



**HAL**  
open science

## **Towards objective and interpretable speech disorder assessment: a comparative analysis of CNN and transformer-based models**

Malo Maisonneuve, Corinne Fredouille, Muriel Lalain, Alain Ghio, Virginie Woisard

### ► **To cite this version:**

Malo Maisonneuve, Corinne Fredouille, Muriel Lalain, Alain Ghio, Virginie Woisard. Towards objective and interpretable speech disorder assessment: a comparative analysis of CNN and transformer-based models. Interspeech 2024, Sep 2024, Kos Island, Greece. hal-04623180

**HAL Id: hal-04623180**

**<https://hal.science/hal-04623180>**

Submitted on 25 Jun 2024

**HAL** is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

# Towards objective and interpretable speech disorder assessment: a comparative analysis of CNN and transformer-based models

Malo Maisonneuve<sup>1</sup>, Corinne Fredouille<sup>1</sup>, Muriel Lalain<sup>2</sup>, Alain Ghio<sup>2</sup>, Virginie Woisard<sup>3</sup>

<sup>1</sup>LIA, Avignon University, France

<sup>2</sup>Aix-Marseille Univ, CNRS, LPL, Aix-en-Provence, France

<sup>3</sup>LNPL, Toulouse University and Toulouse Hospital, Toulouse, France

(malo.maisonneuve, corinne.fredouille)@univ-avignon.fr, muriel.lalain@univ-amu.fr

## Abstract

Head and Neck Cancers (HNC) significantly impact patients' ability to speak, affecting their quality of life. Commonly used metrics for assessing pathological speech are subjective, prompting the need for automated and unbiased evaluation methods. This study proposes a self-supervised Wav2Vec2-based model for phone classification with HNC patients, to enhance accuracy and improve the discrimination of phonetic features for subsequent interpretability purpose. The impact of pre-training datasets, model size, and fine-tuning datasets and parameters are explored. Evaluation on diverse corpora reveals the effectiveness of the Wav2Vec2 architecture, outperforming a CNN-based approach, used in previous work. Correlation with perceptual measures also affirms the model relevance for impaired speech analysis. This work paves the way for better understanding of pathological speech with interpretable approaches for clinicians, by leveraging complex self-learned speech representations.

**Index Terms:** speech disorders, Head and Neck Cancer, deep learning, phone classification, intelligibility, interpretability

## 1. Introduction

Head and Neck Cancers (HNCs) affect the upper respiratory and digestive tracts, including the oral cavity, pharynx, larynx, nasal cavity and salivary glands. Treatment of this cancer, either by radio-therapy, chemotherapy and/or surgery, can significantly impact the speech of HNC patients. Difficulties in communication with others negatively impacts the patients' quality of life. It is essential to correctly assess their speech production, identifying its level of impairment and what makes it atypical, so that they can benefit from the best possible care during rehabilitation sessions or for evaluating the efficacy of a prosthesis following a surgical procedure. Unfortunately, commonly used metrics such as severity or intelligibility measures are subjective, and prone to misjudgments, even by experts [1]. Furthermore, these metrics provide no information, other than a score, on the nature of the measured degradations. This highlights the crucial need for an automated, objective method of speech assessment that offers not only accurate results but also interpretable insights into the specific speech characteristics affected. Some papers investigated the assessment of HNCs patients' speech intelligibility and/or severity with relatively simple networks [2, 3, 4, 5], as well as using deep neural networks [6, 7, 8].

Recently, self-supervised learning (SSL) models have shown to be successful in capturing phonetic patterns and various speech features. In [9], authors have shown that Wav2Vec2 models [10] can learn certain phonetic concepts, and that they correctly preserve the manner and place of articulation. [11, 12] have also shown that phonetic information is contained in the

representations built by these models. Some research has been focused on using these models to automatically assess the speech severity level [13, 14, 15]. Other studies analysed how well diseases can be predicted by these models. For instance, A. Favaro et al. [16] compared interpretable speech features to embeddings produced by SSL models on predicting the presence of Parkinson's disease. They showed that using embeddings provides better detection accuracies at the cost of losing the insight into speech and language deterioration given by interpretable features. While being able to detect a disease and assess its severity is important, we believe it is as important to interpret the output of these models, in order to enhance trust that clinicians can have in these systems. So far, only a limited number of studies focused on interpreting these model outputs. In [17], a model was trained to predict dysarthric speech severity, by adding a bottleneck layer in a feed-forward fully-connected neural network to improve interpretability. Indeed, they relied on transfer learning to learn clinically interpretable labels taken from the Darley classification [18]. Their results show an accuracy improvement in dysarthria speech evaluation, along with justifications based on interpretable characteristics. An extension of this work in [19] avoided perceptual labels (which require costly annotation by experts, and is expensive in terms of recording amount required to train and evaluate models) by shaping the interpretable layer around four automatically-extracted acoustic features characterizing four dysarthria aspects, very close to the previous work. By using a SHapley Additive exPlanations [20], they analysed the contribution of each acoustic feature on the final prediction. Yeo et al. have also shown the efficacy of leveraging the *Goodness of Pronunciation* metric, by using the pre-trained multi-layer convolutional neural network part – frozen – of the Wav2Vec2 model, followed by a phone classifier [21]. Their approach shows the relative impact of each phone on the predicted severity score. Another methodology was used in [8] by using a convolutional neural network (CNN) for a phone classification task. The NCD method (*Neuro-Concept Detector*) along with the ANPS metric (*Artificial Neuron-based Phonological Similarity*) proposed in this work enables predicted severity and intelligibility scores to be interpreted in terms of alterations in the speech produced by patients, by associating the behaviour of neurons from the classifier with the phonetic feature they detect.

The work presented in this paper aims to revisit the overall framework proposed in [8] by replacing the CNN with an SSL model based on the complete Wav2Vec2 architecture. This modification is anticipated to provide not only higher accuracy, but also to bring later on a more nuanced and richer description of the phonetic features of pathological speech. In addition to the change of model architecture, this work will explore

how pre-training corpora, model size, the fine-tuning stage, and datasets associated with this stage affect the choice of a publicly available pre-trained Wav2Vec2 model for phone classification. This exploration is pivotal for optimizing the performance of the model across diverse datasets, contributing to the robustness and generalizability of our approach. Thus, a detailed analysis of phone confusions is conducted, to ensure the generalization ability of our models on other datasets. Subsequently, we analyse how phone classification accuracies correlate with perceptual measures obtained from domain experts on speech from patients who have benefited from care following a HNC. This multidimensional evaluation approach will provide a comprehensive assessment of the proposed SSL model, shedding light on its potential for an accurate and clinically relevant pathological speech analysis, required for Step 2 and 3 of the targeted overall framework [8].

## 2. Corpora

To train and evaluate our models during training, we relied on subsets of BREF [22] and Common Phone [23]. To test our models, we relied on BREF-Int and C2SI [1].

The **BREF** corpus is composed of 120 French speakers reading extracts from the newspaper *Le Monde*. Recordings took place in the 90s, with people recruited in the Paris area. A phone-balanced subset of this corpus was created to ensure that the subsequently fine-tuned models are not showing bias towards a specific phone. It is composed of more than 3 million frames of 127 ms, each aligned on a French phone or silence. We also use the **BREF-Int** dataset, a phone-balanced subset of the BREF test dataset. These datasets are identical to the ones used in [8].

The **Common Phone** corpus is a gender-balanced, multilingual and phonetically-aligned corpus derived from Mozilla’s Common Voice project. Only French recordings were used in this work, including speech from Belgium, Tunisian and Canadian speakers, amongst others. This dataset is also balanced in terms of phones and genders to ensure the phone-wise and gender-wise impartiality of the models.

The **C2SI** corpus consists of 87 patients who have been treated for oral or oropharyngeal cancer, as well as 41 healthy controls (HC). Patients and HC were recorded at the IUCT Oncopole, in Toulouse, France. Patients were faced with multiple tasks: sustained /a/, sentence reading, image description (DES), short text reading (LEC), pseudo-words production (DAP) as well as various other prosodic tasks. Based on the image description task, six experts evaluated the severity (degree of alteration of the speech signal) and intelligibility of the patients’ speech on a scale from 0 – strong alterations – to 10 – perfect speech; the intelligibility being defined here as the “the performance by a listener to recognize the words and/or the sounds of the speech produced by the speaker” [1]. Recordings from both tasks LEC and DAP are used for testing our models. They will be referred here as C2SI-LEC and C2SI-DAP. To match the selection that has been made in a previous work, we will test our models on 24 HC (amongst all available speakers, some of whom have not completed the two LEC and DAP tasks) all recorded under the same conditions. We will compare our results with perceptual evaluations of 81 patients – 6 patients were not evaluated by experts.

All the aforementioned corpora are phonetically aligned with 31 phones and silence. These 31 phones include four archi-phones: /Ê/ = {e, ε}, /Û/ = {œ, ø}, /Ô/ = {o, ɔ}, and /μ/ = {œ̃, ễ}. The use of these archi-phones neutralizes true-mid vowel

oppositions. Table 1 details the amount of phone-aligned frames used from each corpus, as well as the number of hours they represent. Because some frames overlap each other, this number of hours is higher than the sum of recordings lengths. However, it still reflects what the model sees as input.

Table 1: Usage, number of frames, and global duration of audio data for each corpus used in this work.

Corpus	Usage	#frames	#hours
BREF	train, validation	3,118k	110h
Common Phone	train, validation	236k	8.3h
BREF-Int	test	85k	3h
C2SI-LEC (HC)	test	43k	1.5h
C2SI-DAP (HC)	test	73k	2.5h

## 3. Model architectures

The CNN model has been previously trained for phone classification in [8] on the BREF dataset described in section 2. It is composed of two convolution layers combined with maximum pooling layers. The model input is a sliding window of 11 acoustic frames of 20 ms, 10 ms-spaced, where each frame is characterized by 40 Mel-Filter bank feature extracted from audio signal, along with their first and second derivatives. The model thus has a 120 ms-window centred on the phone to be predicted. Once the CNN has been applied, the output is then flattened before being fed to three fully connected layers, detailed below. This model output is considered as the *baseline* later in the paper.

Regarding the Wav2Vec2 architecture, models from LeBenchmark2.0 [24] are used. As we will apply these models on French pathological speech in further works, we targeted Wav2Vec2 models pre-trained exclusively on French speech. Indeed, combining phones from multiple languages could complicate the analysis between healthy and pathological speech (meaning between typical and atypical French phones). LeBenchmark models come in different architectures (6, 12, 24, or 48 hidden layers) as well as with different pre-training corpora sizes. In this work, we will compare results obtained with models containing 6, 12, or 24 hidden layers, respectively referred to as *light*, *base*, and *large*, as well as being pre-trained on 3k or 14k hours of French speech. Wav2Vec2 works with 25 ms-windows, which approximately start every 20 ms, which implies an overlap of about 5 ms. In order to emulate the CNN 120 ms-window, the input of Wav2Vec2 will consist here of audio extracts of six padded-windows of 25 ms. Such an architecture gives us an approximate 127 ms-window, which is very close to the length of the context used for the CNN.

The output of either the CNN or the Wav2Vec2 model, once flattened, are passed through three 1024-dimension fully connected layers, dedicated to the phone classification task. The size of the flattening layer depends on the output size of the encoder used (CNN or Wav2Vec2, and the size of Wav2Vec2). Then, the output phone is chosen using a softmax on the 32 output values.

Datasets used for the training phase were randomly split into two phone-balanced subset: *train* (90% of data) and *validation* (10% left). All of our models were fine-tuned using the SpeechBrain toolkit for 15 epochs, with a training time of approximately one day on a Tesla P100 GPU. The model presenting the best phone error rate on the validation set was chosen for inference on test datasets. The classifier uses an Adadelta

Table 2: *Phone-balanced accuracies (in %) obtained on each test dataset. Confidence intervals were computed with the bootstrapping approach<sup>2</sup>. Results for the CNN are taken from [8] and did not include confidence intervals. Best results are highlighted in bold. \*The CNN model was trained from scratch using BREF.*

Model	Fine-tuning datasets	BREF-Int $\uparrow$	C2SI-LEC $\uparrow$ (HC speakers)	C2SI-DAP $\uparrow$ (HC speakers)
CNN, <i>Baseline</i>	BREF*	81.4	72.2	69.2
14k-large Frozen	BREF	83.5 $\pm$ 0.2	66.9 $\pm$ 0.6	66.9 $\pm$ 0.4
14k-large	BREF	87.6 $\pm$ 0.2	70.2 $\pm$ 0.6	70.6 $\pm$ 0.4
14k-light	BREF	81.8 $\pm$ 0.2	57.0 $\pm$ 0.6	57.3 $\pm$ 0.4
3k-large	BREF	<b>88.3<math>\pm</math>0.2</b>	70.6 $\pm$ 0.6	71.3 $\pm$ 0.4
3k-base	BREF	84.5 $\pm$ 0.2	48.1 $\pm$ 0.6	50.1 $\pm$ 0.4
14k-large	BREF, CP	87.4 $\pm$ 0.2	72.1 $\pm$ 0.5	73.3 $\pm$ 0.4
14k-light	BREF, CP	82.9 $\pm$ 0.3	64.1 $\pm$ 0.6	63.7 $\pm$ 0.4
3k-large	BREF, CP	<b>88.3<math>\pm</math>0.2</b>	<b>72.6<math>\pm</math>0.6</b>	<b>73.9<math>\pm</math>0.4</b>
3k-base	BREF, CP	84.9 $\pm$ 0.2	61.4 $\pm$ 0.6	62.5 $\pm$ 0.4

optimizer with an initial learning rate of 0.9, to improve a cross-entropy loss applied on the phone classification. The Wav2Vec2 architecture – when fine-tuned – relies on an Adam optimizer with an initial learning rate of  $1.10^{-4}$ . Source code containing SpeechBrain recipes is made available on a Github repository<sup>1</sup>.

## 4. Experimental results

### 4.1. Comparison of Wav2Vec2 models

Experiments were run to investigate the impact of the following factors: (1) fine-tuning Wav2Vec2, (2) pre-training datasets, (3) model size, and (4) fine-tuning datasets. Table 2 sums up accuracies obtained for each fine-tuned model. Given that the phone distribution is not balanced in C2SI datasets, accuracies are balanced by phone. Mathematically, it is the average of accuracies obtained for each phone. Thus, even when not specified, accuracies are systematically balanced by phone.

To analyse the impact of fine-tuning Wav2Vec2 (1), two *14k-large* models were used, only one of which has been fine-tuned. By comparing both models *14k-large Frozen* and *14k-large*, we can see that the fine-tuning step significantly improves accuracies on all three test datasets (confidence intervals are not overlapping themselves). Unsurprisingly, fine-tuning Wav2Vec2 models on similar datasets is beneficial regarding accuracy.

Then, the impact of the pre-training datasets (2) was measured by using LeBenchmark models that had different sets of pre-training data. We relied on both 3k and 14k model families, which are pre-trained on respectively 3,000 and 14,000 hours of French speech. They include read, acted, spontaneous, and professional French speech, with standard and accentuated French. Both of these datasets vary greatly in terms of accent variety. Indeed, the 3k pre-training dataset mainly contains audiobooks and French radio programmes, with less than 1% of African-accented speech that we are aware of – included audiobooks and radio programmes do not provide information on the accents of their speakers. However, the 14k pre-training datasets contain around 4,700 hours of speech from the European Parliament, which features at least three additional French accents: Belgian, Swiss, and Aosta Italian, as well as more negligible amounts

of African radio, comprising Malian and Nigerian accents. By comparing models *14k-large* and *3k-large*, our results underline that adding linguistic variability through various French accents does not offer significant improvements regarding classification accuracies.

Regarding model size (3), multiple LeBenchmark model sizes (3) were experimented: *light*, *base* and *large*, involving 6, 12, and 24 hidden layers respectively. Results show that using a *large* model (with 24 hidden layers) offers significantly better accuracies than smaller models (with 6 and 12 hidden layers). Unfortunately, as LeBenchmark offers neither a *14k-base* model nor a *3k-light* model, we cannot compare the results of the *3k-base* and *14k-light* models. Nevertheless, our results show that smaller models have weaker generalization power than larger models on both unseen C2SI datasets.

Finally, to analyse the impact of fine-tuning datasets on transformer-based models (4), we add a new fine-tuning dataset – the Common Phone (CP) corpus – compared with [8], in which only the BREF dataset was involved. All the initial models mentioned above have separately been fine-tuned on a combination of the BREF and Common Phone corpora (*BREF,CP* in Table 2). According to our results, adding read speech from other corpora in the fine-tuning stage significantly improves the generalization to both C2SI datasets on all our fine-tuned models (for instance, an absolute 12.4% improvement is achieved by the *3k-base* model trained on BREF and CP, compared with the *3k-base* model trained on BREF only, on the C2SI-DAP dataset). Accuracies are also similar on the BREF-Int dataset: all confidence intervals overlap each other, except for the *14k-light* model, which is significantly better, once fine-tuned on the combination of BREF and Common Phone. It therefore seems worthwhile to include other datasets in the fine-tuning stage, even if they do not represent a significant percentage of the training data.

### 4.2. Wav2Vec2 and CNN

The best Wav2Vec2 model – *3k-large* fine-tuned on BREF and CP, still according to results shown in table 2, improves by an absolute 6.9% the classification accuracy on BREF-Int compared with the CNN (88.3% versus 81.4%). On C2SI-LEC, the difference is not significant (72.6% versus 72.2%). Comparatively, we note a significant improvement of 4.7% on C2SI-DAP (73.9% versus 69.2%). However, by comparing the number of parameters of these models, we can point out that the best Wav2Vec2 model has 330 million of parameters, against 10 million for the CNN. Furthermore, smaller Wav2Vec2 models *base*

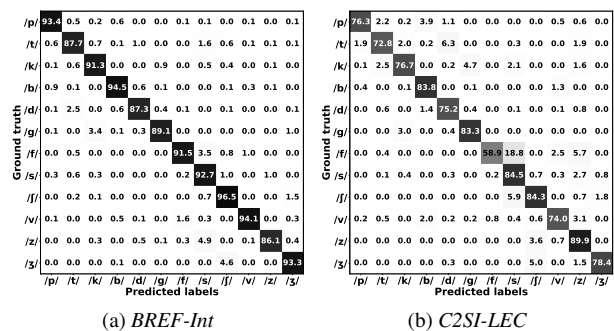


Figure 1: *Confusion matrices of obstruent phones, on BREF-Int and C2SI-LEC (HC speakers only) datasets respectively.*

<sup>1</sup>github.com/MaloMn/wav2vec2-phone-classification

<sup>2</sup>Ferrer, L. and Riera, P. Confidence Intervals for evaluation in machine learning [Computer software]. <https://github.com/luferrier/ConfidenceIntervals>

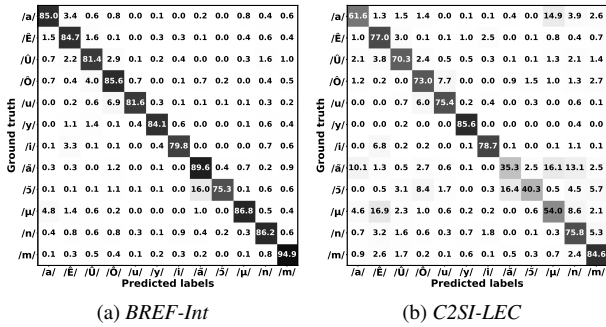


Figure 2: Confusion matrices of oral/nasal phones, on BREF-Int and C2SI-LEC (HC speakers only) datasets respectively.

and *light* – with respectively 90 and 26 million parameters – do not generalize as well as the CNN on C2SI datasets. This gap in parameter count is significant, but both model types do not have the same architecture, and a comparison based only on parameter count is necessarily biased. Nevertheless, this size difference implies that the carbon footprint of our inferences will inevitably be higher when using Wav2Vec2 models.

To ensure that our models do not overfit on certain phones, and that confusions between phones remain explainable, confusion matrices were generated on BREF-Int and C2SI-LEC, using the *3k-large* model fine-tuned on BREF and Common Phone. Figures 1a, 1b, 2a and 2b show specific parts of these matrices, dedicated to obstruent and oral/nasal phones. Restraining our analysis to these phones makes a direct comparison with the analysis done in [8] possible. The comparison between previously obtained results and current results show that Wav2Vec reduces the observed confusions in most cases.

Regarding obstruent phones (figures 1a and 1b), /ʃ/ was confused with /ʒ/ in 9% of cases on BREF-Int in the previous work, against 4.6% now. Confusion between both of these phones is therefore still present, but happens in half the cases compared with previous results. This reduction can also be found on the C2SI dataset: whereas /p/ was previously confused with /t/ in 9.3% of cases, this is only the case 2.2% of the time with our architecture. We also find the same causes for important confusions: either the loss of the place of articulation (acuteness – /f/ → /s/, compactness – /j/ → /s/), or confusion due to the voicing feature (/t/ → /d/).

Regarding oral and nasal phones, we also get strong confusions linked to oral vowels and nasal consonants on C2SI-LEC. Where /ā/ and /a/ were confused in 11.5% of cases, they still are in 10.1% of cases. /ā/ is also more confused now with /n/ (13.1%, against 8.4% previously) but is less confused with /m/ (2.5%, against 6.2% previously). These confusions, previously explained by the change of speakers’ accent (Parisians for BREF-Int, and from the Toulouse area for C2SI), can also be found here, when using another type of model architecture. Our results therefore support those obtained previously, and show that the data used during fine-tuning and the domain differences between datasets remain a non-negligible problem for this type of architecture. On BREF-Int, we also find strong confusions between /ā/ and /ɛ̃/, with a mutual confusion of 16.4%, albeit less intense than when using a CNN – 20.6%. These results are important because they show that our model is also sensitive to atypical pronunciations (such as a regional accent), which is desirable when analysing pathological speech.




Figure 3: Scatter plots of C2SI patients and HC on perceptual evaluations by experts and balanced accuracies obtained.

### 4.3. Application to pathological speech

Having demonstrated the robustness of our model and its ability to generalize well to other datasets, we will now see whether its computed accuracies can correlate positively with the assessments of the six C2SI experts. The severity and intelligibility scores given by the experts on the image description task (DES) for a speech record were averaged and then compared to classification accuracies. Figures 3a and 3b show scatter plots comparing perceptual scores to classification accuracies of phones in recordings of patients and HC from the C2SI corpus. Regression lines are also plotted. High Pearson correlation values – 0.90 for severity and 0.80 for intelligibility – confirm that perceptual measures can be estimated using phone-balanced classification accuracies with our *3k-large* model fine-tuned on BREF and Common Phone. These correlation values are similar to the ones obtained using the CNN [8], where balanced accuracies correlated at 0.91 and 0.81 with severity and intelligibility scores respectively.

The less-intense correlation with perceptual intelligibility can be explained by the fact that experts are known to be biased on evaluating intelligibility, by underestimating the degree of speech disorders. Indeed, when a task involves using expected French words (for instance, using the word “bateau” to describe an image showing a boat), experts tend to understand them better. This was one of the reasons for creating a pseudo-word reading task in [1, 25]. Furthermore, experts have generally less trouble understanding patients. These results confirm that a speech representation based on a Wav2Vec2-based model is well suited for a phone-wise analysis of pathological speech.

## 5. Conclusion

In this work, we showed that a Wav2Vec2-based model outperforms a CNN model on a phone classification task, while preserving certain linguistic specificities, such as regional accents. Our results validate not only the efficacy of our approach based on Wav2Vec2, but also highlights the importance of the choice of the model architecture as well as the diversity of fine-tuning datasets for optimal performance. Future works include applying the NCD concept developed in [8] to analyse hidden layers of our fine-tuned architecture, and analyse how this new model influences the detection of phonetic features, which are crucial for interpreting predicted phones. In turn, this interpretability is needed to have an objective analysis of a patient’s speech, which would improve re-habilitation techniques implemented by clinicians.

## 6. References

- [1] C. Astésano, M. Balaguer, J. Farinas, C. Fredouille, P. Gaillard, A. Ghio, I. Laaridh, M. Lalain, B. Lepage, J. Mauclair, O. Noicaudie, J. Pinquier, O. Pont, G. Pouchoulin, M. Puech, D. Robert, E. Sicard, and V. Woisard, "Carcinologic speech severity index project: A database of speech disorder productions to assess quality of life related to speech after cancer," in *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*. Miyazaki, Japan: European Language Resources Association (ELRA), May 2018.
- [2] A. Maier, H. Tino, F. Stelzle, E. Noeth, E. Nkenke, R. Frank, S. Anne, and M. Schuster, "Automatic speech recognition systems for the evaluation of voice and speech disorders in head and neck cancer," *EURASIP Journal on Audio, Speech, and Music Processing*, vol. 2010, 01 2010.
- [3] C. Middag, R. Clapham, R. van Son, and J.-P. Martens, "Robust automatic intelligibility assessment techniques evaluated on speakers treated for head and neck cancer," *Computer Speech & Language*, vol. 28, no. 2, pp. 467–482, 2014.
- [4] I. Laaridh, C. Fredouille, A. Ghio, M. Lalain, and V. Woisard, "Automatic Evaluation of Speech Intelligibility Based on I-vectors in the Context of Head and Neck Cancers," in *Proc. Interspeech 2018*, 2018, pp. 2943–2947.
- [5] R. Vaysse, J. Farinas, C. Astésano, and R. Andre-Obrecht, "Automatic extraction of speech rhythm descriptors for speech intelligibility assessment in the context of head and neck cancers," 08 2021, pp. 1912–1916.
- [6] L. Bin, M. C. Kelley, D. Aalto, and B. V. Tucker, "Automatic speech intelligibility scoring of head and neck cancer patients with deep neural networks," in *International Congress of Phonetic Sciences (ICPhS' 19)*, Melbourne, Australia, 2019, pp. 3016–3020.
- [7] S. Quintas, A. Abad, J. Mauclair, V. Woisard, and J. Pinquier, "Towards reducing patient effort for the automatic prediction of speech intelligibility in head and neck cancers," in *ICASSP 2023 - 2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2023, pp. 1–5.
- [8] 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," *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, vol. 31, pp. 200–214, 2023.
- [9] T. tom Dieck, P. A. Pérez-Toro, T. Arias, E. Noeth, and P. Klumpp, "Wav2vec behind the Scenes: How end2end Models learn Phonetics," in *Proc. Interspeech 2022*, 2022, pp. 5130–5134.
- [10] A. Baevski, Y. Zhou, A. Mohamed, and M. Auli, "wav2vec 2.0: A framework for self-supervised learning of speech representations," in *Advances in Neural Information Processing Systems*, vol. 33. Curran Associates, Inc., 2020, pp. 12 449–12 460.
- [11] A. Pasad, J.-C. Chou, and K. Livescu, "Layer-wise analysis of a self-supervised speech representation model," in *2021 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*, 2021, pp. 914–921.
- [12] A. Pasad, B. Shi, and K. Livescu, "Comparative layer-wise analysis of self-supervised speech models," in *ICASSP 2023 - 2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2023, pp. 1–5.
- [13] A. Hernandez, P. A. Pérez-Toro, E. Noeth, J. R. Orozco-Arroyave, A. Maier, and S. H. Yang, "Cross-lingual Self-Supervised Speech Representations for Improved Dysarthric Speech Recognition," in *Proc. Interspeech 2022*, 2022, pp. 51–55.
- [14] E. J. Yeo, K. Choi, S. Kim, and M. Chung, "Automatic severity classification of dysarthric speech by using self-supervised model with multi-task learning," in *ICASSP 2023 - 2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2023, pp. 1–5.
- [15] F. Javanmardi, S. R. Kadiri, and P. Alku, "Pre-trained models for detection and severity level classification of dysarthria from speech," *Speech Communication*, p. 103047, 2024.
- [16] A. Favaro, Y.-T. Tsai, A. Butala, T. Thebaud, J. Villalba, N. Dehak, and L. Moro-Velázquez, "Interpretable speech features vs. dnn embeddings: What to use in the automatic assessment of parkinson's disease in multi-lingual scenarios," *Computers in Biology and Medicine*, vol. 166, p. 107559, 2023.
- [17] M. Tu, V. Berisha, and J. Liss, "Interpretable Objective Assessment of Dysarthric Speech Based on Deep Neural Networks," in *Proc. Interspeech 2017*, 2017, pp. 1849–1853.
- [18] F. L. Darley, A. E. Aronson, and J. R. Brown, "Clusters of deviant speech dimensions in the dysarthrias," *Journal of Speech and Hearing Research*, vol. 12, no. 3, pp. 462–496, 1969.
- [19] L. Xu, J. Liss, and V. Berisha, "Dysarthria detection based on a deep learning model with a clinically-interpretable layer," *JASA Express Letters*, vol. 3, no. 1, p. 015201, 01 2023.
- [20] S. M. Lundberg and S.-I. Lee, "A unified approach to interpreting model predictions," in *Advances in Neural Information Processing Systems*, vol. 30. Curran Associates, Inc., 2017.
- [21] E. J. Yeo, K. Choi, S. Kim, and M. Chung, "Speech intelligibility assessment of dysarthric speech by using goodness of pronunciation with uncertainty quantification," in *Proc. INTERSPEECH 2023*, 2023, pp. 166–170.
- [22] L. F. Lamel, J.-L. Gauvain, M. Eskénazi *et al.*, "Bref, a large vocabulary spoken corpus for french," *Eurospeech'91, Italy*, vol. 22, no. 28, p. 50, 1991.
- [23] P. Klumpp, T. Arias, P. A. Pérez-Toro, E. Noeth, and J. Orozco-Arroyave, "Common phone: A multilingual dataset for robust acoustic modelling," in *Proceedings of the Thirteenth Language Resources and Evaluation Conference*. Marseille, France: European Language Resources Association, Jun. 2022, pp. 763–768.
- [24] T. Parcollet, H. Nguyen, S. Evain, M. Zanon Boito, A. Pupier, S. Mdhaffar, H. Le, S. Alisamir, N. Tomashenko, M. Dinarelli, S. Zhang, A. Allauzen, M. Coavoux, Y. Estève, M. Rouvier, J. Goulian, B. Lecouteux, F. Portet, S. Rossato, F. Ringeval, D. Schwab, and L. Besacier, "Lebenchmark 2.0: A standardized, replicable and enhanced framework for self-supervised representations of french speech," *Computer Speech & Language*, vol. 86, p. 101622, 2024.
- [25] M. Lalain, A. Ghio, L. Giusti, D. Robert, C. Fredouille, and V. Woisard, "Design and development of a speech intelligibility test based on pseudowords in french: Why and how?" *Journal of Speech, Language, and Hearing Research*, vol. 63, no. 7, pp. 2070–2083, 2020.