Deep Learning Method: We conducted our experiments using a V100 32G GPU, with the model configuration based on the best results from the AN4 dataset [30] in ESPnet, given that both TIMIT and AN4 are small datasets. The training utilized a single V100 32G GPU with batch bins of 1,000,000, using

Table 3: Landmark detection results for different methods

Method	LER(↓)
Singal Processing Method	
– SpeechMark [13]	56.53
– Auto-Landmark	56.73
Deep Learning Method	
– Conformer [21]	33.0
– ConBiMamba [22]	31.3
– Wav2vec2 + Conformer	48.3
– Wav2vec2 + ConBiMamba	44.6

the Adam optimizer with a learning rate of 0.001 and a warm-up scheduler for 2500 steps. The Conformer/ConExtBimamba encoder was configured with an output size of 256, 2048 linear units, and 6 blocks. The Transformer decoder had 4 attention heads, 2048 linear units, 6 blocks, and a dropout rate of 0.1.

5.2. Baseline Result

Table 3 compares the landmark detection results from the two methods, evaluated by Landmark Error Rate (LER). LER was calculated based on patterns without considering landmark timing. Additionally, following previous conventions, we did not differentiate between onset (+) and offset (-) landmarks [3]. All the deep learning models produced lower LER compared to both signal processing methods, with the ConExtBiMamba model outperforming the others. Typically, the overall performance of models improves when using self-supervised models. However, as shown in Table 3, the performance of both the Conformer and ConExtBiMamba models decreased when combined with wav2vec2, contrary to expectations. This issue arose because several landmarks occur in close temporal proximity to each other, even simultaneously. This means that numerous landmarks can appear within a 10ms window. However, as wav2vec uses a 20ms window with a 10ms stride, it potential misses some landmarks.

5.3. Insights from Self-supervised Weights

Previous research has shown that when using the weighted sum feature of self-supervised models for downstream tasks, the weights can be used to analyze the properties of the models and tasks [31]. Figure 4 shows the feature weights from different layers when using wav2vec2 large combined with the Conformer for landmark detection. Layers 16, 17, and 18 have the highest weights. Previous studies using CCA and mutual information to analyze wav2vec2 found that phonetic information is most salient around layer 18 [32]. This suggests that wav2vec2 treats landmark detection as a task highly related to phonetic information. However, as described by the definition of landmarks, they primarily denote points of abrupt change in the speech signal. Thus, low-level features in speech are also crucial. Due to the limitations of wav2vec’s window size, many consecutive abrupt changes are not captured. From Figure 4, we can see that, unlike the layers associated with phonetic information identified by CCA analysis, the weights of shallow features remain significant for landmark detection. According to [33], the shallow features of wav2vec are more similar to low-level

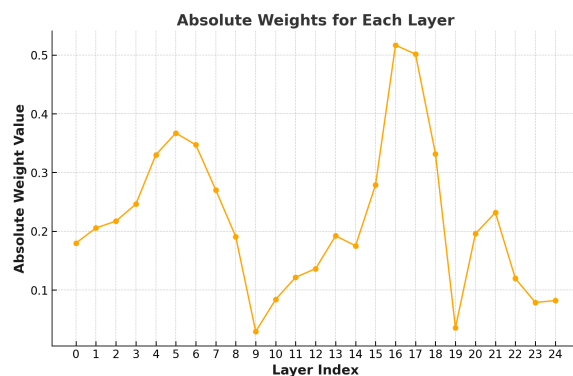


Figure 4: Layer Weight of wav2vec2.0 for Landmark Detection

features, indicating that the model captures effects that are not solely related to phonetic information.

6. Conclusion

In this study, we made several key contributions to the field of acoustic landmark detection. First, we provided the first publicly available dataset with manually annotated acoustic landmarks and precise timing ground truth. Second, we also developed the first Python-based open-source implementation of the landmark detection methods originally proposed by [3]. Finally, we established new baselines for landmark detection with deep learning models and comprehensively benchmarked these methods. Our analysis demonstrated that effective landmark detection requires both phonetic and low-level feature information. This dataset and the accompanying tools will significantly benefit research into conditions that affect speech timing, such as dementia, Parkinson’s, intoxication, speech impairment, and language learning. Previous research has shown that timing information is crucial for the application of landmarks in mental health [5, 10], making our timing-based benchmarking particularly valuable. Moreover, since CTC-based models struggle to capture timing information, developing deep learning models that can more precisely estimate the timing of landmarks will be an important direction for future research.

7. Acknowledgments

This work was supported by the Australian Research Council Discovery Project DP230101184.

8. References

- [1] H. Li, J. Q. Yip, T. Fan, and E. S. Chng, “Speech enhancement using continuous embeddings of neural audio codec,” in *ICASSP 2025-2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2025, pp. 1–5.
- [2] H. Li, Y. Hu, C. Chen, S. M. Siniscalchi, S. Liu, and E. S. Chng, “From kan to gr-kan: Advancing speech enhancement with kan-based methodology,” *arXiv preprint arXiv:2412.17778*, 2025.
- [3] S. A. Liu, “Landmark detection for distinctive feature-based speech recognition,” *The Journal of the Acoustical Society of America*, vol. 100, no. 5, pp. 3417–3430, 1996.
- [4] D. He, X. Yang, B. P. Lim, Y. Liang, M. Hasegawa-Johnson, and D. Chen, “When ctc training meets acoustic landmarks,” in *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2019, pp. 5996–6000.

- [5] Z. Huang, J. Epps, and D. Joachim, "Investigation of speech landmark patterns for depression detection," *IEEE transactions on affective computing*, vol. 13, no. 2, pp. 666–679, 2019.
- [6] Z. Huang, J. Epps, D. Joachim, and M. Chen, "Depression detection from short utterances via diverse smartphones in natural environmental conditions." in *INTERSPEECH*, 2018, pp. 3393–3397.
- [7] K. Ishikawa, M. Pietrowicz, S. Charney, and D. Orbelo, "Landmark-based analysis of speech differentiates conversational from clear speech in speakers with muscle tension dysphonia," *JASA Express Letters*, vol. 3, no. 5, 2023.
- [8] K. Ishikawa, J. MacAuslan, and S. Boyce, "Toward clinical application of landmark-based speech analysis: Landmark expression in normal adult speech," *The Journal of the Acoustical Society of America*, vol. 142, no. 5, pp. EL441–EL447, 2017.
- [9] X. Zhang, H. Liu, K. Xu, Q. Zhang, D. Liu, B. Ahmed, and J. Epps, "When llms meets acoustic landmarks: An efficient approach to integrate speech into large language models for depression detection," *arXiv preprint arXiv:2402.13276*, 2024.
- [10] Z. Huang, J. Epps, D. Joachim, and V. Sethu, "Natural language processing methods for acoustic and landmark event-based features in speech-based depression detection," *IEEE Journal of selected topics in Signal Processing*, vol. 14, no. 2, pp. 435–448, 2019.
- [11] J. S. Garofolo, "Timit acoustic phonetic continuous speech corpus," *Linguistic Data Consortium*, 1993, 1993.
- [12] P. L. Garvin, "Preliminaries to speech analysis: The distinctive features and their correlates," 1953.
- [13] S. Boyce, H. J. Fell, and J. MacAuslan, "Speechmark: Landmark detection tool for speech analysis." in *INTERSPEECH*, 2012, pp. 1894–1897.
- [14] X. Zhang, B. Ahmed, and J. Epps, "Why pre-trained models fail: Feature entanglement in multi-modal depression detection," *arXiv preprint arXiv:2503.06620*, 2025.
- [15] X. Zhang, H. Liu, Q. Zhang, B. Ahmed, and J. Epps, "Speecht-rag: Reliable depression detection in llms with retrieval-augmented generation using speech timing information," *arXiv preprint arXiv:2502.10950*, 2025.
- [16] G. S. Nathan, "The sounds of the world's languages," 1998.
- [17] W. Culbertson, *Fundamentals of the Speech and Language Sciences*. Routledge, 2024.
- [18] P. Boersma and V. Van Heuven, "Speak and unspeak with praat," *Glott International*, vol. 5, no. 9/10, pp. 341–347, 2001.
- [19] S. Watanabe, T. Hori, S. Karita, T. Hayashi, J. Nishitoba, Y. Unno, N. Enrique Yalta Soplín, J. Heymann, M. Wiesner, N. Chen, A. Renduchintala, and T. Ochiai, "ESPnet: End-to-end speech processing toolkit," in *Proceedings of Interspeech*, 2018, pp. 2207–2211. [Online]. Available: <http://dx.doi.org/10.21437/Interspeech.2018-1456>
- [20] S. Watanabe, T. Hori, S. Kim, J. R. Hershey, and T. Hayashi, "Hybrid ctc/attention architecture for end-to-end speech recognition," *IEEE Journal of Selected Topics in Signal Processing*, vol. 11, no. 8, pp. 1240–1253, 2017.
- [21] A. Gulati, J. Qin, C.-C. Chiu, N. Parmar, Y. Zhang, J. Yu, W. Han, S. Wang, Z. Zhang, Y. Wu *et al.*, "Conformer: Convolution-augmented transformer for speech recognition," *arXiv preprint arXiv:2005.08100*, 2020.
- [22] X. Zhang, Q. Zhang, H. Liu, T. Xiao, X. Qian, B. Ahmed, E. Ambikairajah, H. Li, and J. Epps, "Mamba in speech: Towards an alternative to self-attention," *IEEE Transactions on Audio, Speech and Language Processing*, 2025.
- [23] X. Zhang, J. Ma, M. Shahin, B. Ahmed, and J. Epps, "Rethinking mamba in speech processing by self-supervised models," in *ICASSP 2025-2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2025, pp. 1–5.
- [24] M. Chen, Q. Zhang, M. Wang, X. Zhang, H. Liu, E. Ambikairajah, and D. Chen, "Selective state space model for monaural speech enhancement," *IEEE Transactions on Consumer Electronics*, 2025.
- [25] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention is all you need," *Advances in neural information processing systems*, vol. 30, 2017.
- [26] J. Yao, Y. Yang, Y. Pan, Y. Feng, Z. Ning, J. Ye, H. Zhou, and L. Xie, "Fine-grained preference optimization improves zero-shot text-to-speech," *arXiv preprint arXiv:2502.02950*, 2025.
- [27] A. Baevski, Y. Zhou, A. Mohamed, and M. Auli, "wav2vec 2.0: A framework for self-supervised learning of speech representations," *Advances in neural information processing systems*, vol. 33, pp. 12 449–12 460, 2020.
- [28] S. wen Yang, P.-H. Chi, Y.-S. Chuang, C.-I. J. Lai, K. Lakhotia, Y. Y. Lin, A. T. Liu, J. Shi, X. Chang, G.-T. Lin, T.-H. Huang, W.-C. Tseng, K. tik Lee, D.-R. Liu, Z. Huang, S. Dong, S.-W. Li, S. Watanabe, A. Mohamed, and H. yi Lee, "SUPERB: Speech Processing Universal PERFORMANCE Benchmark," in *Proc. Interspeech 2021*, 2021, pp. 1194–1198.
- [29] D. DeVault, R. Artstein, G. Benn, T. Dey, E. Fast, A. Gainer, K. Georgila, J. Gratch, A. Hartholt, M. Lhomme *et al.*, "Simsensei kiosk: A virtual human interviewer for healthcare decision support," in *Proceedings of the 2014 international conference on Autonomous agents and multi-agent systems*, 2014, pp. 1061–1068.
- [30] A. Acero and R. M. Stern, "Environmental robustness in automatic speech recognition," in *International Conference on Acoustics, Speech, and Signal Processing*. IEEE, 1990, pp. 849–852.
- [31] S. S. Li, B. Xu, X. Zhang, H. Liu, W. Chao, and P. Garcia, "A quantitative approach to understand self-supervised models as cross-lingual feature extractors," in *Proceedings of the 6th International Conference on Natural Language and Speech Processing (ICNLSP 2023)*, 2023, pp. 200–211.
- [32] A. Pasad, J.-C. Chou, and K. Livescu, "Layer-wise analysis of a self-supervised speech representation model," in *2021 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*. IEEE, 2021, pp. 914–921.
- [33] A. Pasad, B. Shi, and K. Livescu, "Comparative layer-wise analysis of self-supervised speech models," in *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2023, pp. 1–5.