



An End-to-End Approach to Automatic Speech Assessment for Cantonese-speaking People with Aphasia

Ying Qin¹  · Yuzhong Wu¹ · Tan Lee¹ · Anthony Pak Hin Kong²

Received: 15 February 2019 / Revised: 3 October 2019 / Accepted: 1 December 2019
© Springer Science+Business Media, LLC, part of Springer Nature 2020

Abstract

Conventional automatic assessment of pathological speech usually follows two main steps: (1) extraction of pathology-specific features; (2) classification or regression on extracted features. Given the great variety of speech and language disorders, feature design is never a straightforward task, and yet it is most crucial to the performance of assessment. This paper presents an end-to-end approach to automatic speech assessment for Cantonese-speaking People With Aphasia (PWA). The assessment is formulated as a binary classification task to discriminate PWA with high scores of subjective assessment from those with low scores. The 2-layer Gated Recurrent Unit (GRU) and Convolutional Neural Network (CNN) models are applied to realize the end-to-end mapping from basic speech features to the classification outcome. The pathology-specific features used for assessment are learned implicitly by the neural network model. The Class Activation Mapping (CAM) method is utilized to visualize how the learned features contribute to the assessment result. Experimental results show that the end-to-end approach can achieve comparable performance to the conventional two-step approach in the classification task, and the CNN model is able to learn impairment-related features that are similar to the hand-crafted features. The experimental results also indicate that CNN model performs better than 2-layer GRU model in this specific task.

Keywords Pathological speech assessment · End-to-end · Aphasia · Cantonese · Deep neural network

1 Introduction

Aphasia refers to an acquired neurogenic speech-language disorder resulting from physical damage to specific brain regions. Symptoms of aphasia may adversely affect different modalities of language skills such as auditory comprehension, verbal expression, reading and writing [1].

The impairment could also span over various levels and components of the language system, including phonology, lexicon, syntax, and semantics [2]. Speakers with aphasia may have difficulties in recalling names of objects and/or putting words together into sentences [3]. Symptoms like anomia (word retrieval difficulty), dysfluency, voice disorder and dysprosody may be present in People With Aphasia (PWA) at various severity level and with different combinations [4, 5].

Speech assessment is an essential part of the comprehensive assessment for people with aphasia (PWA), which aims at determining the type and/or severity degree of impairment. It is required to be carried out by a well-trained speech and language therapist, based on subjective evaluation of various abilities of language communication. Subjective assessment of speech from PWA is a challenging task because it requires not only clinical knowledge about the disease but also good understanding of relevant linguistic and cultural background. There are clearly urgent practical needs to develop effective and reliable methods of automatic speech assessment for PWA.

Automatic analysis of PWA speech is expected to facilitate objective and efficient assessment to assist

✉ Ying Qin
yingqin@link.cuhk.edu.hk

Yuzhong Wu
yzwu@link.cuhk.edu.hk

Tan Lee
tanlee@cuhk.edu.hk

Anthony Pak Hin Kong
antkong@ucf.edu

¹ Department of Electronic Engineering, The Chinese University of Hong Kong, Shatin, NT, Hong Kong

² School of Communication Sciences and Disorders, University of Central Florida, Orlando, FL, USA

diagnosis and rehabilitation of PWA. Automatic speech recognition (ASR) technique and machine learning based approaches have been applied to this problem. Various acoustic features and text features were designed and evaluated for PWA speech assessment. Peintner et al. [6] performed an automatic classification on three sub-types of frontotemporal lobar degeneration in a relatively small dataset which includes the progressive non-fluent aphasia. They proposed a number of phone duration features, part-of-speech features as well as linguistic inquiry and word count features, which were extracted from ASR outputs. Fraser et al. [7, 8] performed an automatic classification of sub-types of primary progressive aphasia using acoustic features and text features extracted from manual transcriptions. In their follow-up work [9], an off-the-shelf commercial ASR system was adopted to automatically generate transcriptions for text feature extraction. However, the effectiveness of assessment was limited by the low recognition accuracy on impaired speech. Duc Le et al. [10, 11] improved the PWA speech recognition using discriminative pre-training with out-of-domain dataset and multi-task acoustic model. They achieved good performance in predicting subjective assessment scores based on ASR outputs by analyzing text statistics (e.g. the number of nouns and verbs in spoken utterances), part-of-speech language model etc. Kohlschein et al. [12] proposed to automatically classify four most prevalent aphasia syndromes (Global aphasia, Broca's aphasia, Wernicke's aphasia and amnesic aphasia) based on manual transcriptions. A Long Short-Term Memory (LSTM) based classifier was trained with word vectors which were derived from the manual transcriptions. The peak classification accuracy obtained by their proposed model was 44.3%, and involving more training samples was suggested to further improve the classification performance. The above studies target the automatic speech assessment for English-speaking and German-speaking PWA. In recent years, we made great efforts to develop an automatic assessment system for Cantonese-speaking PWA. In [13], the effectiveness of supra-segmental duration features in speech assessment of Cantonese-speaking PWA was investigated. The features such as duration of pause-delimited speech segments, duration of silence segments and frequency count of silence segments were found to be useful in differentiating speech of PWA from that of unimpaired individuals. The computation of these duration features requires a forced alignment process relying on manual transcriptions, which is undesirable in practical applications. A general-purpose ASR system was applied on PWA speech in [14] and a higher syllable error rate was observed for aphasia subjects compared with unimpaired subjects. In [15], we proposed a framework of fully automatic speech assessment for PWA. Supra-segmental duration features were computed from time alignment

generated by a dedicated ASR system. The text features were derived from the ASR-generated text output using syllable-level embedding technique. With the combination of these extracted features, the prediction of severity of impairment was formulated as a regression task. The prediction results showed a high correlation with subjective assessment scores. We also attempted to improve the ASR performance on PWA speech using multi-task learning strategy, and used N-best lists and confusion networks to mitigate the effect of ASR errors on the text features [16].

From previous studies, it is evident that the exploration of pathology-specific features is an essential part of designing the assessment system and it is usually based on expert knowledge. However, some impairment-related characteristics contained in the raw speech data may be missed after the extraction of hand-crafted features. After extracting a great number of acoustic features and text features, post-processing procedures such as feature selection and classification/regression are required. They are generally complicated and time-consuming. In addition, parameters of feature extraction and assessment model are optimized separately instead of jointly. Recently, the so-called "end-to-end" approach has demonstrated good successes in ASR [17], machine translation [18], and other applications. It requires relatively little human effort on designing tailored features and shows superior performance [17]. Inspired by the "end-to-end" paradigm, we propose the design of a "utterance-to-score" system for PWA speech assessment to boost the efficiency of assessment system. Utterances spoken by PWA are fed directly to a Deep Neural Network (DNN), without requiring explicit feature extraction and selection. The output of the DNN gives the predicted assessment score for each input utterance. The overall assessment score for the impaired speaker is obtained by combining all his/her utterance-level scores.

This paper reports our preliminary results on PWA speech assessment using the end-to-end approach. Among numerous neural network structures, Recurrent Neural Network (RNN) has the ability to model sequential signals and to deal with variable-length utterances. LSTM [19] and Gated Recurrent Unit (GRU) [20] are well-established "gating" models that tackle the problem of gradient vanishing in vanilla RNN. The GRU is able to achieve comparable performance to LSTM, with a more simplified architecture. It has been applied to utterance-based classification tasks, e.g., question detection [21] and emotion classification [22]. On the other hand, Convolutional Neural Network (CNN) based approaches have been investigated in the speech assessment area. It was successfully applied to spoken fluency scoring [23] and differentiate speech of people with Parkinson disease from that of healthy people [24]. In this paper, we focus

on investigating the effectiveness of GRU model and CNN model in PWA speech assessment.

Although the CNN-based models are demonstrated to be effective in speech and audio classification tasks, there have been few attempts to provide interpretations of what types of speech or audio patterns are perceived by a CNN model. In the image classification domain, similar issues have been extensively explored. The Class Activation Mapping (CAM) [25] was proposed to highlight the class-discriminative image regions in CNNs with global average pooling. A generalized version of CAM named the Gradient-weighted Class Activation Mapping (Grad-CAM) was proposed in [26], which can be applied to a wider variety of CNN structures. Recently, the Grad-CAM technique was used to analyze the classification results in the audio scene classification task [27]. Inspired by this, in the present study, we perform the visualization of CAM based on the time-frequency representations of acoustic features, which can be regarded as 2-dimensional images. It allows the comparison between machine perception and human interpretation, as well as the analysis of what acoustic clues are actually important for deciding the severity degree of PWA.

2 Corpus: Cantonese AphasiaBank

Cantonese is an influential Chinese dialect spoken by tens of million of people in Hong Kong, Macau, Southern China as well as overseas Chinese communities. Like Mandarin, Cantonese is a monosyllabic and tone language. Each Chinese character is spoken as a monosyllable carrying a specific tone.

Cantonese AphasiaBank used in this study is a large-scale multi-modal corpus jointly developed by the University of Central Florida and the University of Hong Kong [28]. It aims to support both fundamental and clinical research on Cantonese-speaking aphasia population. The corpus contains recordings of spontaneous speech from 149

unimpaired and 104 aphasic subjects who are all native Cantonese speakers. The speech recordings were elicited following the AphasiaBank protocol, with adaptation to local Chinese culture [29, 30]. Each subject was required to complete 9 narrative tasks, including 4 picture descriptions, 1 procedure description, 2 story telling and 2 personal monologues, with details given as in Table 1. Except personal monologues, the speech produced in each task is expected to be about a specific topic. All the recordings happened in a sound proof booth and were collected using a standardized recording set-up and protocol. Specifically, the audio signal was simultaneously recorded using a solid-state digital recorder, at a sampling frequency of 44.1 kHz with 16-bit quantization, and a headworn condenser microphone. All impaired subjects in the corpus went through a standardized assessment using the Cantonese Aphasia Battery [31]. It involves a series of sub-tests measuring fluency, information content, comprehension, repetition and naming abilities of the subject. The sum of sub-test scores is named the Aphasia Quotient (AQ). The value of AQ ranging from 0 to 100 is an indication of overall severity of impairment. Lower AQ value means higher degree of severity. In [31], it was suggested to use $AQ = 96.4$ as cutoff point for differentiating normal subjects from those with aphasia, as it represented 2 standard deviations below the mean of healthy controls' performance. In this study, the AQ score is regarded as the ground-truth measure of the severity of aphasia.

In this study, about 15.6 hours speech data of 9 tasks from 91 impaired subjects are selected for the following experiments, including 58 Anomic subjects, 6 Transcortical sensory subjects, 12 Transcortical motor subjects, 10 Broca's subjects, 1 Isolation subject, 2 Wernicke's subjects and 2 Global aphasia subjects. Their AQ values are in the range of 11.0 to 99.0. Figure 1 shows the histogram of the AQ scores of 91 aphasic subjects. As the participants we extracted from the Cantonese AphasiaBank were those who could perform open-ended spoken discourse tasks, they tended to be higher-level functioning and less impaired in

Table 1 Recording tasks in the Cantonese AphasiaBank.

Task	Recording	Description
Single picture description	CatRe	Black and white drawing of a cat on a tree being rescued.
	Flood	A color photo showing a fireman rescuing a girl.
Sequential picture description	BrowN	Black and white drawing of a boy accidentally breaking a window.
	RefUm	Black and white drawing of a boy refusing an umbrella from his mother.
Procedure description	EggHm	Procedures of preparing a sandwich with egg, ham and bread.
Story telling	CryWf	Telling a story from a picture book "The boy who cried wolf".
	TorHa	Telling a story from a picture book "The tortoise and the hare".
Personal monologue	ImpEv	Description of an important event in life.
	Stroke	Description of the experience of suffering a stroke.

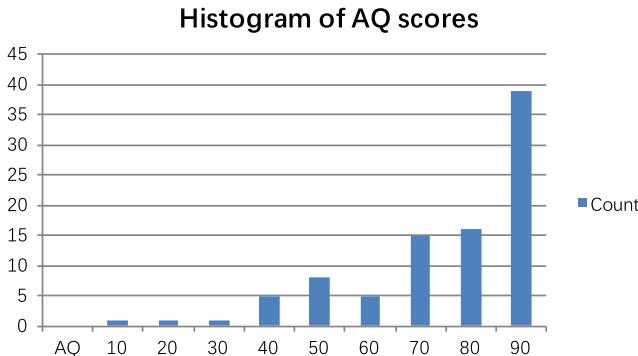


Figure 1 Histogram of AQ scores of 91 impaired speakers.

terms of severity (see details in [32]). Therefore, in our case, it is typical that the AQ scores cluster towards the high scores among the aphasic group. The speakers who had lower AQ scores, i.e., those in the range of 10 or 20, were those with non-fluent aphasia, such as Broca's aphasia, Isolation aphasia, Transcortical Motor aphasia, or Global aphasia.

3 General Framework

We propose to utilize end-to-end approach to differentiating PWA with High-AQ ($AQ \geq 90$) from those with Low-AQ ($AQ < 90$). The cutoff value of 90 is set to reach balanced number of subjects in two groups. Figure 2 illustrates the general framework of proposed classification system. Fundamental frame-level acoustic features (e.g. Mel-frequency cepstral coefficients, filterbank features) are extracted from the utterance of PWA and subsequently fed to a Neural Network-based (NN-based) classifier. Under the “end-to-end” framework, neither pathology-specific feature extraction nor feature selection procedures are required. They are expected to be automatically carried out by the NN-based classifiers. The classification label of the input utterance is inherited from the impaired speaker. For the

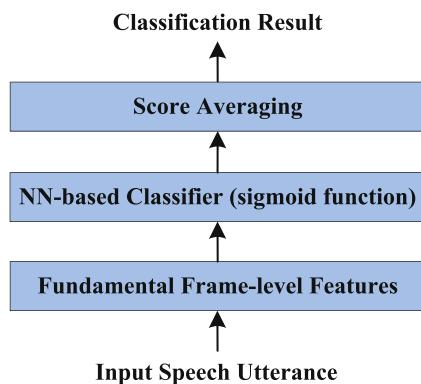


Figure 2 General framework of the classification system.

output of classifier, an assessment score for each utterance can be obtained after the sigmoid function of the neural network. The speaker-level overall assessment score is computed by taking the average of all utterance-level scores from the test speaker. The higher overall assessment score indicates a higher possibility of the speaker being in the High-AQ group.

4 Methods

4.1 Models

In this study, we experiment with two mainstream structures of neural network, namely the 2-layer GRU and CNN for the classification purpose. GRU model is known for learning features along a temporal sequence of speech frames, while CNN model is good for exploring features from time-frequency matrix representation of speech. These two different mechanisms of pattern modeling are compared on the intended aphasia speech assessment task.

4.1.1 2-layer GRU

GRU is specialized in sequential modeling, which is expected to characterize the impairment of the entire utterance spoken from PWA. The architecture of proposed 2-layer GRU model is detailed in Fig. 3. Frame-level acoustic features o_1, o_2, \dots, o_T of an utterance are sequentially fed to the neural network. A hidden representation is generated at each time step, summarizing the past acoustic information. The hidden representation h at the t^{th} time step is given by

$$h_t = \text{GRU}(o_1, o_2, \dots, o_t; \theta_h), \quad (1)$$

where θ_h denotes the model parameters and the function $\text{GRU}(\cdot)$ indicates two uni-directional GRU layers with hidden size of 200 per layer. Our preliminary experiments

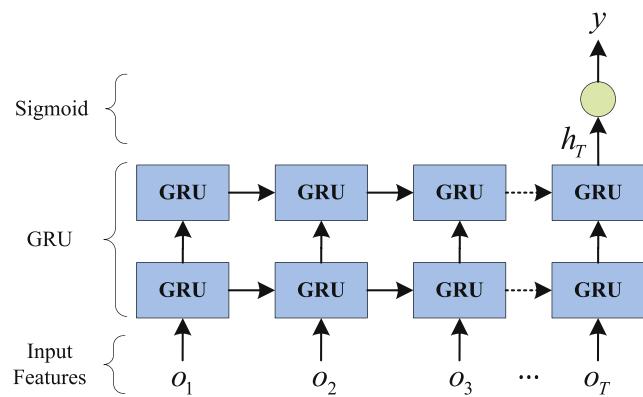


Figure 3 Architecture of the 2-layer GRU.

show no significant improvement using the bi-directional GRU layers, thus a uni-directional structure is adopted in this study. A predicted score y for the input utterance is produced by a sigmoid activation function $\text{Sigmoid}(\cdot)$ at the output layer, which is given by

$$y = \text{Sigmoid}(\mathbf{W}_o \mathbf{h}_T + \mathbf{b}_o), \quad (2)$$

where \mathbf{h}_T is the hidden representation at the last time step which summarizes the information of entire utterance. \mathbf{W}_o and \mathbf{b}_o are the trainable parameters connecting \mathbf{h}_T with y .

4.1.2 Convolutional Neural Network

CNN is known as to learn local high-level features from the spatial or temporal data. It performs discrete convolution between a local region of the input image f and a set of filters w with the shapes of $m \times n$. After the input image f passes through the convolutional layer, the output is given by

$$(f * w)[i, j] = \sum_m \sum_n f[i + m, i + n] w[m, n]. \quad (3)$$

For the input utterance, the short-time acoustic features are extracted and concatenated sequentially as a 2-dimensional time-frequency representation. The CNN model structure described in Table 2 is inspired by the AlexNet [33] and VGG [34] models. It is constructed with a Global Average Pooling (GAP) layer instead of a fully-connected layer after the last convolutional layer. GAP has been proved to be a good regularizer for CNNs in image classification [35] and audio scene classification [27] tasks. Batch normalization (BN) layer and ReLU activation function are added after each convolutional layer.

Table 2 Architecture of the CNN model.

Input 1 × 300 × 128	
1	3 × 3 Convolution (pad-1, stride-1)-64-BN-ReLU
2	3 × 3 Max Pooling (stride-2)
3	3 × 3 Convolution (pad-1, stride-1)-192-BN-ReLU
4	3 × 3 Max Pooling (stride-2)
5	3 × 3 Convolution (pad-1, stride-1)-384-BN-ReLU
6	3 × 3 Convolution (pad-1, stride-1)-256-BN-ReLU
7	3 × 3 Convolution (pad-1, stride-1)-256-BN-ReLU
8	3 × 3 Max Pooling (stride-2)
9	Global Average Pooling
10	Sigmoid

4.2 Class Activation Mapping

Given a CNN model, the technique of Class Activation Mapping (CAM) [25] is capable of visualizing the class-discriminative regions in the input images. Therefore, the CAM can improve understanding and interpretability of CNN models. Given a trained CNN model with GAP, the mean values of K feature maps at the penultimate layer are linearly transformed to generate a score y^c for each class c ,

$$y^c = \sum_k w_k^c \frac{1}{Z} \sum_x \sum_y f_k(x, y), \quad (4)$$

where $f_k(x, y)$ denotes the point (x, y) in the k^{th} feature map before GAP. The weight w_k^c connecting k^{th} feature map to output class c indicates the importance of k^{th} feature map for the target class c . Z is a normalization term, representing the number of pixels in the feature map. The spatial elements of class activation map L_{CAM}^c for class c are obtained by exchanging the order of summation in Eq. 4,

$$L_{\text{CAM}}^c(x, y) = \sum_k w_k^c f_k(x, y). \quad (5)$$

The above CAM is only applicable to CNNs with GAP. As a strict generalization of CAM, the Gradient-weighted Class Activation Mapping (Grad-CAM) can be applied to a wider variety of CNN models, such as those with fully-connected layers (e.g. VGG, AlexNet). It is implemented by replacing the weight of each feature map w_k^c with the average gradient (α_k^c) of the score for class c with respect to each feature map, which is computed by

$$\alpha_k^c = \frac{1}{Z} \sum_x \sum_y \frac{\partial y^c}{\partial f_k(x, y)}. \quad (6)$$

Then we have

$$L_{\text{Grad-CAM}}^c(x, y) = \text{ReLU}(\sum_k \alpha_k^c f_k(x, y)), \quad (7)$$

where ReLU is applied to identify the features that have a positive influence on the target class. Note that f_k in this case is no longer limited to the last convolutional layer in CNN but can be from any convolutional layer.

We propose to use Grad-CAM technique to analyze a trained CNN model for the assessment task. It is expected to help the understanding of the feature patterns learned by CNN in the decision of severity degree. Note that the class number c is equal to 1 and the sigmoid function is used in this study. Equation 7 is used to highlight the positive regions to the High-AQ class. By negating the weighted combination before the ReLU function, the negative regions to the High-AQ class (positive to Low-AQ class) could be visualized.

Table 3 The number of variable-length utterances in each fold by manual segmentation.

	Fold 1	Fold 2	Fold 3	Fold 4	Fold 5
Train	8,000	8,278	8,207	8,535	8,368
Test	2,623	2,001	2,389	2,261	2,276
Valid	927	1,271	954	754	906

5 Experimental Setup

5.1 Data

To perform severity assessment for all 91 PWA, the binary classification experiment is carried out with the arrangement of 5-fold cross validation. In each fold, 80% of the subjects are used for training and the rest 20% subjects are used for test. 10% of subjects are randomly selected from training subjects as the validation data. **There are 39 PWA in High-AQ group and 52 PWA in Low-AQ group respectively.**

To provide more training samples, the long recordings are segmented into short utterances. Each short utterance will be given a score of severity degree by neural networks independently, depending on the learned features related to speech impairment. We consider to use two types of segmentation methods to process speech recordings. The first method is manual segmentation according to the sentence boundaries marked in the Cantonese AphasiaBank. In this way, the speech recordings of 91 PWA are cut into 11,550 utterances. A total number of 6,172 utterances spoken from High-AQ speakers are assigned with label 1, while 5,378 utterances from Low-AQ group are assigned with label 0. The lengths of utterances range from 0.07 second to 70.03 seconds. Table 3 details the 5-fold data sets with the manual segmentation.

The second segmentation method is to cut the speech recordings of 91 PWA into non-overlapping segments of 3 seconds, leading to 18,149 equal-length short utterances for training models. There are 7,322 utterances and 10,827 utterances being labeled as 1 and 0 respectively. The number of equal-length utterances in 5 folds is listed in Table 4.

Table 4 The number of equal-length utterances of 3 seconds in each fold.

	Fold 1	Fold 2	Fold 3	Fold 4	Fold 5
Train	12,839	13,270	13,155	13,984	12,669
Test	4,100	3,224	3,725	2,940	4,160
Valid	1,210	1,655	1,269	1,225	1,320

5.2 Features

Mel-Frequency Cepstral Coefficient (MFCC) and Log-Mel filterbank features are explored to train the proposed models. The 39-dimensional MFCC+ Δ + $\Delta\Delta$ features are extracted from each input utterance, with 25 milliseconds window length and 10 milliseconds window shift. The number of points in Discrete Fourier Transform (DFT) is set to 2,048, and the number of filters in the Mel filterbank is set to 128. Dimension-wise normalization of the MFCC features is performed within each fold using the mean and variance computed from the training set. For the extraction of Log-Mel filterbank features, the settings of DFT and Mel filterbank are the same as those used for extracting the MFCC features. The Log-Mel filterbank features are derived from short-time frames (same window size and window shift as mentioned before) from input utterances, with the dimension of 128. Subsequently, the frame-level features are stacked together as a time-frequency matrix for the input utterance. The feature normalization is also applied to Log-Mel filterbank features in each fold.

In addition, topic information is taken into consideration in the feature design. With the topic information, the content-related acoustic patterns are expected to be learned by the NN-based models. For instance, if an utterance is from the topic 1, the first element of a topic vector (topicvec) is set to 1 and other 8 elements are set to zeros. For training the 2-layer GRU model, the 9-dimensional one-hot topic vector is concatenated with a 128-dimensional Log-Mel filterbank feature, resulting in a 137-dimensional vector for each frame as a kind of topic adaptation. For the CNN model, the one-hot topic vectors are utilized in a different way. The topic vector is concatenated with the embedding of time-frequency representation of the utterance. Specifically, the 256-dimensional high-level features derived from the Log-Mel filterbank features can be obtained after the GAP layer, which is shown in Table 2. They are concatenated with 9-dimensional one-hot topic vectors and then fed to a fully-connected neural network containing two hidden layers (the vector size of each hidden layer is 256). It is followed by a sigmoid layer as the final output layer.

In the following experiments, we consider to use MFCC, Log-Mel filterbank and topic-adapted features derived from two types of segmented speech data to train the 2-layer GRU model and CNN model.

5.3 Hyperparameters for Model Training

The training parameters are set empirically. The mini-batch size is 64 for training the 2-layer GRU model and the CNN model. The initial learning rate is set to 10^{-4} for the 2-layer GRU model and it is set to 10^{-3} for the CNN model. Model training aims at minimizing the binary

cross-entropy loss using the Adam optimizer [36]. Weight decay coefficient is set to 5×10^{-4} for both models to tackle the overfitting problem. All models are implemented using the deep learning toolkit Pytorch [37].

6 Experiments and Results

6.1 Utterance-level Classification Accuracy

For the binary classification on test utterances from High-AQ and Low-AQ subjects, the Area Under receiver operating characteristic Curve (AUC) [38] is adopted as the performance metric. In the case of binary classification, AUC can be viewed as the probability that the classifier ranks a randomly chosen positive sample higher than a negative one [39]. An AUC value 0.5 means a random guess and 1.0 represents a perfect classification.

Table 5 compares the AUC results of 5-fold cross-validation experiments with 2-layer GRU model (GRU model) and CNN model. The “-var” and “-eq” denote the acoustic features extracted from variable-length and equal-length utterances respectively. The weighted averaged of AUC (W-AUC) across all folds is also compared, where the weight is the proportion of test data of each fold in the whole test data. From the AUC results based on the GRU model, it can be seen that higher AUCs are obtained using the Log-Mel filterbank (Log-Mel) features than using the MFCC features. This suggests that the Log-Mel features are more suitable for the intended classification task. The result comparison between Log-Mel-var and Log-Mel-eq shows that the equal-length utterances perform better than the variable-length ones for training the GRU model. It may attribute to a larger number of utterances used for training. With the Log-Mel-eq features, a significant AUC improvement can be observed using the CNN model compared with using the GRU model. This probably because the CNN model performs better on capturing

localized time-frequency information of impairment-related acoustic patterns. Nevertheless, the CNN model does not benefit much from the topic vectors for the severity degree classification, whilst for the 2-layer GRU model, using the concatenation of Log-Mel and topic vector features achieves better performance than using the Log-Mel features only. This shows that the topic information may provide additional benefits to the classification performance to some extent.

The 2-layer GRU model trained with Log-Mel+topicvec-eq features attains the best performance among all 2-layer GRU models, thus it is used for the analysis in Sections 6.2 and 7.1. For the CNN model, it is noted that the topic vectors show little improvement on the classification performance. Therefore, the CNN model trained with Log-Mel-eq without topic vectors is used for the following analysis of experimental results and Grad-CAM visualization.

Although the utterance-level AUC is a direct metric to measure the performance of classifiers, it is unable to reflect the performance of severity degree assessment for PWA. A score fusion procedure is required to give a speaker-level classification decision.

6.2 Speaker-level Classification Accuracy

We consider to compare the speaker-level classification accuracy given by the 2-layer GRU model with the best performance (trained with Log-Mel+topicvec-eq features), the CNN model (trained with Log-Mel-eq features) and a baseline system. As shown in Fig. 2, for the NN-based models, after obtaining utterance-level scores from a test speaker, a score fusion is performed to combine them as an overall score for the speaker-level classification. The overall score for each test speaker is obtained by taking the average of speaker’s utterance-level scores. The threshold for binary classification is set to 0.5. If the overall score is higher than 0.5, the test speaker is

Table 5 Performance comparison between 2-layer GRU and CNN models. MFCC, Log-Mel and topic-adapted features derived from variable-length (-var) and equal-length (-eq) utterances are used for training the models. W-AUC refers to the weighted average of AUC. The highest AUC value for each fold is marked in bold

Model	Feature	AUC					W-AUC
		Fold 1	Fold 2	Fold 3	Fold 4	Fold 5	
GRU	MFCC-var	0.58	0.62	0.58	0.61	0.80	0.64
	Log-Mel-var	0.62	0.65	0.62	0.61	0.80	0.67
	Log-Mel-eq	0.62	0.69	0.64	0.61	0.81	0.68
	Log-Mel+topicvec-eq	0.64	0.75	0.65	0.63	0.83	0.70
CNN	Log-Mel-eq	0.74	0.81	0.78	0.70	0.90	0.79
	Log-Mel+topicvec-eq	0.70	0.82	0.76	0.70	0.89	0.78

Table 6 Performance comparison between the baseline system, 2-layer GRU and CNN model. It is measured in terms of the accuracy, average F1 score, recall, specificity and equal error rate (EER). The values in bold indicate best classification results

Model	Accuracy	F1	Recall/Specificity	EER
Baseline	0.813	0.810	0.795/0.827	—
GRU	0.791	0.789	0.795/0.789	0.212
CNN	0.824	0.822	0.821/0.827	0.173

classified as High-AQ, otherwise the speaker is classified as Low-AQ.

The baseline assessment system in this study follows a conventional two-step assessment approach proposed in our previous study [15]. A 5-dimensional feature vector of supra-segmental duration features is evaluated on the same task of binary classification using a random forest classifier. The leave-one-out cross validation strategy is adopted. The feature vector includes: (1) duration ratio between non-speech part and speech part; (2) average duration of silence segments (longer than 0.5 second); (3) average duration of speech segments (speech region between two silence segments); (4) ratio of silence segment count to syllable count; (5) syllable count per second. All of the features are generated from the time alignment of a dedicated ASR system. The time-delay layers stacked with bidirectional long short term memory layers (TDNN-BLSTM) are used as acoustic model of the ASR system and it is trained using multi-task learning strategy [16]. These ASR-generated features were shown to be effective to classify High-AQ speakers from Low-AQ ones in the aspect of acoustic impairment of PWA speech [15].

Table 6 lists the speaker-level binary classification results on 91 test speakers with the threshold of 0.5. The classification performance is measured in terms of the accuracy, average F1 score, recall and specificity. It is seen that the CNN model performs the best among the three models, while the performance of 2-layer GRU model is worse than that of the baseline system. This demonstrates that the proposed CNN model is preferred than the 2-layer GRU model for this specific assessment task. The CNN model outperforms the baseline system in terms of the recall, with the value of 0.821 (32/39) versus 0.795 (31/39). Table 7 shows the confusion matrix

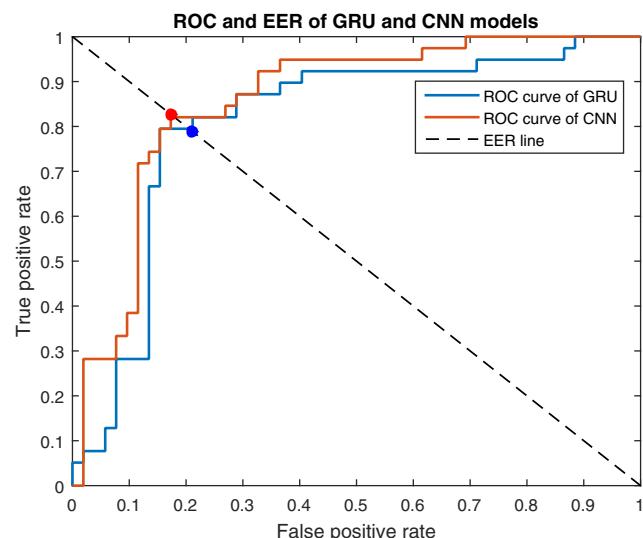


Figure 4 ROC curves attained with 2-layer GRU and CNN models.

of classification result using the CNN model. The result suggests that the end-to-end approach has great potential to replace the traditional two-step method to perform speech assessment for PWA. With all possible classification thresholds, we further plot ROC curves (see Fig. 4) and compute equal error rate (EER) (see Table 6) to measure the performance of GRU and CNN models. From Fig. 4, the same conclusion of CNN outperforming GRU can be drawn.

Pairwise McNemar's Test [40, 41] is performed on the classification results given by the three systems in Table 6. The results are shown in Table 8. It is noted that there is no statistically significant difference among them.

7 Discussion

7.1 2-layer GRU vs. CNN

To investigate the performance of the 2-layer GRU and CNN models, we analyze twenty typical impaired subjects for whom the classification results are correct. Ten of the selected speakers are with High-AQ while the others are with Low-AQ. Figure 5 illustrates the histograms of utterance-level scores given by the 2-layer GRU and CNN for High-AQ speakers respectively. Although two models

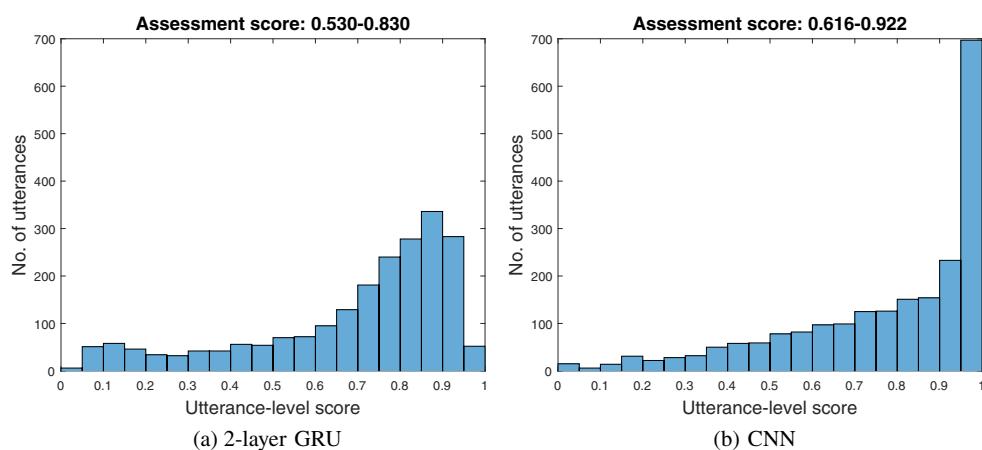
Table 7 Confusion matrix of classification result using CNN model.

	True: High-AQ	True: Low-AQ	Total
Predicted: High-AQ	32	9	41
Predicted: Low-AQ	7	43	50
Total	39	52	91

Table 8 Pairwise McNemar's Test on classification results given by the three systems. $p < .05$ indicates statistical significance.

Model comparison	Results of McNemar's Test
Baseline vs. GRU model	$\chi^2(1) = .062, p = .803$
Baseline vs. CNN model	$\chi^2(1) = .000, p = 1.000$
CNN model vs. GRU model	$\chi^2(1) = .444, p = .505$

Figure 5 Comparison between histogram of utterance-level scores from 2-layer GRU model and that from CNN model for ten speakers with High-AQ.

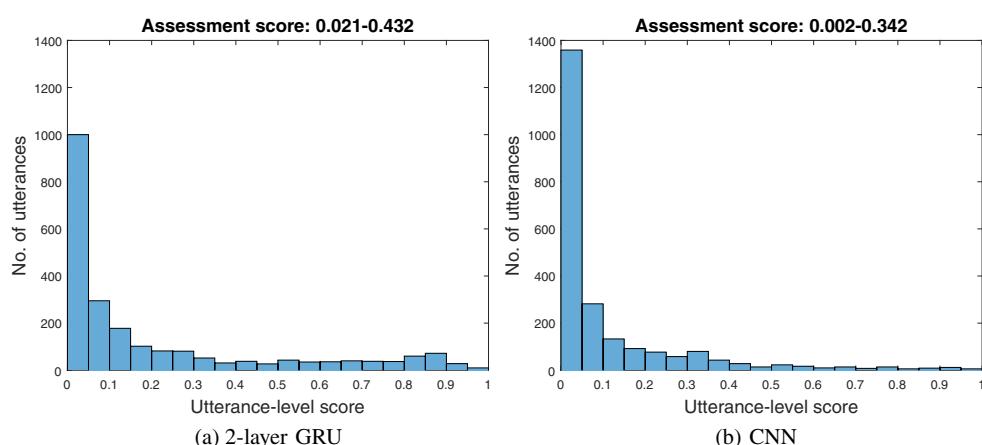


make correct decisions on these speakers, their overall assessment scores generated from the CNN (range from 0.616 to 0.922) is much nearer to 1.0 than those from the 2-layer GRU (range from 0.530 to 0.830). Specifically, most of the utterance-level scores are higher than 0.9 with the CNN model, while those are mainly distributed in the range of 0.7 to 0.9 based on the 2-layer GRU model. Similar to the previous case, the CNN performs better than the 2-layer GRU in generating overall scores for ten speakers with Low-AQ (0.002 – 0.342 versus 0.021 – 0.432). The distribution of utterance-level scores produced by the CNN model is even denser in the low-score region as shown in Fig. 6. Generally speaking, utterance-level scores from the CNN model tend to be more polarized than those from the 2-layer GRU model, meaning the CNN model has a higher classification confidence. This is also consistent with the AUC results as shown in Table 5.

7.2 Visualization with Class Activation Maps

Given a trained CNN model, the Grad-CAM technique is applied to explore what types of features in impaired speech are learned in the CNN classifier to perform assessment.

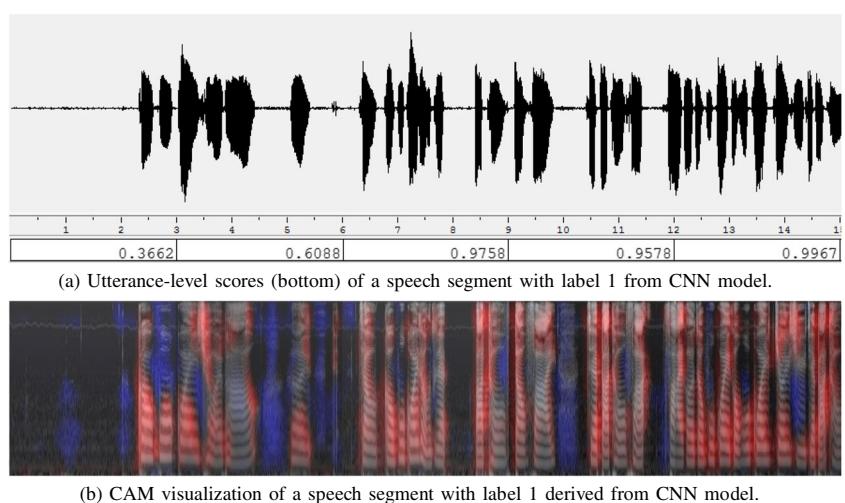
Figure 6 Comparison between histogram of utterance-level scores from 2-layer GRU model and that from CNN model for ten speakers with Low-AQ.



Recall that 3-second Log-Mel features are as the input to the CNN model, and they can be viewed as time-frequency gray-scale images. After simply up-sampling the class activation maps to the same size of the input image, the transformed class activation maps are further combined with the input image for localizing the class-discriminative regions. The proposed CAM visualization consists of three image components. The base one is the gray-scale Log-Mel image. The second component is a semi-transparent image in red derived from the class activation map. It contains the time-frequency regions that positively influence the classification score of the High-AQ class. For the last component, it is a semi-transparent image of blue color, in which the blue areas indicate regions having a negative influence on the class of High-AQ. The detailed computations of class activation maps for creating the second and third components are described in section 4.2.

Figures 7 and 8 give utterance-level scores and CAM visualizations of two 15-second speech segments from a High-AQ speaker (AQ: 94.7) and a Low-AQ speaker (AQ: 74.2) respectively. These two speakers are correctly classified by the CNN model. The CAM visualizations are derived from the last convolutional layer (7th layer) in

Figure 7 Utterance-level scores and CAM visualization of the same speech segment from a High-AQ speaker. They are derived from the CNN model with Log-Mel-eq input. The positive and negative activation for target class are highlighted in red and blue colours respectively.



the CNN structure, since the last convolutional layer was expected to keep balance between high-level semantics and detailed spatial information [26]. It is clearly observed that the area of positive activation (red) in speech from the High-AQ speaker is larger than that from the Low-AQ speaker. In general, we observe that the CNN model is able to learn the following features to determine the severity degree of PWA:

1. Duration-related features:

As illustrated in Figs. 7b and 8b, silence parts are usually highlighted as negative activation in blue, while speech parts are as positive activation in red. Meanwhile, 3-second utterances with longer pauses tend to have lower classification scores from CNN model, which can be observed in Figs. 7a and 8a. This indicates that milder PWA (with higher AQ) tend to exhibit less dysfluency and vice versa. This also confirms that the CNN is able to automatically learn similar duration-related features as we designed in the baseline system, such as “average duration of silence segments” and “average duration of speech segments”.

2. Transitions between speech parts and silence parts:

As we can see from Figs. 7b and 8b, the magnitudes of positive activation (red) are higher at the transition regions between speech parts and silence parts. Thus, speakers who have high speaking rate with a large number of transitions in their speech have a higher probability to be classified as High-AQ. This CNN-learned feature is also matched with one of the conventional features named “syllable count per second” used in the baseline system.

3. Significant variation of formants:

It can be seen that the positive activation (red) concentrate more on the time-frequency regions with significant variation of formants. Speakers who produce utterances with a variety of formant changes are more likely to be classified into the High-AQ group. This reveals that the fundamental formant information is important to the classification result and it can be implicitly learned by the CNN model. As mentioned in Section 6.2, the formant feature is not considered in the baseline system, which may be the reason of the poorer

Figure 8 Utterance-level scores and CAM visualization of the same speech segment from a Low-AQ speaker. They are derived from the CNN model with Log-Mel-eq input. The positive and negative activation for target class are highlighted in red and blue colours respectively.

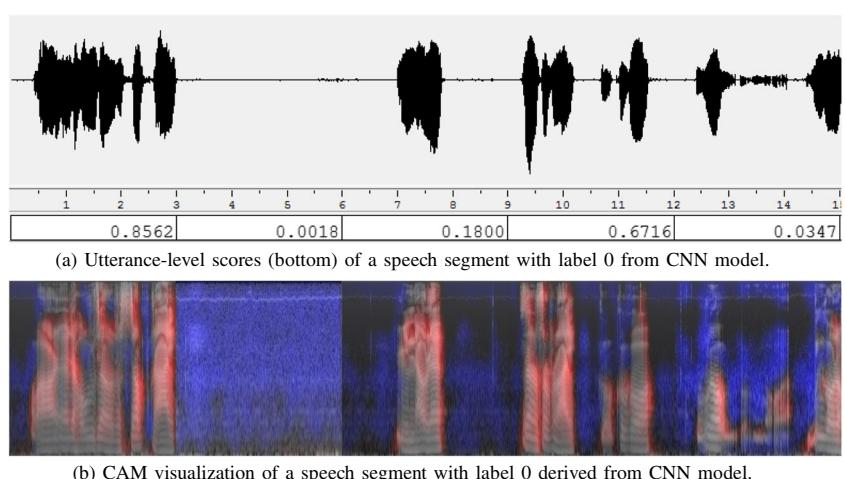
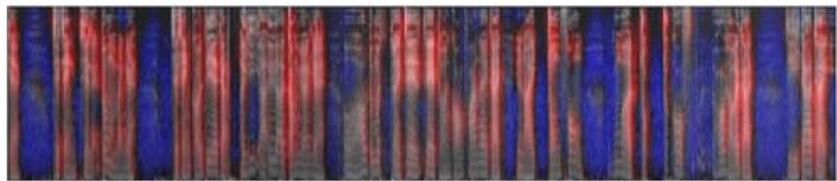


Figure 9 CAM-visualization of a speech segment from the over-estimated speaker.



classification performance than the CNN model. The lack of formant changes could be related to phonetic disorder, voice disorder and limited vocabulary in PWA speech. This suggests that using formant features for the PWA speech assessment is promising and it will be investigated in the future.

Overall, with the CAM visualization, the CNN model is demonstrated to be capable of implicitly learning impaired acoustic patterns for the PWA speech assessment.

7.3 Limitations of Study

We analyze one test speaker (AQ: 66.8) who is misclassified into High-AQ group with the CNN model. It is found that there are a large number of function words and filler words but few topic-specific words in his utterance, even though his utterances are quite fluent. Figure 9 shows the CAM visualization of a speech segment from this over-estimation impaired subject, in which most of the speech parts are in red. This reveals a major limitation of proposed model: it fails to sufficiently learn the semantic content of utterances but mainly focuses on acoustic impairment of PWA. In the future, we propose to establish another neural network that aims at characterizing language impairment of PWA. A combined neural network considering both acoustic and language impairments will be investigated afterwards.

8 Conclusions

This paper presents an investigation on end-to-end approaches to automatic speech assessment for Cantonese-speaking PWA. The CNN model trained with equal-length utterances using Log-Mel filterbank features shows comparable performance to the conventional two-step assessment method. The experimental results confirm the effectiveness of using CNN model for automatic extraction of pathological-related features. With the CAM visualization technique, it is observed that the CNN-learned features show similar physical meaning to hand-crafted features. It suggests that applying end-to-end approach is able to improve the efficiency of developing assessment system and save significant amount of manual work.

Acknowledgements This research was partially supported by a GRF project grant (Ref: CUHK14227216) from the Hong Kong Research Grants Council, a Direct Grant from the CUHK Research Committee, the CUHK Research Sustainability Fund, and the CUHK Shenzhen Research Institute. The Cantonese AphasiaBank project was supported by a fund from the National Institutes of Health (project number: NIH-R01-DC010398).

References

1. Benson, D.F., Benson, D.F., Ardila, A. (1996). *Aphasia: A Clinical Perspective*, (pp. 89–98). Oxford: Oxford University Press.
2. Adam, H. (2014). Dysprosody in aphasia: an acoustic analysis evidence from palestinian arabic. *Journal of Language and Linguistic Studies*, 10(1), 153–162.
3. National Aphasia Association (2018). Aphasia definitions. <https://www.aphasia.org/aphasia-definitions/>, accessed 9 August 2018.
4. Wikipedia contributors (2018a). Aphasia — Wikipedia, the free encyclopedia. <https://en.wikipedia.org/wiki/Aphasia>, accessed 10 September 2018.
5. Wikipedia Contributors (2018). Anomic aphasia — Wikipedia, the free encyclopedia. https://en.wikipedia.org/wiki/Anomic_aphasia, accessed 10 September 2018.
6. Peintner, B., Jarrold, W., Vergyri, D., Richey, C., Tempini, M.L.G., Ogar, J. (2008). Learning diagnostic models using speech and language measures. In *Proceedings of Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBS)* (pp. 4648–4651). Vancouver: IEEE.
7. Fraser, K.C., Rudzicz, F., Rochon, E. (2013). Using text and acoustic features to diagnose progressive aphasia and its subtypes. In *Proceedings of Annual Conference of the International Speech Communication Association (INTERSPEECH)* (pp. 2177–2181). Lyon: ISCA.
8. Fraser, K.C., Meltzer, J.A., Graham, N.L., Leonard, C., Hirst, G., Black, S.E., Rochon, E. (2014). Automated classification of primary progressive aphasia subtypes from narrative speech transcripts. *Cortex*, 55, 43–60.
9. Fraser, K.C., Rudzicz, F., Graham, N., Rochon, E. (2013). Automatic speech recognition in the diagnosis of primary progressive aphasia. In *Proceedings of the 4th Workshop on Speech and Language Processing for Assistive Technologies (SLPAT)* (pp. 47–54). Grenoble: ACL/ISCA Special Interest Group.
10. Le, D., & Provost, E.M. (2016). Improving automatic recognition of aphasic speech with Aphasiabank. In *Proceedings of Annual Conference of the International Speech Communication Association (INTERSPEECH)* (pp. 2681–2685). San Francisco: ISCA.
11. Le, D., Licata, K., Provost, E.M. (2018). Automatic quantitative analysis of spontaneous aphasic speech. *Speech Communication*, 100, 1–12.
12. Kohlschein, C., Klischies, D., Meisen, T., Schuller, B.W., Werner, C.J. (2018). Automatic processing of clinical aphasia data collected during diagnosis sessions: challenges and prospects.

In *Proceedings of resources and processing of linguistic, para-linguistic and extra-linguistic Data from people with various forms of cognitive/psychiatric impairments (RaPID-2), satellite of the 11th Language Resources and Evaluation Conference (LREC)* (pp. 11–18). Miyazaki: ELRA.

13. Lee, T., Kong, A., Chan, V., Wang, H. (2013). Analysis of auto-aligned and auto-segmented oral discourse by speakers with aphasia: a preliminary study on the acoustic parameter of duration. *Procedia, Social and Behavioral Sciences*, 94, 71–72.
14. Lee, T., Liu, Y., Huang, P.W., Chien, J.T., Lam, W.K., Yeung, Y.T., Law, T.K., Lee, K.Y., Kong, A.P.H., Law, S.P. (2016). Automatic speech recognition for acoustical analysis and assessment of cantonese pathological voice and speech. In *Proceedings of International Conference on Acoustics, Speech and Signal Processing (ICASSP)* (pp. 6475–6479). Shanghai: IEEE.
15. Qin, Y., Lee, T., Kong, A.P.H. (2018a). Automatic speech assessment for aphasic patients based on syllable-level embedding and supra-segmental duration features. In *Proceedings of International Conference on Acoustics, Speech, and Signal Processing (ICASSP)* (pp. 5994–5998). Calgary: IEEE.
16. Qin, Y., Lee, T., Feng, S., Kong, A.P.H. (2018b). Automatic speech assessment for people with aphasia using TDNN-BLSTM with multi-task learning. In *Proceedings of Annual Conference of the International Speech Communication Association (INTERSPEECH)* (pp. 3418–3422). Hyderabad: ISCA.
17. Graves, A., & Jaitly, N. (2014). Towards end-to-end speech recognition with recurrent neural networks. In *Proceedings of International Conference on Machine Learning (ICML)* (pp. 1764–1772). Beijing: IMLS.
18. Wu, Y., Schuster, M., Chen, Z., Le, Q.V., Norouzi, M., Macherey, W., Krikun, M., Cao, Y., Gao, Q., Macherey, K., et al. (2016). Google's neural machine translation system: Bridging the gap between human and machine translation. *arXiv:cs.CL/1609.08144*.
19. Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. *Neural Computation*, 9(8), 1735–1780.
20. Cho, K., Van Merriënboer, B., Bahdanau, D., Bengio, Y. (1259). On the properties of neural machine translation: Encoder-decoder approaches. *arXiv:cs.CL/1409*.
21. Tang, Y., Huang, Y., Wu, Z., Meng, H., Xu, M., Cai, L. (2016). Question detection from acoustic features using recurrent neural network with gated recurrent unit. In *Proceedings of International Conference on Acoustics, Speech, and Signal Processing (ICASSP)* (pp. 6125–6129). Shanghai: IEEE.
22. Rana, R. (2016). Gated Recurrent Unit (GRU) for emotion classification from noisy speech. *arXiv:cs.HC/1612.07778*.
23. Chung, H., Lee, Y.K., Lee, S.J., Park, J.G. (2017). Spoken english fluency scoring using convolutional neural networks. In *Proceedings of Conference of the Oriental Chapter of the International Coordinating Committee on Speech Databases and Speech I/O Systems and Assessment (O-COCOSDA)* (pp. 1–6). Seoul: COCOSDA.
24. Vásquez-Correa, J., Orozco-Arroyave, J.R., Nöth, E. (2017). Convolutional neural network to model articulation impairments in patients with Parkinson's disease. In *Proceedings of Annual Conference of the International Speech Communication Association (INTERSPEECH)* (pp. 314–318). Stockholm: ISCA.
25. Zhou, B., Khosla, A., Lapedriza, A., Oliva, A., Torralba, A. (2016). Learning deep features for discriminative localization. In *Proceedings of Conference on Computer Vision and Pattern Recognition (CVPR)* (pp. 2921–2929). Las Vegas: IEEE.
26. Selvaraju, R.R., Cogswell, M., Das, A., Vedantam, R., Parikh, D., Batra, D. (2017). Grad-CAM: Visual explanations from deep networks via gradient-based localization. In *Proceedings of International Conference on Computer Vision (ICCV)* (pp. 618–626). Venice: CVF.
27. Wu, Y., & Lee, T. (2019). Enhancing sound texture in CNN-based acoustic scene classification. *arXiv:cs.SD/1901.01502*.
28. Kong, A.P.H., & Law, S.P. (2018). Cantonese Aphasiabank: An annotated database of spoken discourse and co-verbal gestures by healthy and language-impaired native cantonese speakers. *Behavior research methods*, pp. 1–14.
29. MacWhinney, B., Fromm, D., Forbes, M., Holland, A. (2011). Aphasiabank: Methods for studying discourse. *Aphasiology*, 25(11), 1286–1307.
30. Kong, A.P.H., Law, S.P., Kwan, C.C.Y., Lai, C., Lam, V. (2015). A coding system with independent annotations of gesture forms and functions during verbal communication: Development of a database of speech and gesture (DoSaGE). *Journal of Nonverbal Behavior*, 39(1), 93–111.
31. Yiu, E.M. (1992). Linguistic assessment of Chinese-speaking aphasics: Development of a Cantonese Aphasia Battery. *Journal of Neurolinguistics*, 7(4), 379–424.
32. Kong, A.P.H. (2016). Analysis of neurogenic disordered discourse production: From theory to practice. Routledge.
33. Krizhevsky, A., Sutskever, I., Hinton, G.E. (2012). ImageNet classification with deep convolutional neural networks. In *Proceedings of Annual Conference on Neural Information Processing Systems (NIPS)* (pp. 1097–1105). Lake Tahoe: NIPS Foundation.
34. Simonyan, K., & Zisserman, A. (1556). Very deep convolutional networks for large-scale image recognition. *arXiv:cs.CV/1409*.
35. Lin, M., Chen, Q., Yan, S. (2013). Network in network. *arXiv:cs.NE/1312.4400*.
36. Kingma, D.P., & Ba, J. (2014). Adam: A method for stochastic optimization. *arXiv:cs.LG/1412.6980*.
37. Paszke, A., Gross, S., Chintala, S., Chanan, G., Yang, E., DeVito, Z., Lin, Z., Desmaison, A., Antiga, L., Lerer, A. (2017). Automatic differentiation in pytorch In *Conference on Neural Information Processing Systems Workshop (NIPS-w)*, NIPS Foundation, Long Beach.
38. Vuk, M., & Cuk, T. (2006). ROC Curve, lift chart and calibration plot. *Metodoloski zvezki*, 3(1), 89.
39. Fawcett, T. (2006). An introduction to ROC analysis. *Pattern Recognition Letters*, 27(8), 861–874.
40. McNemar, Q. (1947). Note on the sampling error of the difference between correlated proportions or percentages. *Psychometrika*, 12(2), 153–157.
41. Dietterich, T.G. (1998). Approximate statistical tests for comparing supervised classification learning algorithms. *Neural Computation*, 10(7), 1895–1923.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.