



# Speech Intelligibility Assessment of Dysarthric Speech by using Goodness of Pronunciation with Uncertainty Quantification

Eun Jung Yeo<sup>\*1</sup>, Kwanghee Choi<sup>\*2</sup>, Sunhee Kim<sup>1</sup>, Minhwa Chung<sup>1</sup>

<sup>1</sup>Seoul National University, Republic of Korea

<sup>2</sup>Carnegie Mellon University, United States of America

ej.yeo@snu.ac.kr, kwanghec@andrew.cmu.edu, sunhkim@snu.ac.kr, mchung@snu.ac.kr

## Abstract

This paper proposes an improved Goodness of Pronunciation (GoP) that utilizes Uncertainty Quantification (UQ) for automatic speech intelligibility assessment for dysarthric speech. Current GoP methods rely heavily on neural network-driven overconfident predictions, which is unsuitable for assessing dysarthric speech due to its significant acoustic differences from healthy speech. To alleviate the problem, UQ techniques were used on GoP by 1) normalizing the phoneme prediction (entropy, margin, maxlogit, logit-margin) and 2) modifying the scoring function (scaling, prior normalization). As a result, prior-normalized maxlogit GoP achieves the best performance, with a relative increase of 5.66%, 3.91%, and 23.65% compared to the baseline GoP for English, Korean, and Tamil, respectively. Furthermore, phoneme analysis is conducted to identify which phoneme scores significantly correlate with intelligibility scores in each language.

**Index Terms:** dysarthric speech, speech intelligibility, automatic assessment, goodness of pronunciation, uncertainty quantification

## 1. Introduction

Dysarthria is a motor speech disorder caused by weakness or paralysis of the articulators [1]. People with dysarthria often suffer from degraded speech intelligibility, repeated communication failures, and ultimately low quality of life. Accordingly, dysarthric speech assessments regarding speech intelligibility are conducted to check the patient's status and track the effectiveness of treatments [2]. While the common way of dysarthric speech assessment is perceptual evaluation, the method is often subjective and laborious. Therefore, automatic speech assessment with objective and rapid results can assist clinicians in diagnosis and treatment planning.

There are two main approaches to automatic assessment of dysarthric speech. The first approach is to propose a list of hand-crafted features that are expected to capture the characteristics of dysarthric speech. Explored feature sets include voice quality features [3], prosody features [4], articulation or pronunciation features [5, 6], and their combinations [7, 8]. This approach has the benefit of having medical implications, as it provides a transparent understanding of the features employed for automatic assessment. Nevertheless, this approach has the drawback that features that could be valuable in the assessment could be discarded during feature extraction. The second approach involves leveraging the capabilities of neural networks (NNs), which can achieve better results by using raw inputs. [9, 10]. However, due to the black-box nature of NNs, the

\* Equal contribution.

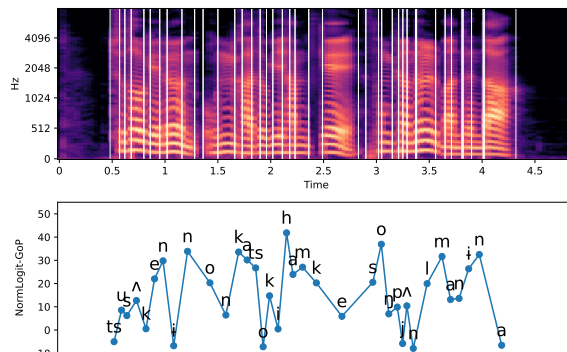


Figure 1: Example of GoP scores for each phone within an utterance.

approach limits interpretability, which clinicians often crave for.

Recent studies have attempted to integrate the benefits of both approaches, by enforcing the neural networks to learn the intermediate labels used for perceptual assessment, such as voice quality, articulation precision, nasality, and prosody [11, 12]. Furthermore, certain studies focused on decoding misarticulation characteristics, which are the prominent aspect of dysarthric speech across languages [8] and a significant factor that influences speech intelligibility [13]. For instance, the framework that measures the level of phonetic impairment was proposed by utilizing the activations of the hidden neuron [14, 15]. While the method can provide overall phonetic characteristics of utterances, it is unable to provide assessments at the level of individual phonemes, which can help clinicians to pinpoint specific phonemes that require pronunciation training.

A common approach of phoneme-level speech assessment is to use the parallel NN that employs parallel datasets. The NN was trained using the same set of utterances recorded by both healthy speakers and patients to learn how to distinguish whether each phone in the utterances was from healthy or disordered speech [16, 17]. However, obtaining parallel datasets is a challenging task, especially for disordered speech. Moreover, this approach often constrains the analysis to pre-defined speech materials, which may not capture the natural speech patterns utilized in everyday communication.

Another conventional approach of phoneme-level pronunciation evaluation is Goodness of Pronunciation (GoP) [18]. GoP, which is defined as the degree of similarity between produced and canonical pronunciation of phonemes, has two advantages in automatic speech assessments. First, it provides detailed information on which phonemes are mispronounced and to what extent each phoneme is atypical. As demonstrated in Figure 1, one can analyze the scores of each phoneme within a single utterance. Second, it does not necessitate a parallel dataset for model training. While GoP has often been applied

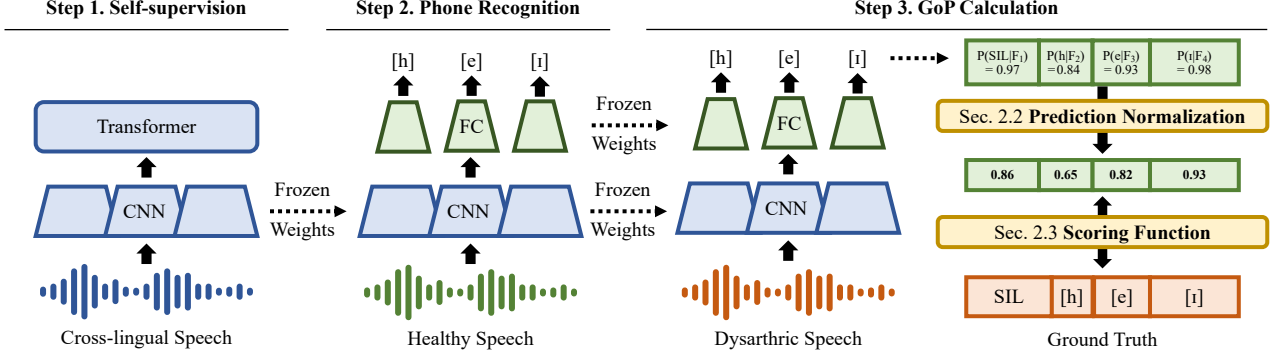


Figure 2: Overview of the improved GoP with UQ methods.

to assess non-native (L2) speech pronunciation, some studies have also verified its potential use in assessing speech disorders. [19, 20].

With the development of neural networks (NNs), variants of GoP which employ probabilities from the state-of-the-art neural networks have been suggested [12, 21, 22]. However, using these probabilities without taking into account the modern NNs’ tendency towards overconfidence can result in inaccurate conclusions: NNs often generate probabilities close to 1.0 even when their predictions are incorrect [23]. Since GoP relies heavily on probabilities, this can be especially problematic. The issue is compounded when the model encounters out-of-distribution (OOD) inputs, which are data that differ significantly from the training data’s distribution [24], such as dysarthric speech for healthy speech. To alleviate such issues, Uncertainty Quantification (UQ) techniques, techniques to combat OOD problems [23, 24], can be applied.

This paper proposes improved GoP for automatic speech intelligibility assessment for dysarthric speech by employing Uncertainty Quantification (UQ) methods in two ways: (1) normalizing the phoneme prediction and (2) modifying the scoring function. As pathological speech greatly differs in acoustics from healthy speech [25], dysarthric speech can be also understood as OOD input. Therefore, we employ conventional UQ methods to improve GoP calculations for dysarthric speech assessment. To assess the effectiveness of the improved GoP with UQ techniques, three dysarthric datasets, namely, UASpeech English, QoLT Korean, and SSNCE Tamil dataset, are utilized. To summarize, this paper redefines the current versions of GoP from a UQ standpoint and evaluates the effectiveness of UQ methods in improving GoPs.

## 2. Proposed approach

The study applies various conventional UQ methods to calculate GoP scores, which are demonstrated in Figure 2. We release the source code of all the experiments.<sup>1</sup>

### 2.1. Prerequisite: Goodness of Pronunciation (GoP)

In this subsection, we provide a succinct summary of the existing GoP-based methods from the uncertainty quantification (UQ) perspective. Starting from GMM-GoP, Equation (1) presents the definition of the corresponding method: Given a phone  $p$  during frames  $f \in F$  with frame-wise phone probability and its logits as  $P^f(p|f)$  and  $L^f(p|f)$ , and the total phone

<sup>1</sup><https://github.com/juice500ml/dysarthria-gop>

set as  $Q$ , **GMM-GoP** [18] is defined as an averaged log probability across the phone duration,  $\mathbb{E}_F[\log P^f(p|f)]$ :

$$s_{\text{GMM-GoP}}(p) = \frac{1}{|F|} \sum_{f \in F} \log \frac{e^{L^f(p|f)}}{\sum_{q \in Q} e^{L^f(q|f)}}. \quad (1)$$

Averaging the log probabilities can be seen as a form of temporal ensembling [26], which is a well-known UQ method, as it combines the predictions from multiple frames into a single estimate. Further, directly using the probability output is often used as a baseline for OOD [24].

Another popular baseline is **NN-GoP** [21], which improved GMM-GoP by leveraging the development of deep neural networks and modifying the score function:

$$\bar{P}(p|F) = \mathbb{E}_F[P^f(p|f)] = \frac{1}{|F|} \sum_{f \in F} P^f(p|f), \quad (2)$$

$$s_{\text{NN-GoP}}(p) = \log \bar{P}(p|F) - \max_{q \in Q} \log \bar{P}(q|F). \quad (3)$$

Finally, **DNN-GoP** [21] normalizes the phone probability with the phone prior [27]:

$$s_{\text{DNN-GoP}}(p) = \bar{P}(p|F) / P(p). \quad (4)$$

### 2.2. Normalizing the phoneme predictions

There are two commonly used ways to calibrate the posterior prediction  $P(p|F)$  by modifying its logit  $L(p|F)$ . One is to normalize by removing the influence of the prior  $P(p)$  [21, 27] (**Prior**), and the other is to reduce the peakiness by temperature scaling [23] **Scale**:

$$L^f(p, f) = L^f(p|f) - \log P(p), \quad (5)$$

$$L_T^f(p|f) = L^f(p|f) / T, \quad (6)$$

where  $T$  is the hyperparameter and the modified predictions are the softmax function’s output of the corresponding logits.

Normalizing via the prior is the same with the idea of DNN-GoP, where it is commonly applied to disentangle the training distribution of the phone recognizer, where majority classes are often overconfident than minority classes [27]. Temperature scaling is also commonly used as a baseline for UQ, as it avoids the peaky distribution of posterior probabilities.

### 2.3. Modifying the scoring function

We first employ one of the most common methods to measure the data uncertainty: Entropy  $H$  and Margin  $M$ . **Entropy** mea-

sure the uncertainty associated with a set of possible outcomes:

$$s_H(p) = - \sum_{q \in Q} \bar{P}(q|F) \log \bar{P}(p|F). \quad (7)$$

Specifically, the entropy represents the average amount of information needed to specify which outcome was actually observed. Entropy does not require ground truth labels, so it is often used when obtaining the labels is expensive.

**Margin** refers to the difference between the true class and the highest class probabilities for a given prediction:

$$s_M(p) = \bar{P}(p|F) - \max_{q \in Q - \{p\}} \bar{P}(q|F). \quad (8)$$

Note that the equation is similar to that of NN-GoP except the above definition excludes the true phone as  $Q - \{p\}$ .

On the other hand, the **MaxLogit** [28] and **LogitMargin** involves directly utilizing the logits. Softmax function is often known to squash the useful information inside logits, so that it can normalize the sum into one. Hence, one can apply the idea of both GMM-GoP (directly using the probability) and NN-GoP (using the Margin):

$$s_{\text{MaxLogit}}(p) = \frac{1}{|F|} \sum_{f \in F} L^f(p|f) = \bar{L}(p|F), \quad (9)$$

$$s_{\text{LogitMargin}}(p) = \bar{L}(p|F) - \max_{q \in Q - \{p\}} \bar{L}(q|F). \quad (10)$$

### 3. Experimental setting

#### 3.1. Datasets

To train the acoustic model to learn the distribution of healthy phonemes, we use the Common Phone dataset [29] and the L2-ARCTIC dataset [30]. To evaluate the efficacy of our proposed approach, three dysarthric datasets are utilized: UASpeech English dataset [31], QoLT Korean dataset [32], and SSNCE Tamil dataset [33]. All three datasets contain dysarthric speech from speakers suffering from Cerebral Palsy. We focus the analysis on sentences from QoLT and SSNCE dysarthric datasets, since the Common Phone and L2-ARCTIC datasets solely consist of sentences. Word materials are analyzed for the UASpeech dataset, since the dataset contains words only.

##### 3.1.1. Common Phone dataset and L2-ARCTIC dataset

The acoustic model is trained on healthy speech using the Common Phone dataset and the L2-ARCTIC dataset, which were selected for their extensive phoneme coverage and phonetic annotations. These datasets are expected to cover most of the phonemes used in English, Korean, and Tamil. **Common Phone dataset** [29] is a gender-balanced, multilingual corpus with six languages. Comprised of more than 11,000 speakers, the dataset includes around 116 hours of speech. **L2-ARCTIC dataset** [30] is a speech corpus often used for detecting mispronunciations in non-native English speakers. It includes recordings from 24 speakers with a balanced distribution of gender and first language, representing six different countries. On average, each speaker has around 67.7 minutes of speech, which has a total duration of approximately 27.1 hours.

##### 3.1.2. UASpeech English dysarthric dataset

UASpeech dataset [31] is a publicly-available English dysarthric speech dataset, which contains 15 dysarthria speakers (11 males, 4 females) and 13 aged-matched healthy speakers (9 males, 4 females). Speakers were classified based on the

scores on the Frenchay Dysarthria Assessment (FDA) [34]: 5 mild speakers (score 1), 3 moderate-to-severe speakers (score 2), 3 moderate-to-severe speakers (score 3), and 4 severe speakers (score 4). Montreal Forced Aligner (MFA) [35] is employed to extract phoneme-level alignments.

##### 3.1.3. QoLT Korean dysarthric dataset

Quality of Life Technology (QoLT) dataset [32] is a privately held dataset of Korean dysarthric speech. The corpus consists of 70 dysarthric speakers (45 males, 25 females) and 10 healthy speakers (5 males, 5 females). Each speaker recorded five phonetically balanced sentences twice. Five speech pathologists were asked to determine the intelligibility levels of the speakers on a 5-point Likert scale. With a score of 0 considered healthy, the dataset holds 25 mild (score 1), 26 mild-to-moderate (score 2), 12 moderate-to-severe (score 3), and 7 severe (score 4) intelligibility level speakers. Accordingly, 100 healthy utterances and 700 dysarthric utterances are used for the experiment. After using MFA to align the phonemes, two linguists further fixed the automated alignment for better quality.

##### 3.1.4. SSNCE Tamil dysarthric dataset

SSNCE dataset [33] is a Tamil dysarthric speech corpus available by request. The dataset includes recordings from 20 dysarthric speakers (13 males and 7 females) and 10 healthy speakers (5 males and 5 females). The dataset groups dysarthric speakers based on their speech intelligibility scores, which were marked by two speech pathologists on a 7-point Likert scale. A score of 0 considered healthy, score 1 and 2 are grouped into mild (score 1), score 3 and 4 into moderate (score 2), and score 5 and 6 into severe (score 3). There were different numbers of speakers in each score category: 7 with mild, 10 with moderate, and 3 with severe. For the experiment, we used a total of 5,200 utterances from the dysarthric speakers and 2,600 utterances from the healthy speakers, with 260 unique sentences recorded from each speaker. For forced alignments, we use the time-aligned phonetic transcriptions provided by the dataset.

#### 3.2. Experimental details

In this study, we evaluated GoP performance by following the approach of the previous study [20]. Concretely, rather than evaluating the models based on their accuracy in mispronunciation detection, our objective was to calculate the average GoP score for each utterance and compare their correlation with the intelligibility scores.

##### 3.2.1. Phoneme Prediction

To fairly compare the performances between various GoP scoring functions, we extract the posterior probabilities from the common cross-lingual Wav2Vec 2.0 XLS-R model [36] instead of using the acoustic model from Kaldi, following recent literature [12]. We slightly modify the architecture by attaching the linear phone prediction head to the convolutional layer, not the transformer layer, to avoid extensive computational overhead and preserve the phonetic characteristics in convolutional features [37]. Also, the loss function is simplified by removing the adaptive pooling, where we observed that the final performance difference was negligible. AdamW optimizer [38] is used with the default learning rate of 0.001 for three epochs. As we only trained the linear prediction head, the final performance was not sensitive to other hyperparameters. Refer to [12] and our source code<sup>1</sup> for more details.

Table 1: Kendall’s rank coefficient between GoP & intelligibility severity levels. A higher absolute value indicates a stronger correlation between the two variables.

Method	Norm.	Scoring Func.	English	Korean	Tamil
Baseline	None	GMM [18, 39]	-0.2049	-0.5237	-0.3571
	None	NN [21]	-0.1536	-0.4687	-0.4003
	Prior	DNN-GoP [21]	-0.1836	-0.4237	-0.4681
Proposed	None	Entropy	-0.1831	-0.2643	-0.3251
		Margin	-0.1628	-0.4434	-0.4445
		MaxLogit	-0.2164	-0.5440	-0.5786
		LogitMargin	-0.1732	-0.4753	-0.5158
	Scale	Entropy	-0.1755	-0.1974	-0.2263
		Margin	-0.1260	-0.4444	-0.4210
		MaxLogit	-0.2164	-0.5440	-0.5786
		LogitMargin	-0.1732	-0.4753	-0.5158
	Prior	Entropy	-0.1833	-0.2645	-0.3254
		Margin	-0.1630	-0.4432	-0.4447
		<b>MaxLogit</b>	<b>-0.2165</b>	<b>-0.5442</b>	<b>-0.5788</b>
		LogitMargin	-0.1733	-0.4753	-0.5160

### 3.2.2. Baselines

Three baseline experiments are conducted: GMM-GoP [18], NN-GoP [21], and DNN-GoP. We compare the baselines with the UQ methods introduced in Section 2.2 and Section 2.3 by using the same phoneme probabilities. As we aim to see the correlations between GoP scores (continuous) and intelligibility scores (ordinal), we utilize the Kendall Rank Coefficient  $\tau$  to compare the performances. Kendall’s  $\tau$  measures the strength of the relationships, with a higher absolute coefficient indicating higher correlations between the two variables.

## 4. Experimental results

### 4.1. Correlation between GoPs and intelligibility scores

Table 1 demonstrates the performances of both baseline and proposed experiments, with the best performance indicated in bold. GoP with prior normalized MaxLogit perform the best on all the languages among the baselines and the UQ methods, achieving  $-21.65\%$ ,  $-54.42\%$ ,  $-57.88\%$  correlation, for English, Korean, and Tamil, respectively.

The results of the baseline experiments show that GMM-GoP has the highest correlation for English and Korean at  $-20.49\%$  and  $-52.37\%$ , respectively, while DNN-GoP performs best for Tamil with a correlation of  $-46.81\%$ . For the proposed experiments, GoP without normalization generally shows lower performance than the baseline, except for MaxLogit-based GoP. Additionally, while scaling normalization has minimal impact, prior normalization has a positive effect on GoP performance for all languages. Furthermore, when entropy-based, probability-based (Margin), and logit-based (MaxLogit, LogitMargin) GoP variants are compared, the logit-based GoPs show the highest correlations to the intelligibility scores. Additionally, performances on English are notably lower than that of Tamil and Korean. We suspect that automatically generated alignment causes the degradation to occur [40], for the severe cases where alignment becomes much more challenging. We aim to mitigate this issue in our future work.

### 4.2. Analysis on phonemes

Figure 3 illustrates the GoP distribution between two Korean phonemes /i/ and /m/. While the distribution of /i/ differs significantly, the distribution of /m/ is similar across all severity levels. This finding suggests that certain phonemes have more impact power for severity levels based on speech intelligibil-




Figure 3: Kendall’s  $\tau$  distributions for two phonemes by severity. 0:healthy, 1:mild, 2:mild-to-mod., 3:mod.-to-sev., 4:severe.

ity, which is consistent with previous findings [17]. Identifying which phoneme pronunciation scores highly correlate to speech intelligibility can be useful in diagnosis and treatment planning.

We conducted a quantitative analysis on which phoneme scores highly correlate to speech intelligibility by languages. Kendall’s  $\tau$  is calculated for each phoneme between our best-performing Prior+MaxLogit GoP score and the intelligibility scores. The top-5 phonemes based on correlation are as follows: English: /aɪ/, /ɛ/, /aʊ/, /z/, /ʧ/; Korean: /i/, /s/, /n/, /a/, /ʌ/; Tamil: /ɛ/, /h/, /tʃ/, /z/, /aɪ/. In summary, fricative sounds (/s/, /ɛ/, /ʃ/, /z/) strongly correlate to speech intelligibility across languages, consistent with the previous results [41]. Affricates (/tʃ/, /ʧ/) and diphthongs (/aɪ/, /aʊ/) are shared as the top-5 phoneme list for English and Tamil. This can be explained by the complexity in the articulation of affricates and diphthongs leading to difficulties in correct pronunciation for speakers with lower speech intelligibility. On the other hand, Korean show higher correlations for nasal (/n/) and monophthongs (/a/, /i/, /ʌ/). This may be again related to the movement of the articulators, such as the tongue, velum, and jaw. We additionally provide the correlation scores of all the phonemes in our repository.<sup>1</sup>

## 5. Conclusion

This paper proposes an improved GoP for dysarthria speech intelligibility assessment by using UQ methods. Expected to alleviate the problem of modern NN’s overconfidence, especially for disordered speech, tested UQ methods include (1) normalization of phoneme prediction and (2) modification of the scoring function. The experiments were carried out on dysarthric speech datasets in English, Korean, and Tamil. According to the experimental results, the prior normalized MaxLogit GoP showed the best performance, outperforming both the traditional GoPs and other proposed GoP variants. Furthermore, to verify the usefulness of our proposed method, an analysis of which phoneme pronunciation scores highly correlate to speech intelligibility was conducted.

## 6. Acknowledgements

This work was supported by Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea government (MSIT) (No.2022-0-00223, Development of digital therapeutics to improve communication ability of autism spectrum disorder patients).

## 7. References

- [1] F. L. Darley, J. R. Brown, and N. P. Goldstein, "Dysarthria in multiple sclerosis," *Journal of Speech and Hearing research*, vol. 15, no. 2, pp. 229–245, 1972.
- [2] R. D. Kent, G. Weismer, J. F. Kent, and J. C. Rosenbek, "Toward phonetic intelligibility testing in dysarthria," *Journal of Speech and Hearing Disorders*, vol. 54, no. 4, pp. 482–499, 1989.
- [3] N. Narendra and P. Alku, "Automatic assessment of intelligibility in speakers with dysarthria from coded telephone speech using glottal features," *Computer Speech & Language*, vol. 65, p. 101117, 2021.
- [4] A. Hernandez, S. Kim, and M. Chung, "Prosody-based measures for automatic severity assessment of dysarthric speech," *Applied Sciences*, vol. 10, no. 19, p. 6999, 2020.
- [5] M. J. Kim and H. Kim, "Automatic assessment of dysarthric speech intelligibility based on selected phonetic quality features," in *ICCHP*. Springer, 2012, pp. 447–450.
- [6] Y. Liu, N. Penttilä, T. Ihalainen, J. Lintula, R. Convey, and O. Räsänen, "Language-independent approach for automatic computation of vowel articulation features in dysarthric speech assessment," *TASLP*, vol. 29, pp. 2228–2243, 2021.
- [7] J. C. Vásquez-Correa, J. Orozco-Arroyave, T. Bocklet, and E. Nöth, "Towards an automatic evaluation of the dysarthria level of patients with parkinson's disease," *Journal of communication disorders*, vol. 76, pp. 21–36, 2018.
- [8] E. J. Yeo, S. Kim, and M. Chung, "Multilingual analysis of intelligibility classification using english, korean, and tamil dysarthric speech datasets," in *Oriental-COCOSDA*, 2022, pp. 1–6.
- [9] E. J. Yeo, K. Choi, S. Kim, and M. Chung, "Automatic severity assessment of dysarthric speech by using self-supervised model with multi-task learning," in *ICASSP*, 2023.
- [10] A. A. Joshy and R. Rajan, "Dysarthria severity classification using multi-head attention and multi-task learning," *Speech Communication*, vol. 147, pp. 1–11, 2023.
- [11] M. Tu, V. Berisha, and J. Liss, "Interpretable objective assessment of dysarthric speech based on deep neural networks," in *Interspeech*, 2017, pp. 1849–1853.
- [12] X. Xu, Y. Kang, S. Cao, B. Lin, and L. Ma, "Explore wav2vec 2.0 for mispronunciation detection," in *Interspeech*, 2021, pp. 4428–4432.
- [13] M. S. De Bodt, M. E. H.-D. Huici, and P. H. Van De Heyning, "Intelligibility as a linear combination of dimensions in dysarthric speech," *Journal of communication disorders*, vol. 35, no. 3, pp. 283–292, 2002.
- [14] S. Abderrazek, C. Fredouille, A. Ghio, M. Lalain, C. Meunier, and V. Woisard, "Interpreting deep representations of phonetic features via neuro-based concept detector: Application to speech disorders due to head and neck cancer," *TASLP*, vol. 31, pp. 200–214, 2022.
- [15] —, "Validation of the neuro-concept detector framework for the characterization of speech disorders: A comparative study including dysarthria and dysphonia," in *Interspeech*, 2022.
- [16] G. F. Miller, J. C. Vásquez-Correa, and E. Nöth, "Assessing the dysarthria level of parkinson's disease patients with gmm-ubm supervectors using phonological posteriors and diadochokinetic exercises," in *TSD*, 2020, pp. 356–365.
- [17] S. Quintas, J. Maclair, V. Woisard, and J. Pinquier, "Automatic assessment of speech intelligibility using consonant similarity for head and neck cancer," in *Interspeech*, 2022.
- [18] 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.
- [19] T. Pellegrini, L. Fontan, J. Maclair, J. Farinas, and M. Robert, "The goodness of pronunciation algorithm applied to disordered speech," in *Fifteenth Annual Conference of the International Speech Communication Association*, 2014.
- [20] L. Fontan, T. Pellegrini, J. Olcoz, and A. Abad, "Predicting disordered speech comprehensibility from goodness of pronunciation scores," in *SLPAT@Interspeech*, 2015.
- [21] 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.
- [22] S. Cheng, Z. Liu, L. Li, Z. Tang, D. Wang, and T. F. Zheng, "Asr-free pronunciation assessment," in *Interspeech*, 2020.
- [23] C. Guo, G. Pleiss, Y. Sun, and K. Q. Weinberger, "On calibration of modern neural networks," in *ICML*, 2017, pp. 1321–1330.
- [24] D. Hendrycks and K. Gimpel, "A baseline for detecting misclassified and out-of-distribution examples in neural networks," in *ICLR*, 2017.
- [25] J. Wilson, Bronagh Blaney, "Acoustic variability in dysarthria and computer speech recognition," *Clinical Linguistics & Phonetics*, vol. 14, no. 4, pp. 307–327, 2000.
- [26] S. Laine and T. Aila, "Temporal ensembling for semi-supervised learning," in *ICLR*, 2017.
- [27] Y. Hong, S. Han, K. Choi, S. Seo, B. Kim, and B. Chang, "Disentangling label distribution for long-tailed visual recognition," in *CVPR*, 2021, pp. 6626–6636.
- [28] D. Hendrycks, S. Basart, M. Mazeika, A. Zou, J. Kwon, M. Mostajabi, J. Steinhardt, and D. Song, "Scaling out-of-distribution detection for real-world settings," in *ICML*. PMLR, 2022, pp. 8759–8773.
- [29] P. Klumpp, T. Arias-Vergara, P. A. Pérez-Toro, E. Nöth, and J. R. Orozco-Arroyave, "Common phone: A multilingual dataset for robust acoustic modelling," *arXiv preprint arXiv:2201.05912*, 2022.
- [30] G. Zhao, S. Sonsaat, A. Silpachai, I. Lucic, E. Chukharev-Hudilainen, J. Levis, and R. Gutierrez-Osuna, "L2-arctic: A non-native english speech corpus," in *Interspeech*, 2018.
- [31] H. Kim, M. Hasegawa-Johnson, A. Perlman, J. Gunderson, T. S. Huang, K. Watkin, and S. Frame, "Dysarthric speech database for universal access research," in *Ninth Annual Conference of the International Speech Communication Association*, 2008.
- [32] D.-L. Choi, B.-W. Kim, Y.-W. Kim, Y.-J. Lee, Y. Um, and M. Chung, "Dysarthric speech database for development of qolt software technology," in *LREC*, 2012, pp. 3378–3381.
- [33] M. TA, T. Nagarajan, and P. Vijayalakshmi, "Dysarthric speech corpus in tamil for rehabilitation research," in *Region TENCON*. IEEE, 2016, pp. 2610–2613.
- [34] P. Enderby, "Frenchay dysarthria assessment," *British Journal of Disorders of Communication*, 1980.
- [35] M. McAuliffe, M. Socolof, S. Mihuc, M. Wagner, and M. Sonderegger, "Montreal forced aligner: Trainable text-speech alignment using kald," in *Interspeech*, 2017, pp. 498–502.
- [36] A. Babu, C. Wang, A. Tjandra, K. Lakhotia, Q. Xu, N. Goyal, K. Singh, P. von Platen, Y. Saraf, J. Pino, A. Baevski, A. Conneau, and M. Auli, "XLS-R: self-supervised cross-lingual speech representation learning at scale," in *Interspeech*, 2022.
- [37] K. Choi and E. J. Yeo, "Opening the black box of wav2vec feature encoder," *arXiv preprint arXiv:2210.15386*, 2022.
- [38] I. Loshchilov and F. Hutter, "Decoupled weight decay regularization," in *ICLR*, 2019.
- [39] J. Zhang, "Gmm-based gop (goodness of pronunciation) using kald," <https://github.com/jimbozhang/kaldi-gop>, 2020.
- [40] V. C. Mathad, T. J. Mahr, N. Scherer, K. Chapman, K. C. Hustad, J. Liss, and V. Berisha, "The impact of forced-alignment errors on automatic pronunciation evaluation," in *Interspeech*, 2021, pp. 1922–1926.
- [41] A. Hernandez, H.-y. Lee, and M. Chung, "Acoustic analysis of fricatives in dysarthric speakers with cerebral palsy," *Phonetics and Speech Sciences*, vol. 11, no. 3, pp. 23–29, 2019.