



# Nonwords Pronunciation Classification in Language Development Tests for Preschool Children

*Ilja Baumann<sup>1</sup>, Dominik Wagner<sup>1</sup>, Sebastian Bayerl<sup>1</sup>, Tobias Bocklet<sup>1,2</sup>*

<sup>1</sup>Technische Hochschule Nürnberg Georg Simon Ohm, Germany

<sup>2</sup>Intel Labs

{ilja.baumann, dominik.wagner, sebastian.bayerl, tobias.bocklet}@th-nuernberg.de

## Abstract

This work aims to automatically evaluate whether the language development of children is age-appropriate. Validated speech and language tests are used for this purpose to test the auditory memory. In this work, the task is to determine whether spoken nonwords have been uttered correctly. We compare different approaches that are motivated to model specific language structures: Low-level features (FFT), speaker embeddings (ECAPA-TDNN), grapheme-motivated embeddings (wav2vec 2.0), and phonetic embeddings in form of senones (ASR acoustic model). Each of the approaches provides input for VGG-like 5-layer CNN classifiers. We also examine the adaptation per nonword. The evaluation of the proposed systems was performed using recordings from different kindergartens of spoken nonwords. ECAPA-TDNN and low-level FFT features do not explicitly model phonetic information; wav2vec2.0 is trained on grapheme labels, our ASR acoustic model features contain (sub-)phonetic information. We found that the more granular the phonetic modeling is, the higher are the achieved recognition rates. The best system trained on ASR acoustic model features with VTLN achieved an accuracy of 89.4% and an area under the ROC (Receiver Operating Characteristic) curve (AUC) of 0.923. This corresponds to an improvement in accuracy of 20.2% and AUC of 0.309 relative compared to the FFT-baseline. **Index Terms:** children's speech, speech assessment, transfer learning

## 1. Introduction

Language competence is one of the key skills for later academic success. A need for support should be identified as soon as possible [1]. For this reason, language assessment as part of a school enrollment examination is mandatory in Germany. Professional speech therapists perform the assessment. One of the aspects to be tested is auditory memory, with phonological representability and short-term memory span being of particular importance. In German-speaking countries, among others, the SETK (Sprachentwicklungstest für Kinder – Language Development Test for Children) [2] test is used to diagnose language processing abilities and auditory memory capacity. The test is designed for children in the critical age range 3;0 – 5;11 years or older children with known developmental disabilities. The SETK is divided into subtests, each of which can be assigned to language comprehension, language production or language memory. In order to determine the level of language attained and to relate it causally to auditory memory performance, phonological working memory for nonwords or made-up words is a crucial factor [3]. The fact that children with speech disorders repeat nonwords less accurately than children of the same age with normal speech development was proven by various studies [4, 5]. The phonological working memory subtest for nonwords as part of the SETK test is evaluated in this paper. The aim is to classify children's recited nonwords to

determine whether the pronunciation was correct or incorrect. Different systems for automatic classification are compared and evaluated to support speech therapists during a test in a screening setup. Beginning with modeling of raw features, utterance embeddings, grapheme-motivated embeddings, and finally phonetic embeddings.

Convolutional Neural Networks (CNNs) proved to be very effective in modeling aspects of speech and have been successfully applied to downstream tasks, such as speech emotion recognition [6]. In [7], architectures such as AlexNet, VGG, Inception, and ResNet from the image domain were primarily used to classify audio. Results indicate that neural embeddings extracted from these models perform better than raw feature inputs. In [8], VGGish, a pre-trained deep CNN was used as a feature extractor to train a classifier based on the AudioSet dataset [9]. The extracted features are vastly superior over MFCCs with 30 times fewer features to represent the data. Phoneme pronunciation classification was investigated in [10], using CNN architectures and fine-tuning approaches. In [11], voice pathology detection is performed using CNN models. X-Vectors [12] or ECAPA-TDNN as enhanced architecture, were originally introduced for speaker verification. The architectures have been successfully used in the areas of speaker diarization [13], dialect classification [14], pathologic speech assessment [15, 16] and other paralinguistic tasks, e.g., emotion recognition [17]. Wav2vec 2.0 [18] models have been successfully applied to downstream speech processing tasks such as emotion recognition [19] or detection of Alzheimer's speech [20].

We use a VGG-like CNN system as a backend to classify the pronunciation of spoken nonwords. We also investigate fine-tuning for word-optimized classification. As input, we use mel spectrograms and outputs from pretrained state-of-the-art architectures in related domains. These are motivated to extract and model different levels of speech structure. We in particular use ECAPA-TDNN [21], wav2vec 2.0 [18] and a Kaldi-based [22] TDNN acoustic model for ASR. Whole utterance embeddings are extracted with a pre-trained ECAPA-TDNN for speaker verification. The embeddings represent a holistic expression without explicit modeling of phonetic information. Grapheme level embeddings are extracted via wav2vec 2.0; embeddings after each transformer block are compared. Finally, we use a phonetically motivated embedding generator in form of an ASR acoustic model, i.e., the senone outputs from a Kaldi-based time-delay neural network (TDNN) acoustic model. The embeddings of this modeling contain (sub-)phonetic information.

## 2. Data

The SETK subtest for phonological working memory used in this work consists of nonwords that have to be repeated by the children. Performance is evaluated based on the number of correct and incorrect repetitions. The ability to capture unfamiliar phonetic patterns, store them short-term in working

memory, and recall them immediately is tested. Phonological awareness is the underlying process that describes children’s ability to identify sublexical units based on holistic representations. In addition, the repetition of nonwords tests the phonological short-term memory since phonological patterns which are not known have no entry in the lexical memory [23]. Seven nonwords must be repeated during the test, which differ in complexity, length, and number of syllables. The test begins with 2-syllable nonwords and gradually progresses to 5 syllables. The nonwords are: *Maluk* [ma:l-ɔk], *Bilop* [bi:l-ɔp], *Ronterklabe* [ʁɔn-tɐ-kla:bə], *Glösterkeit* [glø:s-tɐ-kait], *Seregropist* [ze:ʁe-ɡʁo:pist], *Pristobierichkeit* [pʁis-to:bi:ʁiç-kait] and *Kabusaniker* [ka-bu:za-ni:kɛ].

Audio recordings were collected by the authors of [24] in various German kindergartens. The recordings were obtained by experienced speech therapists in the respective kindergartens on site. Each child was tested and recorded individually. For each nonword, a recording of the spoken nonword by a speech therapist was played to provide the same conditions for each recording. The final dataset consists of 984 utterances or nonwords from 140 children. The data was labeled by two expert raters. Each nonword was rated as correct or incorrect, with an additional confidence rating, confident or uncertain, since some statements were very unclear. Only utterances for which both raters were confident that they labeled correctly and agreed on the rating were selected, resulting in 651 utterances from 139 children. We divide the data into a training and test set (random 75% / 25% split). 25% of the test set is taken for validation. We removed silence, other voices, and noise from the recordings using rVAD [25] before and after each nonword is repeated.

### 3. Modeling Techniques

#### 3.1. Classification system

Our baseline is a CNN system trained on spectrograms, inspired by the VGG [26] network used for image classification. The network consists of a series of convolution and activation layers, followed by batch normalization and a max-pooling layer. Similar to VGG, the convolutional layers have small filters with a kernel size of 3×3 and a stride of 1. ReLU was used as activation after convolution. Our best-performing base model had 5 layers and 4 filter kernels in the first layer; filter kernels are doubled layer by layer. The left-hand side of Figure 1 shows the CNN architecture. Convolutional layers are followed by fully connected layers containing the number of channels corresponding to the number of filter kernels of the last convolutional layer.

The proposed network structure was used to train with different features in each case, extracted with ECAPA-TDNN, wav2vec 2.0, and Kaldi TDNN acoustic models. Figure 1 illustrates, beside the CNN classifier architecture, the proposed classifiers trained on different features, modeling specific aspects of speech. Only the input layer was modified for the corresponding feature dimensions.

In addition, different fine-tuning methods were explored in the baseline setup. The fine-tuning was performed on the respective word-independent models trained on all words. We trained word-dependent models in 3 different ways:

- 1) The last linear layer was retrained for each nonword, the remaining layers were frozen.
  - 2) The whole pre-trained model was fine-tuned with a lower learning rate of  $10^{-5}$ .
  - 3) An additional linear layer with 16 nodes was appended to the model, all other layers of the pre-trained model were frozen.
- For comparison, we trained a word-dependent model for each nonword from scratch, i.e., without pre-training.




Figure 1: The proposed systems use identical CNN architectures to classify the nonwords. Features are extracted more and more granular according to the speech structure.

#### 3.2. Features

##### 3.2.1. Low-level features

We first used low-level features for the training of the baseline system. We converted the input audio waveform of the whole utterance into a sequence of 128-dimensional mel spectrograms. The mel spectrograms were obtained through a short-time Fourier transform (STFT) using a 25ms frame size, 10ms frame hop, and a Hann window function. The resulting sequences were zero-padded to the same length according to the longest element. By using the spectrograms, we attempt to model each nonword as a holistic representation that does not contain any underlying information of the speech.

##### 3.2.2. Utterance embeddings

ECAPA-TDNN [21] is an architecture based on the x-vectors [12] topology, a Time Delay Neural Network (TDNN) that applies statistics pooling to project variable-length utterances into fixed-length speaker embeddings. Several enhancements were introduced to create more robust speaker embeddings. A channel- and context-dependent attention mechanism performs pooling, allowing the network to attend to different frames per channel. In addition, global context information is added to local operating convolutional blocks by using 1-dimensional squeeze-excitation (SE) blocks that rescale the channels of intermediate frame-level feature maps. 1-dimensional Res2Net blocks improve performance and reduce the total number of parameters. Features are concatenated with intermediate outputs from previous layers before the attentive statistics pooling layer.

The ECAPA-TDNN [27] was trained on the development part of the VoxCeleb2 dataset with 5994 speakers. 80-dimensional MFCCs have been generated as input features. Different augmentations were applied to the training data, including time-domain SpecAugment, reverb, and noise. Using ECAPA-TDNN we extracted 192-dimensional embeddings over the entire utterances. We trained a simple classifier with two linear layers and a 256-dimensionality of output space. The whole pre-trained classifier was fine-tuned for each nonword. Similar to spectrograms, we use these embeddings to model a holistic representation of the words, as ECAPA-TDNN was trained on speaker labels without phonetic information.

##### 3.2.3. Grapheme embeddings

Wav2vec 2.0 uses self-supervised learning and subsequent fine-tuning to make speech recognition more robust using less labeled data. The raw waveform is normalized and fed into a feature encoder, consisting of several blocks of temporal convolutions, layer normalization, and gaussian error linear unit (GELU) activation function. The output of the feature encoder is fed into a transformer network. A convolutional layer is used to derive the relative position instead of positional embeddings.

Using product quantization, the output of the feature encoder is also discretized to a finite set of speech representations. Since a fraction of time steps in the latent feature encoder space is masked, a contrastive task is given. The model has to identify the true quantized representation for a masked time step. Fine-tuning of pre-trained models for speech recognition is achieved by adding a randomly initialized linear projection on top of the context network.

We used a wav2vec 2.0 model pre-trained on unlabeled audio from the LibriSpeech dataset [28] and fine-tuned for ASR on the same audio with the corresponding transcripts. We extracted 768-dimensional features for every 20ms of audio. Wav2vec 2.0 is trained on unlabeled data and fine-tuned with an additional layer for ASR to predict characters. This results in embeddings of nonwords at the grapheme-motivated level of speech.

### 3.2.4. Phonetic embeddings

We finally use the ASR recipe Kaldi Tuda-De provided in [29] available for German to extract phonetic embeddings. The model is trained on freely available data resources, including the Spoken Wikipedia Corpora (SWC), the German subset of m-aillabs, the German Commonvoice Corpus, and Tuda-De. The final result is about 1050 hours of training data. We trained a hybrid DNN-HMM model for ASR and used a TDNN as an acoustic model following the chain recipe in [29]. Since our goal is to classify the nonwords, we do not use the whole ASR system but only the outputs of the acoustic model to train a classifier based on these. The input to the TDNN acoustic model is composed of two types of features: 40-dimensional MFCCs extracted from a 25ms window and 10ms frameshift and 100-dimensional i-vectors from chunks of 150 consecutive frames. The final input is a 220-dimensional concatenated vector consisting of three consecutive MFCC vectors and i-vector. For decorrelation, LDA is applied on the features without changing the dimensionality. The TDNN stack consists of 6 blocks that operate on the 2048 dimensional output of the previous block. Each TDNN block is a succession of 1-dimensional temporal convolutions except in the first block and a consecutive process of ReLU activation and batch normalization. The contexts differ from block to block and use sub-sampling in the last 3 blocks, skipping temporal frames to access a broader context. The network is trained with two outputs for better performance, a chain-based LF-MMI [30] criterion and cross-entropy criterion. However, only the chain-based output without softmax is used for inference, since the model is trained with sequence objective. The output of the acoustic model is 3510-dimensional and consists of posterior probabilities of the acoustic states. Based on the extracted senones we trained a word-independent classifier and fine-tuned it for word-dependent models.

Due to the shorter vocal tract and smaller vocal folds, children’s utterances have a higher fundamental frequency than adults [31]. To normalize the spectral distribution of different speakers, vocal tract length normalization (VTLN) was used as a feature-level transformation. In this process, the positions of the center frequencies of the triangular frequency bins are moved. We applied VTLN to match the fundamental frequency of the input to the training data as closely as possible. Using FCN-f0 [32] we calculated the average fundamental frequency on the original training data of the Kaldi ASR model. Afterward, we calculated a warp factor per speaker to obtain more robust scores based on the short utterances. We finally applied the warp factor to each speaker’s utterances before extracting the senones. The resulting features represent the most granular unit of speech with the phonetic embeddings in this work.

## 4. Experiments and Results

The classifiers were trained with TensorFlow [33] using the Adam [34] optimizer. Batch normalization [35] was applied after each convolutional layer. In the last linear layers, dropout was used with a probability of 0.5. A sigmoid layer was used as the final layer in all models, binary cross-entropy as the loss function. Class weights were applied according to the respective distribution of labels in training to compensate data imbalance, especially in word-specific models. The test set was constant across all experiments. We report accuracy, precision, recall, and AUC, the area under the Receiver Operating Characteristic (ROC) curve [36] across the 7 nonwords.

We first trained word-independent classifiers on the different features and fine-tuned these further. The SETK test is established as a screening test, and one nonword is tested after the other by a speech therapist. Due to this, it is always known which nonword is being tested, and therefore word-dependent classifiers can be used. We investigated different methods to implement word-dependent systems in 3.1. Thus, for each type of feature, one word-independent system and 7 word-dependent systems were created.

### 4.1. Results

#### 4.1.1. Low-level features

We trained the baseline classifier on spectrograms with a learning rate of  $10^{-3}$  for 100 epochs and early stopping with patience of 10 epochs. The entire training set was used as input, which provides the word-independent model. The results are listed in Table 1, including the proposed fine-tuning methods for word-dependent systems. Word-dependent models performed best, with fine-tuning of the word-independent model. Training from scratch for each nonword provided similar performance. Only the models fine-tuned on all layers provided an AUC over 0.7. Nevertheless, the precision of the best models remains low, with 0.693 across all nonwords. The best approach, fine-tuning all layers of a pre-trained word-independent model, was used for further experiments with different embedding types of speech.

Table 1: Baseline CNN systems trained on mel spectrograms.

System	Accuracy	Precision	Recall	AUC
Word-independent	0.741	0.627	0.722	0.640
Word-dependent				
- FT last layer	0.682	0.660	0.843	0.525
- FT all layers	0.744	0.693	0.737	0.705
- FT add. layer	0.704	0.611	0.586	0.554
From scratch	0.720	0.669	0.758	0.689

#### 4.1.2. Utterance embeddings

We used ECAPA-TDNN extracted embeddings to train a word-independent model and fine-tuned it. Results in Table 2 show that the fine-tuned models are similar to the baseline in accuracy, precision, and recall. However, low AUC indicates that due to the short utterances, not enough meaningful information could be extracted using the pre-trained ECAPA-TDNN. An examination of the individual nonword classes revealed that AUC was particularly low for short words. It is noticeable for *Bilop* with an AUC of 0.523, whereas *Pristobierlichkeit* and *Kabusaniker* result in 0.758 and 0.803.

Table 2: ECAPA-TDNN word-independent and fine-tuned word-independent systems.

System	Accuracy	Precision	Recall	AUC
Word-independent	0.610	0.630	0.632	0.505
Word-dependent	0.727	0.746	0.790	0.689




Figure 2: Wav2vec 2.0 transformer blocks 1-12 with the respective weighted averaged metrics for word-independent models.

#### 4.1.3. Grapheme embeddings

Using grapheme-motivated embeddings from wav2vec 2.0, we trained one word-independent model for each output of the 12 transformer blocks to find the most relevant for the classification of nonwords. Figure 2 shows the result for each transformer block output features for word-independent models. We got the best results from the last transformer block. The higher the processing hierarchy of the respective transformer blocks, the better were the results. This is consistent with the processing hierarchy, where the higher-level modeling is more relevant to phonetic structures. We used the wav2vec 2.0 outputs of the last transformer block to pre-train the word-independent model, on which the word-specific fine-tuning was applied. Table 3 shows the results of the wav2vec 2.0 feature models. By fine-tuning to word-dependent classifiers, a precision, recall, and AUC around 0.8 is achieved. The lowest AUC for the fine-tuned models is 0.683 for one of the longest nonwords *Seregropist*. The variability of the performance depending on the nonword class is smaller compared to the baseline system.

Table 3: Wav2vec 2.0 word-independent and fine-tuned word-independent systems.

System	Accuracy	Precision	Recall	AUC
Word-independent	0.766	0.607	0.633	0.707
Word-dependent	0.827	0.800	0.794	0.794

#### 4.1.4. Phonetic embeddings

We used the senones of the Kaldi-based acoustic model to train two word-independent models, with and without VTLN applied, and fine-tuned these. Results in Table 4 show that the word-independent model outperforms the other modeling techniques considered in this work. Fine-tuning leads to further improvements. Between the word-independent models with and without VTLN, the main improvement was in AUC from 0.783 to 0.848. Especially in fine-tuning, the use of VTLN showed a major impact. Also, more complex and long nonwords like *Pristobierichkeit* yield solid results in precision, recall, and AUC with 0.8, 1.0, 0.96. Compared to the word-dependent models with wav2vec 2.0 features, phonetic embeddings lead

to improvements in precision, recall, and AUC from 6-18% relative.

Table 4: Word-independent and fine-tuned word-independent systems based on TDNN ASR acoustic model features.

System	Accuracy	Precision	Recall	AUC
Word-independent	0.815	0.849	0.705	0.783
Word-dependent	0.842	0.863	0.835	0.848
VTLN				
Word-independent	0.814	0.854	0.727	0.848
Word-dependent	0.894	0.849	0.934	0.923

## 4.2. Discussion

Our findings show that the classification accuracy increases gradually with the phonetic modeling capabilities of the different systems. At the same time, the variability in the evaluation of individual nonwords decreases. The more detailed phonetic representations are modeled, the more meaningful the features and the more robust the recognizers are. For the task of evaluating specific nonwords, a precise modeling of phonetic information is more important than for speaker-holistic tasks, e.g., accent, emotion, or presence/absence of certain diseases. Spectrograms and ECAPA-TDNN embeddings, modeling the holistic statement, have provided the lowest results on our task with the highest performance variability across the different nonwords (AUC standard deviation (SD) of 0.218 and 0.132). The use of grapheme-motivated features (wav2vec 2.0) results in improvements of all metrics, especially AUC. The variability of individual nonwords performance is lower compared to ECAPA-TDNN embeddings and spectrograms (AUC SD of 0.114). This is because more detailed structures of speech are modeled over time. Senones as (sub-)phonetic level features from the ASR acoustic model performed best. Probably also because more significant information can be obtained from an even more granular modeling of speech compared to wav2vec 2.0 features. The metrics variability across nonwords is also the lowest (AUC SD of 0.050). The more detailed the phonetic information seems to be modeled, the better the results.

## 5. Conclusion

We explored different ways of extracting and modeling features for nonword classification. Low-level features, speaker embeddings, grapheme-motivated embeddings, and phonetically-motivated embeddings. We proposed a 5 layer VGG-like CNN as the classifier for the different feature levels of the extracted speech. We evaluated the models on nonwords spoken by children and showed that wav2vec 2.0 and especially TDNN acoustic models for ASR provide robust features for nonword classification. TDNN acoustic model with VTLN at feature-level outperformed other embeddings while using an identical CNN classifier system. The best model trained with phonetic embeddings achieved an accuracy of 89.4%, precision of 0.849, recall of 0.934, and AUC of 0.923. We have shown the progress of performance over the different levels of speech modeling and how it increases the more granular the speech structure is. To the best of our knowledge, this is the first comparison of modeling different detail levels of speech and using senones for classification.

## 6. Acknowledgements

We thank Birgit Müller-Kolmstetter for labeling and the support from the perspective of a speech therapist. This work is funded by Bavarian Ministry of Science and Arts.

## 7. References

- [1] C. F. Norbury, D. Gooch, G. Baird, T. Charman, E. Simonoff, and A. Pickles, "Younger children experience lower levels of language competence and academic progress in the first year of school: evidence from a population study," *Journal of Child Psychology and Psychiatry*, vol. 57, no. 1, pp. 65–73, 2016.
- [2] H. Grimm, M. Aktas, and S. Frevert, "Setk 3–5: Sprachentwicklungstest für 3-bis 5-jährige kinder," *Hogrefe: Göttingen, Germany*, 2001.
- [3] A. Baddeley, "Working memory and language: An overview," *Journal of communication disorders*, vol. 36, no. 3, pp. 189–208, 2003.
- [4] S. E. Gathercole and A. D. Baddeley, "Phonological memory deficits in language disordered children: Is there a causal connection?" *Journal of memory and language*, vol. 29, no. 3, pp. 336–360, 1990.
- [5] J. W. Montgomery, "Sentence comprehension in children with specific language impairment: The role of phonological working memory," *Journal of Speech, Language, and Hearing Research*, vol. 38, no. 1, pp. 187–199, 1995.
- [6] Z. Huang, M. Dong, Q. Mao, and Y. Zhan, "Speech emotion recognition using cnn," in *Proceedings of the 22nd ACM international conference on Multimedia*, 2014, pp. 801–804.
- [7] S. Hershey, S. Chaudhuri, D. P. Ellis, J. F. Gemmeke, A. Jansen, R. C. Moore, M. Plakal, D. Platt, R. A. Saurous, B. Seybold *et al.*, "Cnn architectures for large-scale audio classification," in *2017 IEEE international conference on acoustics, speech and signal processing (icassp)*. IEEE, 2017, pp. 131–135.
- [8] M. El Hajji, M. Daniel, and L. Gelin, "Transfer learning based audio classification for a noisy and speechless recordings detection task, in a classroom context," in *SLaTE*, 2019, pp. 109–113.
- [9] J. F. Gemmeke, D. P. Ellis, D. Freedman, A. Jansen, W. Lawrence, R. C. Moore, M. Plakal, and M. Ritter, "Audio set: An ontology and human-labeled dataset for audio events," in *2017 IEEE international conference on acoustics, speech and signal processing (ICASSP)*. IEEE, 2017, pp. 776–780.
- [10] A. Asif, H. Mukhtar, F. Alqadheeb, H. F. Ahmad, and A. Alhumam, "An approach for pronunciation classification of classical arabic phonemes using deep learning," *Applied Sciences*, vol. 12, no. 1, p. 238, 2021.
- [11] M. A. Mohammed, K. H. Abdulkareem, S. A. Mostafa, M. Khanapi Abd Ghani, M. S. Maashi, B. Garcia-Zapirain, I. Oleagordia, H. Alhakami, and F. T. Al-Dhief, "Voice pathology detection and classification using convolutional neural network model," *Applied Sciences*, vol. 10, no. 11, p. 3723, 2020.
- [12] D. Snyder, D. Garcia-Romero, G. Sell, D. Povey, and S. Khudanpur, "X-vectors: Robust dnn embeddings for speaker recognition," in *2018 IEEE international conference on acoustics, speech and signal processing (ICASSP)*. IEEE, 2018, pp. 5329–5333.
- [13] N. Dawalatabad, M. Ravanelli, F. Grondin, J. Thienpondt, B. Desplanques, and H. Na, "Ecapa-tdnn embeddings for speaker diarization," in *Interspeech*, 2021.
- [14] R. Kethireddy, S. R. Kadiri, and S. V. Gangashetty, "Deep neural architectures for dialect classification with single frequency filtering and zero-time windowing feature representations," *The Journal of the Acoustical Society of America*, vol. 151, no. 2, pp. 1077–1092, 2022.
- [15] R. Pappagari, J. Cho, L. Moro-Velázquez, and N. Dehak, "Using State of the Art Speaker Recognition and Natural Language Processing Technologies to Detect Alzheimer's Disease and Assess its Severity," in *Proc. Interspeech 2020*, 2020, pp. 2177–2181. [Online]. Available: <http://dx.doi.org/10.21437/Interspeech.2020-2587>
- [16] R. Scheuerer, T. Haderlein, E. Nöth, and T. Bocklet, "Applying x-vectors on pathological speech after larynx removal," in *2021 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*, 2021, pp. 1079–1086.
- [17] R. Pappagari, T. Wang, J. Villalba, N. Chen, and N. Dehak, "X-vectors meet emotions: A study on dependencies between emotion and speaker recognition," in *ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2020, pp. 7169–7173.
- [18] A. Baeveski, Y. Zhou, A. Mohamed, and M. Auli, "wav2vec 2.0: A framework for self-supervised learning of speech representations," *Advances in Neural Information Processing Systems*, vol. 33, pp. 12 449–12 460, 2020.
- [19] L. Pepino, P. E. Riera, and L. Ferrer, "Emotion recognition from speech using wav2vec 2.0 embeddings," in *Interspeech*, 2021.
- [20] A. Balagopalan and J. Novikova, "Comparing acoustic-based approaches for alzheimer's disease detection," in *Interspeech 2021, 22nd Annual Conference of the International Speech Communication Association*. ISCA, 2021, pp. 3800–3804.
- [21] B. Desplanques, J. Thienpondt, and K. Demuynck, "Ecapa-tdnn: Emphasized channel attention, propagation and aggregation in tdnn based speaker verification," in *INTERSPEECH*, 2020.
- [22] D. Povey, A. Ghoshal, G. Boulianne, L. Burget, O. Glembek, N. Goel, M. Hannemann, P. Motlicek, Y. Qian, P. Schwarz *et al.*, "The kaldı speech recognition toolkit," in *IEEE 2011 workshop on automatic speech recognition and understanding*, no. CONF. IEEE Signal Processing Society, 2011.
- [23] J. L. Metsala, "Young children's phonological awareness and non-word repetition as a function of vocabulary development," *Journal of educational psychology*, vol. 91, no. 1, p. 3, 1999.
- [24] T. Bocklet, C. Winterholler, A. M. Maier, M. Schuster, and E. Nöth, "An automatic screening test for preschool children: theory and data collection," in *Proceedings of the 2nd Workshop on Child, Computer and Interaction*, 2009, pp. 1–4.
- [25] Z.-H. Tan, N. Dehak *et al.*, "rvad: An unsupervised segment-based robust voice activity detection method," *Computer speech & language*, vol. 59, pp. 1–21, 2020.
- [26] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," *CoRR*, vol. abs/1409.1556, 2015.
- [27] M. Ravanelli, T. Parcollet, P. Plantinga, A. Rouhe, S. Cornell, L. Lugosch, C. Subakan, N. Dawalatabad, A. Heba, J. Zhong *et al.*, "Speechbrain: A general-purpose speech toolkit," *arXiv preprint arXiv:2106.04624*, 2021.
- [28] V. Panayotov, G. Chen, D. Povey, and S. Khudanpur, "Librispeech: an asr corpus based on public domain audio books," in *2015 IEEE international conference on acoustics, speech and signal processing (ICASSP)*. IEEE, 2015, pp. 5206–5210.
- [29] B. Milde and A. Köhn, "Open source automatic speech recognition for German," in *Proceedings of ITG 2018*, Oldenburg, Germany, 2018, pp. 251–255.
- [30] D. Povey, V. Peddinti, D. Galvez, P. Ghahremani, V. Manohar, X. Na, Y. Wang, and S. Khudanpur, "Purely sequence-trained neural networks for asr based on lattice-free mmi," in *Interspeech*, 2016, pp. 2751–2755.
- [31] W. T. Fitch and J. Giedd, "Morphology and development of the human vocal tract: A study using magnetic resonance imaging," *The Journal of the Acoustical Society of America*, vol. 106, no. 3, pp. 1511–1522, 1999.
- [32] L. Ardaillon and A. Roebel, "Fully-convolutional network for pitch estimation of speech signals," in *Interspeech 2019*, 2019.
- [33] M. Abadi, P. Barham, J. Chen, Z. Chen, A. Davis, J. Dean, M. Devin, S. Ghemawat, G. Irving, M. Isard *et al.*, "{TensorFlow}: A system for {Large-Scale} machine learning," in *12th USENIX symposium on operating systems design and implementation (OSDI 16)*, 2016, pp. 265–283.
- [34] D. Kingma and J. Ba, "Adam: A method for stochastic optimization," *International Conference on Learning Representations*, 12 2014.
- [35] S. Ioffe and C. Szegedy, "Batch normalization: Accelerating deep network training by reducing internal covariate shift," in *International conference on machine learning*. PMLR, 2015, pp. 448–456.
- [36] T. Fawcett, "Roc graphs: Notes and practical considerations for researchers," *Machine learning*, vol. 31, no. 1, pp. 1–38, 2004.