

Zero-Shot Automatic Pronunciation Assessment

Hongfu Liu, Mingqian Shi, Ye Wang

School of Computing, National University of Singapore, Singapore

{hongfu,m-shi,wangye}@comp.nus.edu.sg

Abstract

Automatic Pronunciation Assessment (APA) is vital for computer-assisted language learning. Prior methods rely on annotated speech-text data to train Automatic Speech Recognition (ASR) models or speech-score data to train regression models. In this work, we propose a novel zero-shot APA method based on the pre-trained acoustic model, HuBERT. Our method involves encoding speech input and corrupting them via a masking module. We then employ the Transformer encoder and apply k-means clustering to obtain token sequences. Finally, a scoring module is designed to measure the number of wrongly recovered tokens. Experimental results on *speechocean762* demonstrate that the proposed method achieves comparable performance to supervised regression baselines and outperforms non-regression baselines in terms of Pearson Correlation Coefficient (PCC). Additionally, we analyze how masking strategies affect the performance of APA.

Index Terms: automatic pronunciation assessment, zero-shot learning, self-supervised learning, HuBERT

1. Introduction

Learning a second language (L2) is a common requirement in bilingual or multilingual communities. However, L2 learners often struggle with achieving good proficiency in pronunciation. Computer-assisted pronunciation training (CAPT) is a notable application that enables language learners to effectively learn the pronunciation of new languages [1, 2]. CAPT provides feedback containing evaluation results, which can be automatically generated based on pronunciation, facilitating L2 learners in adjusting their pronunciation for improvement. Therefore, providing an overall assessment of pronunciation automatically is one of the primary objectives of CAPT.

Automatic pronunciation assessment has been extensively investigated over a prolonged period. Existing pronunciation assessment methods are implemented in the supervised setting. These approaches involve the usage of collected speech data with text annotations for training ASR models. Then the evaluation can be conducted based on the recognition results of ASR models. Goodness of Pronunciation (GoP) is one of the most commonly used metrics, aiming to provide phoneme-level scores for a given utterance. GoP requires calculating the log-posterior probability for each reference phoneme based on the contextual information [3, 4, 5]. On the other hand, there is an alternative research line that involves using speech data from non-native speakers with pronunciation scores annotated by domain experts to train regression models. Various features of speech data have been explored in this line, one of which is the phone-level features of speech [6, 7]. To enhance regression performance, [8] propose to use deep features transferred from

the acoustic models of ASR. Using speech representations of pre-trained acoustic models such as *wav2vec 2.0* or *HuBERT* also contributes to improving the regression performance by fine-tuning [9, 10]. Furthermore, multi-aspect pronunciation assessment at multiple granularities [11, 12] has been explored with multi-task supervised learning. However, there is a lack of unsupervised assessment approaches in the literature. All current pronunciation assessment methods require supervised signals to obtain the evaluation results.

Resource-efficient methods have been widely investigated for the low-resource scenario in the speech community [13, 14]. Nevertheless, it remains challenging to evaluate the quality of pronunciation using few or no data samples. Recent advances in Self-Supervised Learning (SSL) pre-trained language models (PLMs) have demonstrated strong few-shot and zero-shot learning abilities in the natural language processing community [15, 16] due to the knowledge acquired during the pre-training stage. PLMs are capable of performing downstream tasks via appropriate prompting with limited or even no data samples. However, the zero-shot ability has not been fully explored for SSL pre-trained acoustic models. This is because they learn at the acoustic level and it is challenging to learn linguistic representations from raw audio [17, 18], making it difficult to adapt them to downstream tasks without fine-tuning. While fine-tuning SSL pre-trained acoustic models with supervised data has been shown to be effective in automatic pronunciation assessment [9, 10], zero-shot pronunciation assessment has yet to be explored. Nevertheless, the acoustic-level knowledge acquired by SSL pre-trained acoustic models presents a viable option for zero-shot pronunciation assessment based on the unlabelled speech data observed during pre-training.

In this work, we propose a zero-shot pronunciation assessment approach that requires no annotated speech data. This is achieved by leveraging the SSL pre-trained acoustic model, *HuBERT* [19], for conducting the masked token prediction task. Our method involves encoding the waveform speech input into frame sequences and transforming them into corrupted sequences via a masking module. We then employ the Transformer Encoder of *HuBERT* and apply k-means clustering to obtain tokens of frame sequences and recovered tokens of corrupted sequences. Finally, a scoring module is designed to evaluate the pronunciation of a given speech by measuring the number of wrongly recovered tokens. Our proposed method is unsupervised and requires no fine-tuning. We conduct experiments on the *speechocean762* dataset [7]. The experimental results demonstrate that the proposed method achieves comparable performance compared to supervised baselines and outperforms non-regression baselines in terms of the Pearson Correlation Coefficient.

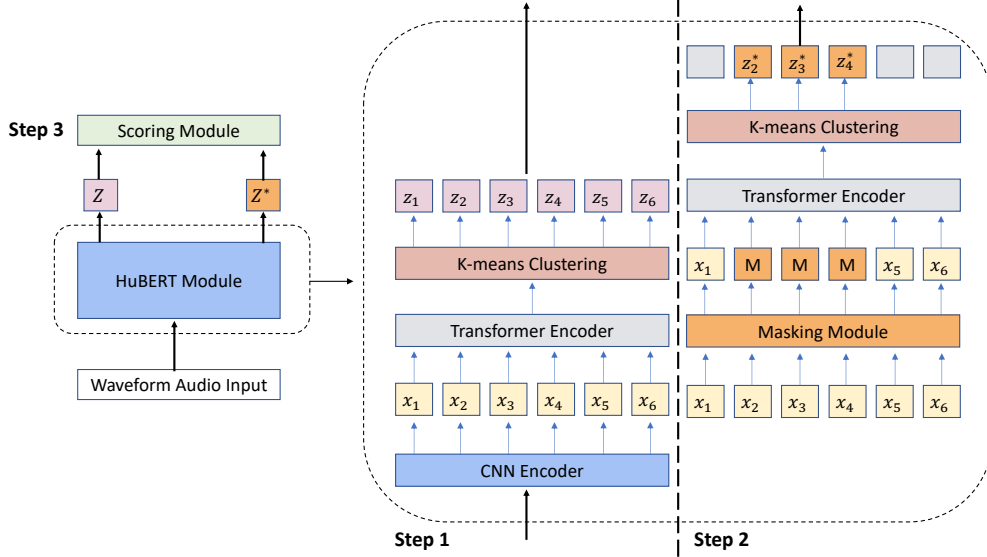


Figure 1: Overview of the zero-shot automatic pronunciation assessment.

2. Method

2.1. Overview

An overview of our proposed method is shown in Figure 1. We developed three main steps to achieve this assessment as shown in Figure 1. The first step is to input the waveform speech audio to the convolutional neural network (CNN) encoder to get a frame sequence. The Transformer encoder takes as input the frame sequence and the k-means clustering is utilized to obtain the token sequences. The second step is to apply a mask module on the frame sequence gained in Step 1 and input the masked sequence to the Transformer encoder followed by the k-means clustering to obtain the recovered tokens of masked spans. Finally, a scoring module is employed to measure the number of wrongly recovered tokens based on the outputs of Step 1 and Step 2. The intuition behind this is that for well-pronounced speech, recovered tokens would be similar to tokens of corresponding positions obtained from uncorrupted input. On the contrary, for mispronunciation speech, recovered tokens would differ from their counterparts a lot.

2.2. HuBERT Module

The HuBERT Module is adapted from the original HuBERT architecture [19]. This module consists of one CNN encoder, one Transformer encoder, one k-means clustering, and one masking module. Let $X = \{x_1, \dots, x_T\}$ denote the output of the CNN encoder with T frames. Then the Transformer encoder is employed to get latent representations of X that are further utilized to obtain the token sequences $Z = \{z_1, \dots, z_T\}$ through k-means clustering, where $z_t \in [C]$ is a C -class categorical variable and $t \in [T]$. z_t is also known as the hidden acoustic unit.

2.3. Masking Module

To construct the masked token prediction task, we employ a masking strategy r on X . If the set of indices to be masked is denoted by $M \subset [T]$ for a sequence X with T frames, then the corrupted version is denoted by $X^* = r(X, M)$, where

x_m is replaced by a mask embedding x^* for $m \in M$. Then we feed X^* into the same Transformer encoder and use the same k-means clustering. As a consequence, the output token sequence of masked spans is denoted by $Z^* = \{z_m^* | m \in M\}$.

Masking strategy is of great importance in the proposed method. Basically, we aim to mask mispronounced segments and expect the SSL pre-trained acoustic model can recover them with correctly pronounced tokens. However, whether the speech is mispronounced and where the mispronunciation occurs are unknown due to our unsupervised setting. To address this issue, we propose two strategies that are used to mask out the mispronunciation segments.

2.3.1. Random Masking

Random masking is a direct approach that is based on the masking strategy employed in pre-training. However, a single instance of random masking may have a lower probability of covering the mispronunciation component. To address this concern, we propose to repeat random masking k times for a given sequence X . Specifically, we randomly select $p\%$ of the frames in X as starting indices, and subsequently mask spans of l for each start index. These spans are mutually exclusive, with no overlap between them. By increasing the value of k , it is possible to ensure that each frame is masked at least once.

2.3.2. Regular Masking

Regular masking is an alternative approach that masks frames in a rule-based way. This strategy involves segmenting the input into k slices of equal length. We then proceed to mask one of those segments at a time and perform inference. The process is repeated until every segment has been masked at least once. The number k of segmented slices determines the granularity of the segmentation.

2.4. Scoring Module

In order to assess the quality of speech pronunciation, we introduce the scoring module, which measures the number of incor-

rectly recovered tokens based on Z and Z^* . Specifically, the average Mis-Recovered Token (aMRT) is proposed as a metric to measure the performance of pronunciation. Formally,

$$\text{aMRT} = \frac{1}{k} \sum_{j=1}^k \sum_{i \in M_j} \delta(z_i, z_i^*)$$

where $M_j \subset [T]$ represents the j -th set of indices to be masked, and function δ is defined as:

$$\delta(z, z^*) = \begin{cases} 0, & z = z^* \\ 1, & z \neq z^* \end{cases}$$

A higher aMRT value corresponds to a greater number of mis-recovered tokens and thus a lower quality of pronunciation. To obtain the PCC results between our proposed metrics and ground-truth scores, we adopt the negative values of aMRT as our final metrics.

3. Experiments

3.1. Dataset

We conduct experiments on the dataset *speechocean762* [7], which is specifically designed for pronunciation assessment. This open-source speech corpus is composed of 5,000 English utterances collected from 250 non-native speakers, half of whom are children. The corpus provides rich label information including phoneme, word, and sentence levels, and includes assessment scores ranging from 0 to 10 annotated by five experts. Our proposed approach is evaluated at the sentence level on the test set, which contains 2,500 utterances. We choose this public dataset for easy reproduction and comparison.

3.2. Baseline Models

We compare our proposed method with regression-based and non-regression-based baselines. The regression-based baselines include GoP [3, 20], DeepFeature¹ [8], and the state-of-the-art GOPT [12], all of which are supervised with human-annotated pronunciation scores. The non-regression-based baseline, on the other hand, utilizes the average phoneme-level GoP over the entire sentence as the measurement, and is referred to as non-reg GoP. This method does not require score annotations but instead uses a supervised ASR model.

3.3. Experimental Setup

We utilize the HuBERT-Base² model and adopt the CNN encoder, Transformer encoder, and k-means clustering in the experiments. HuBERT-Base is pre-trained on the LibriSpeech-960 [21], and the k-means with 100 clusters is fitted on LibriSpeech train-clean-100 split as per [22] using intermediate representations from HuBERT-Base. The output of the 7th layer of the Transformer Encoder is chosen as the default feature for clustering, as the resulting acoustic units perform well in discrimination tests [19, 17, 23]. We set masking probability $p = 20\%$, masking length $l = 5$, and repeating times $k = 50$ as the default. Each experiment is repeated three times with three different random seeds $\{13, 21, 100\}$, and the mean and standard deviation of the results are reported. Prior to performing the inference steps, all input audios are resampled with 16000 as

¹DeepFuture refers to the methods in [8] using deep features of ASR acoustic model

²<https://github.com/pytorch/fairseq>

the sample rate. The non-reg GoP is computed using Kaldi [24] to obtain the average phoneme-level GoP of the entire sentence. The ASR model utilized in this calculation is LibriSpeech ASR Chain Model³, as per [12].

3.4. Main Results

Two comparative studies are conducted to assess the effectiveness of the proposed method. The first study involves PCC performance comparison between our proposed method with regression-based and non-regression-based baselines, while the second study compares the PCC performance of different masking strategies.

The performances of various regression-based baselines and non-regression-based baselines are presented in Table 1. The results indicate that, compared to regression-based baselines, the proposed method lags behind the basic supervised baseline by a small margin of 0.04 PCC, despite a large performance gap of 0.14 PCC with the state-of-the-art supervised baseline. Notably, the proposed method is achieved by leveraging only the acoustic knowledge of HuBERT-Based acquired during pretraining, without the usage of annotated scores.

Furthermore, in comparison with the non-regression-based baseline, our proposed method shows a performance improvement of 0.03 PCC over the non-reg GoP. It is noteworthy that non-reg GoP requires an ASR model, while our method does not, underscoring the effectiveness of our ASR-free approach.

Table 1: Comparison between our method with regression-based and non-regression-based baselines on *speechocean762*

Model	PCC
Regression based	
GoP [3]	0.64
GoP(2BLSTM+MLP) [20]	0.67
DeepFeature [6]	0.72
GOPT [12]	0.74
Non-regression based	
non-reg GoP	0.57
Ours	0.60

Table 2 presents the performance comparison of two masking strategies employed in this study. The results show that random masking achieves superior performance with an improvement of 0.014 PCC over regular masking. We conjecture that this may be due to the fact that the input distribution with random masking is closer to the input distribution during pretraining, leading to enhanced performance. In addition, the experimental results reveal that random masking exhibits a low variance, indicating the stability of the method.

Table 2: Comparison of two masking strategies. The standard deviation of Random Masking is reported.

Masking Strategy	PCC
Random Masking	0.595 ± 0.002
Regular Masking	0.581

³<https://kaldi-asr.org/models/m13>

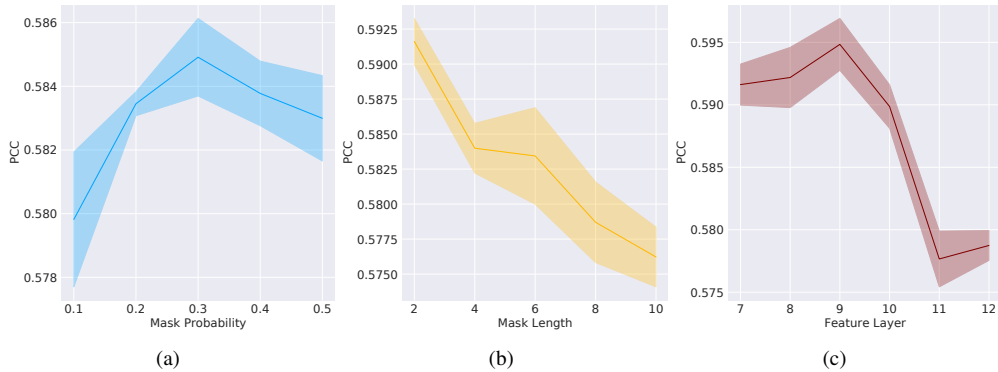


Figure 2: Impact of (a) mask probability, (b) mask length, and (c) feature layer on PCC results

3.5. Impact of masking hyperparameters

3.5.1. Random Masking

In order to further examine the impact of various hyperparameters of random masking, including masking probability, masking length, and feature layers used for clustering on the final results, three additional experiments are carried out. The results are presented in Figure 2.

The initial subfigure 2(a) illustrates the impact of mask probability on PCC results, with mask probability ranging from 0.1 to 0.5 with an interval of 0.1. The mask length is set to 5, and the feature layer is set to 7. The results indicate that the mask probability of 0.3 yields the best performance, while both higher and lower mask probabilities produce inferior outcomes. This observation may be attributed to the fact that the high mask probability may discard essential information that is required for reconstruction, whereas the low mask probability may decrease the possibility of masking mispronunciation parts.

Subfigure 2(b) showcases how the length of each masked span affects the PCC results. The mask length ranges from 2 to 10 with an interval of 2, while the mask probability is set to 0.2, and the feature layer is set to 7. The curve of this figure suggests a linear decrease in performance as the length increases. This phenomenon may stem from the pre-trained HuBERT-Base’s inadequate ability to recover a long masked span given the context.

Apart from the aforementioned factors, this study also investigates the degree to which the features used for clustering can contribute to pronunciation assessment. Therefore, the features from various layers of the Transformer Encoder ranging from 7 to 12 are examined. The outcomes presented in subfigure 2(c) reveal that using features from the 9th layer results in the best PCC performance. Generally, features from the 7th to 10th layer are considered useful for pronunciation assessment, whereas deeper features lead to poorer performance.

3.5.2. Regular Masking

For regular masking, we mainly investigate the impact of the slice number on the PCC results, namely how the mask granularity affects the PCC results. The results are presented in Figure 3. Our finding suggests that the more refined granularity of a single mask span does not necessarily lead to improved performance. One potential explanation for this outcome is that the use of a single mask span causes a shift from the input distribution, leading to poor performance. In addition, shorter masked spans may fail to cover entire words or even phonemes, which can have an adverse impact on the results.




Figure 3: Impact of slice number on PCC results

4. Discussion

While our zero-shot method achieves results comparable to supervised methods, it is essential to acknowledge that our method differs from the canonical-text-based pronunciation assessment. Our method draws on the acoustic knowledge obtained during pre-training, and thus, even if the transcription is different from the canonical text, a speech that is accurately pronounced may still receive a high score. Moreover, our method is limited to sentence-level assessment, and the exploration of unsupervised pronunciation assessment at the phoneme and word levels will be left as future work. The objective of this study is to establish a baseline and provide a pilot study of unsupervised pronunciation assessment.

5. Conclusion

In this paper, we present a zero-shot automatic pronunciation assessment approach. Instead of training regression models or using ASR models to compute GoP, we directly utilize a SSL pre-trained acoustic model and use the acoustic knowledge acquired from pre-training. To perform the ASR-free pronunciation assessment, we design two masking strategies and a novel evaluation metric to score the pronunciation of given speeches at the sentence level. Experimental results on speechocean762 achieve comparable performance to the supervised regression-based baseline and outperform the non-regression-based baseline. In the future, we hope to extend this research line of unsupervised pronunciation assessment to phoneme and word levels.

6. Acknowledgements

The authors would like to thank anonymous reviewers for their valuable suggestions. This project is funded in part by a research grant MOESOL-2021-0017 from the Ministry of Education in Singapore.

7. References

- [1] M. Eskenazi, "An overview of spoken language technology for education," *Speech Communication*, vol. 51, no. 10, pp. 832–844, 2009.
- [2] P. M. Rogerson-Revell, "Computer-assisted pronunciation training (capt): Current issues and future directions," *RELC Journal*, vol. 52, no. 1, pp. 189–205, 2021.
- [3] S. M. Witt and S. J. Young, "Phone-level pronunciation scoring and assessment for interactive language learning," *Speech communication*, vol. 30, no. 2-3, pp. 95–108, 2000.
- [4] S. Witt and S. Young, "Computer-assisted pronunciation teaching based on automatic speech recognition," pp. 25–35, 2014.
- [5] S. Sudhakara, M. K. Ramanathi, C. Yarra, and P. K. Ghosh, "An improved goodness of pronunciation (gop) measure for pronunciation evaluation with dnn-hmm system considering hmm transition probabilities," in *INTERSPEECH*, 2019, pp. 954–958.
- [6] W. Hu, Y. Qian, F. K. Soong, and Y. Wang, "Improved mispronunciation detection with deep neural network trained acoustic models and transfer learning based logistic regression classifiers," *Speech Communication*, vol. 67, pp. 154–166, 2015.
- [7] J. Zhang, Z. Zhang, Y. Wang, Z. Yan, Q. Song, Y. Huang, K. Li, D. Povey, and Y. Wang, "speechocean762: An open-source non-native english speech corpus for pronunciation assessment," *arXiv preprint arXiv:2104.01378*, 2021.
- [8] B. Lin and L. Wang, "Deep feature transfer learning for automatic pronunciation assessment," in *Interspeech*, 2021, pp. 4438–4442.
- [9] X. Xu, Y. Kang, S. Cao, B. Lin, and L. Ma, "Explore wav2vec 2.0 for mispronunciation detection," in *Interspeech*, 2021, pp. 4428–4432.
- [10] E. Kim, J.-J. Jeon, H. Seo, and H. Kim, "Automatic pronunciation assessment using self-supervised speech representation learning," *arXiv preprint arXiv:2204.03863*, 2022.
- [11] B. Lin, L. Wang, X. Feng, and J. Zhang, "Automatic scoring at multi-granularity for l2 pronunciation," in *Interspeech*, 2020, pp. 3022–3026.
- [12] Y. Gong, Z. Chen, I.-H. Chu, P. Chang, and J. Glass, "Transformer-based multi-aspect multi-granularity non-native english speaker pronunciation assessment," in *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2022, pp. 7262–7266.
- [13] X. Li, S. Dalmia, D. Mortensen, J. Li, A. Black, and F. Metze, "Towards zero-shot learning for automatic phonemic transcription," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 34, no. 05, 2020, pp. 8261–8268.
- [14] J. Xu, X. Tan, Y. Ren, T. Qin, J. Li, S. Zhao, and T.-Y. Liu, "Lrspeech: Extremely low-resource speech synthesis and recognition," in *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, 2020, pp. 2802–2812.
- [15] T. Brown, B. Mann, N. Ryder, M. Subbiah, J. D. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell *et al.*, "Language models are few-shot learners," *Advances in neural information processing systems*, vol. 33, pp. 1877–1901, 2020.
- [16] J. Wei, M. Bosma, V. Y. Zhao, K. Guu, A. W. Yu, B. Lester, N. Du, A. M. Dai, and Q. V. Le, "Finetuned language models are zero-shot learners," *arXiv preprint arXiv:2109.01652*, 2021.
- [17] T. A. Nguyen, M. de Seyssel, P. Rozé, M. Rivière, E. Kharitonov, A. Baevski, E. Dunbar, and E. Dupoux, "The zero resource speech benchmark 2021: Metrics and baselines for unsupervised spoken language modeling," *arXiv preprint arXiv:2011.11588*, 2020.
- [18] E. Dunbar, M. Bernard, N. Hamilakis, T. A. Nguyen, M. De Seyssel, P. Rozé, M. Rivière, E. Kharitonov, and E. Dupoux, "The zero resource speech challenge 2021: Spoken language modelling," *arXiv preprint arXiv:2104.14700*, 2021.
- [19] W.-N. Hsu, B. Bolte, Y.-H. H. Tsai, K. Lakhotia, R. Salakhutdinov, and A. Mohamed, "Hubert: Self-supervised speech representation learning by masked prediction of hidden units," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 29, pp. 3451–3460, 2021.
- [20] Z. Yu, V. Ramanarayanan, D. Suendermann-Oeft, X. Wang, K. Zechner, L. Chen, J. Tao, A. Ivanou, and Y. Qian, "Using bidirectional lstm recurrent neural networks to learn high-level abstractions of sequential features for automated scoring of non-native spontaneous speech," in *2015 IEEE Workshop on Automatic Speech Recognition and Understanding (ASRU)*. IEEE, 2015, pp. 338–345.
- [21] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, "Librispeech: an asr corpus based on public domain audio books," in *2015 IEEE international conference on acoustics, speech and signal processing (ICASSP)*. IEEE, 2015, pp. 5206–5210.
- [22] B. van Niekerk, M.-A. Carbonneau, J. Zaïdi, M. Baas, H. Seuté, and H. Kamper, "A comparison of discrete and soft speech units for improved voice conversion," in *ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2022, pp. 6562–6566.
- [23] B. van Niekerk, L. Nortje, M. Baas, and H. Kamper, "Analyzing speaker information in self-supervised models to improve zero-resource speech processing," *arXiv preprint arXiv:2108.00917*, 2021.
- [24] D. Povey, A. Ghoshal, G. Boulianne, L. Burget, O. Glembek, N. Goel, M. Hannemann, P. Motlicek, Y. Qian, P. Schwarz *et al.*, "The kaldi speech recognition toolkit," in *IEEE 2011 workshop on automatic speech recognition and understanding*, no. CONF. IEEE Signal Processing Society, 2011.