



Exploring linguistic feature and model combination for speech recognition based automatic AD detection

Yi Wang, Tianzi Wang, Zi Ye, Lingwei Meng, Shoukang Hu, Xixin Wu, Xunying Liu, Helen Meng

The Chinese University of Hong Kong, Hong Kong SAR, China

{ywang, twang, zye, lmeng, skhu, wuxx, xyliu, hmmeng}@se.cuhk.edu.hk

Abstract

Early diagnosis of Alzheimer's disease (AD) is crucial in facilitating preventive care and delay progression. Speech based automatic AD screening systems provide a non-intrusive and more scalable alternative to other clinical screening techniques. Scarcity of such specialist data leads to uncertainty in both model selection and feature learning when developing such systems. To this end, this paper investigates the use of feature and model combination approaches to improve the robustness of domain fine-tuning of BERT and Roberta pre-trained text encoders on limited data, before the resulting embedding features being fed into an ensemble of backend classifiers to produce the final AD detection decision via majority voting. Experiments conducted on the ADReSS20 Challenge dataset suggest consistent performance improvements were obtained using model and feature combination in system development. State-of-the-art AD detection accuracies of 91.67% and 93.75% were obtained using manual and ASR speech transcripts respectively on the ADReSS20 test set consisting of 48 elderly speakers.

Index Terms: AD detection, speech recognition, uncertainty, model ensemble, pre-trained feature representation

1. Introduction

Alzheimer's disease (AD), the most common form of dementia often found in aged people, is characterized by progressive degradation of the memory, cognition, and motor skills, and consequently decline in the speech and language skills of patients [1, 2]. Currently, there is no effective cure for AD [3], but a timely intervention approach can delay its progression and reduce the negative physical and mental impact on patients [4]. To this end, speech-based automatic AD diagnosis has advantages of being non-intrusive, more scalable and less costly, when compared with conventional approaches based on clinical tests. In recent years there has been a burgeoning interest in developing such systems, notably represented by the recent ADReSS challenges [5, 6].

Current speech-based AD detection systems consist of front-end feature extraction and back-end classification components. Commonly used forms of acoustic features extracted from the audio include, but not limited to, spectral features, prosodic features, vocal quality, and pre-trained deep neural network (DNN) embedding features [7–15]. Handcrafted linguistic features [7, 10, 16–18] capturing lexical or syntactical cues, or those pre-trained text neural embedding [7–9, 19–21] can be extracted from the transcripts of AD assessment speech recordings. To fully automate the AD detection process, ASR systems [12–15, 17, 18, 21–25] have been increasingly often used to produce speech transcripts, in place of ground truth manual transcripts [7, 8, 16, 19, 22, 26, 27] that are costly to annotate in large quantities. In addition, disfluency measures of speech were found useful for AD detection tasks [13, 18, 19, 21, 28]. Diverse forms of back-end classifiers range from DNN vari-

ants [13, 20, 25, 29] to SVM, decision trees and further combination among these [30]. Recent researches have widely reported state-of-the-art AD detection accuracy performance can be obtained using linguistic features [7, 8, 19, 21, 22, 31, 32]. Some of them further used model ensembling methods to exploit linguistic features and enhance integration among inner-modal information [8, 19, 21, 32].

One major challenge confronting the development of automatic speech-based AD detection systems is the scarcity of the required specialist data due to the difficulty in collecting large quantities of transcribed AD assessment speech recordings and shortage of clinical professionals. The resulting "small data" learning task leads to model uncertainty encountered in multiple stages of system development. These include, but not limited to, the following considered in this paper: a) cross-domain transfer text embedding DNNs, for example, state-of-the-art BERT [33] and Roberta [34] pre-trained models, by parameter fine-tuning on limited in-domain AD detection data, leading to over-confident parameters; b) possible errors in the transcripts produced by ASR systems; and c) the suitable form of back-end classifier design and parameter estimation.

To this end, feature and ensemble approaches are adopted in this paper as overarching solutions to address the above model uncertainty issues. First, the robustness of domain fine-tuning of BERT and Roberta pre-trained text encoders on limited AD detection speech transcripts is improved using model averaging. Second, the uncertainty over the suitable form of text embedding features is accounted for using both BERT and Roberta text embedding representations for back-end classifier construction. Third, the uncertainty over ASR transcript's quality is further handled using two-pass decoding based speech recognition system combination, or later fusion of their respective AD detection outputs. Lastly, the uncertainty over the back-end classifier is addressed using decision voting among an ensemble of classifiers.

Experiments conducted on the ADReSS20 Challenge dataset suggest consistent performance improvements were obtained using model and feature combination in system development. State-of-the-art AD detection accuracies of 91.67% and 93.75% were obtained using manual and state-of-the-art hybrid CNN-TDNN ASR system produced speech transcripts respectively on the ADReSS20 test set consisting of 48 elderly speakers.

The main contributions of this paper are summarized below. First, this paper presents a systematic application of feature and model combination approaches to address the inherent data sparsity and model uncertainty issues in speech recognition based AD detection systems. In contrast, prior research primarily investigated model combination over back-end classifiers [8, 10, 19, 21, 27, 32], while its application to cross-domain fine-tuning of pre-trained text embedding and ASR system outputs fusion remains limited. Second, this paper presents a

state-of-the-art AD detection accuracy of 91.67% obtained using manual transcripts and best published accuracy of 93.75% using ASR speech transcripts respectively on the ADReSS20 test set. To the best of our knowledge, the closest published accuracy performance to date on the same task are 91.67% [32] and 87.50% [22] using manual and ASR speech transcripts, respectively.

The rest of this paper is organized as follows. Section 2.1 introduces the AD detection task and the ADReSS20 [5] dataset used in this work. Section 2.2 briefly describes the hybrid CNN-TDNN and end-to-end Conformer ASR systems, BERT and Roberta text embedding models, and back-end classifiers. Section 3 suggests an approach to fine-tune the pre-trained text encoders. Section 4 presents AD detection systems with feature combination and classifier combination methods on manual transcripts, followed by Section 5 showing AD detection with those methods using ASR system outputs. Finally, the conclusions are drawn, and future works are discussed in Section 6.

2. Task Description

This subsection describes the audio and text data used in this paper and the baseline system structure.

2.1. Dataset

In this paper, the ADReSS20 Challenge dataset [5] is used for the AD detection system training and evaluation. It is selected from the Cookie Theft picture description part of the Pitt Corpus in the DementiaBank database [35] with balanced age, gender and AD label distribution among subjects. The data consists of speech recordings, corresponding manual transcripts annotated by the CHAT coding system [36] and binary AD labels. Both the participant and the interviewer parts are provided. It is divided into the training set of 2 hours 9 minutes from 108 subjects and the test set of 1 hour 6 minutes from 48 subjects.

2.2. Baseline System

A baseline tandem BERT text embedding and SVM classifier based system structure was used [22] to produce AD detection results with either ground truth manual transcripts or ASR system generated transcripts.

ASR Systems: The hybrid CNN-TDNN system is based on a recently proposed lattice-free MMI trained factored TDNN system [37] with additional 2-dimensional convolutional layers being employed as the first 6 layers, and is further optimized with data re-segmentation, domain adaptation, and Bayesian LHUC speaker adaptation [38]. The end-to-end ASR system used is a Conformer specially developed for elderly speech by training with augmented elderly speech data, incorporating manual re-configuration of Conformer hyper-parameters and cross-domain adaptation. Two-pass decoding based system combination ASR outputs involving the CNN-TDNN ASR system to produce initial N-best recognition hypotheses before being rescored using the Conformer system using a 2-way cross-system score interpolation [39, 40] are also considered.

BERT and Roberta text embedding: we combine all participant speech segments corresponding to a single picture description for extracting text features. The principal 768-dimensions of the last hidden layer BERT or Roberta model output are taken as the single vector representation for each participant’s AD classification. This is then normalized by removing mean and scaling to unit variance before being passed into the classifiers. The BERT pre-trained model `bert-base-uncased` is based on the open-source model library ¹, with 12 transformer

¹<https://huggingface.co/bert-base-uncased>

blocks, 768 hidden layers, and 12 attention heads per attention layer. The max input length is 512. The standard BERT tokenizer transforms input texts into tokens. This paper also explored the use of Roberta pre-trained text feature embedding model and the accompanying tokenizer².

Classifier: The SVM classifier used SVC API from the Sklearn package, with the regularization parameter C for soft margin set to 1 and linear kernel as the kernel function. The other four classification models used are Linear Discriminant Analysis (LDA), Gaussian Process (GP), Multilayer Perceptron (MLP)³, and extreme gradient boost (XGB)⁴. Each classifier is optimized with the best hyper-parameters found via greedy search on cross-validation data using BERT derived text features of manual transcripts. The best-performed classifier hyper-parameters are adopted and fixed in all the experiments of this paper. We evaluated the system performance by 10-fold cross-validation (CV)⁵ on the training set and tested on the test set.

3. Fine-tuning Pre-trained Text Encoder

To reduce the mismatch between pre-trained BERT/Roberta models built on general texts and elderly speech transcripts, fine-tuning was performed on the 11591 words manual transcripts of the ADReSS training set. We fine-tuned the BERT or Roberta models with the Masked Language Modelling and Next Sentence Prediction tasks used in pre-training. To circumvent the risk of over-fitting, a selected set of three fine-tuned pre-trained models on the ADReSS training data obtained at the final three update epochs were used to produce separate text embedding features for SVM based AD detection outputs majority voting, as is shown in Figure 1 (A). Ablation studies on alternative update epoch selection schemes, for example, based on random epoch sampling, larger or non-even inter-epoch intervals, suggest the above provides the best AD detection performance in practice. The performance of their corresponding AD detection accuracy on the training set 10-fold CV and the test set are shown in Table 1. Notably, during the evaluation, we aggregated the scores from ten runs of CV by performing decision voting from models trained in ten runs, as it generated higher stability and reduced the uncertainty in model selection. Accuracy scores are primarily used to evaluate task performance.

Several trends can be found in Table 1. First, irrespective of the precise form of epoch selection being used prior to their combination shown in Figure 1, a trend is that performance improvements were obtained over the non-fine-tuned baseline BERT models, and the directly fine-tuned models without inter-epoch combination (Sys.3 vs. Sys.1-2). Second, epoch selection using a fixed stride with the optimal epoch selection stride 1 (Sys.3 vs. Sys.4-5) outperformed the others selection methods (Sys.3 vs. Sys.6-7). More specifically, the best performing fine-tuned BERT model produced features with merging over the last three update epochs outperformed the baseline BERT features without fine-tuning and model merging by 11.36% absolute on the CV data (Sys.3 vs. Sys.1). A more balanced performance across the CV and test sets were also obtained. Similar trends of accuracy improvements were also found on the fine-tuned Roberta features with model combination (Sys.9 vs. Sys.8). Based on these results, BERT and Roberta text encoder fine-tuning with model combination over the last update snapshots

²<https://huggingface.co/roberta-base>

³<https://scikit-learn.org/stable/>

⁴<https://xgboost.readthedocs.io/en/latest/python/index.html>

⁵In ADReSS20 dataset, each participant only has one recording, so CV is conducted on participant-level

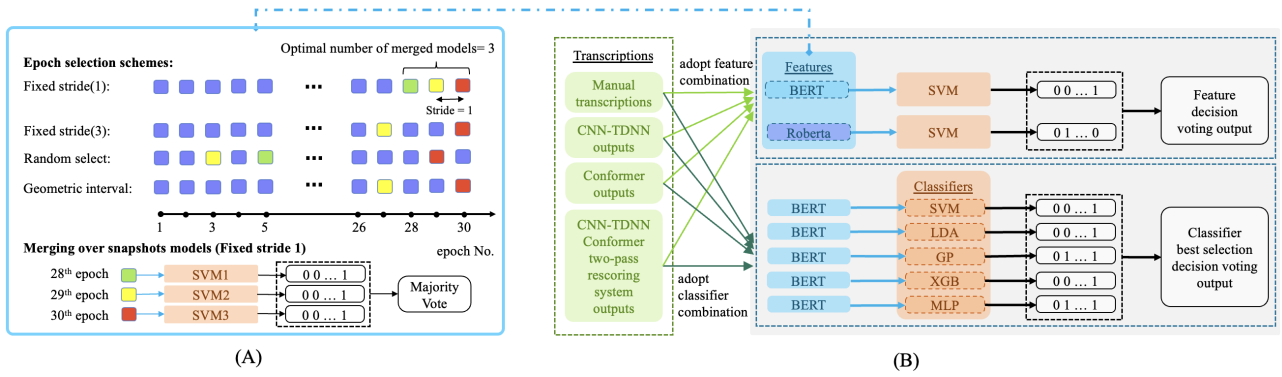


Figure 1: Example of (A) an ensemble of 3 fine-tuned BERT/Roberta pre-trained models on the ADRess training data obtained at different update epochs (top, left) and being used to produce separate text embedding features for SVM based AD detection outputs majority voting (in green/yellow/red, bottom, left); and (B) AD detection system pipeline using manual, or one of three ASR systems (CNN-TDNN, Conformer, or their combination via two-pass recognition [39]) produced transcripts (green) for text embedding features extraction (blue) and back-end classification (orange). Feature combination feeds fine-tuned BERT and Roberta text embedding of (A) into a single back-end SVM classifier for outputs voting (top right). Classifiers combination feeds fine-tuned BERT or Roberta text features into five classifiers for outputs voting (bottom right).

Table 1: AD detection performance evaluated by 10-fold CV on the ADRess training set and test set using BERT/Roberta as the text feature extractor w/o model combination based fine-tuning of Fig. 1(A) using SVM classifier. All BERT/Roberta models fine-tuned on ground truth transcript (GT). All back-end classifiers trained and evaluated on ground truth transcript (GT). Sys.3-7 and Sys.9 show the effects of various update epochs selection schemes for model combination in Section 3. Random result (Sys.6) is averaged from 10 random selections. Fine-tuning without merging result (Sys.2) is averaged from performances of the selected single snapshot models.

Sys.	Feature	Fine-tuning Data	Combination	Epoch selection	Epochs merged	CV Acc.(%)	Test Acc.(%)
1	BERT	N/A	×	-	-	74.75	87.50
2		GT	×	fixed stride(1)	[28,29,30]	78.38	84.70
3			fixed stride(1)	[28,29,30]	86.11	85.42	
4			fixed stride(3)	[24,27,30]	86.11	79.17	
5			fixed stride(10)	[10,20,30]	84.26	81.25	
6			random	-	81.39	79.58	
7			geo. interval	[18, 27, 30]	82.41	83.33	
8			N/A	×	-	-	70.79
9		Roberta	GT	✓	fixed stride(1)	[28,29,30]	80.56

are used in all the following experiments of this paper.

4. Feature and Classifier Combination

This section investigates the feature combination of two text encoders (BERT and Roberta), and the back-end classifiers combination, with the manual transcriptions as inputs.

4.1. Feature Combination

The two types of fine-tuned text embedding DNNs, BERT and Roberta, can capture linguistic information with discrepancy. To further reduce the uncertainty underlying the resulting text embedding features, we exploited a late feature fusion via majority voting over decisions made by a single type of classifier based on either BERT and Roberta features respectively (Feature decision voting, Table 2), which is shown in Figure 1 (B, top right). The effectiveness of feature decision voting with different types of back-end classifiers is examined. We also explored alternative methods of fusing the two text embedding features by concatenation before feeding them into SVMs.

Two trends are observed from the manual transcripts based results in Table 2 (column 5-6). First, the method of feature fusion by decision-level majority voting produced improvements using all 5 back-end classifiers on the CV data over the comparable baselines using either BERT or Roberta features alone (Sys.13-17 vs. Sys.1-5 & Sys.7-11). On the test set, similar improvements were obtained for SVM and LDA classifiers (Sys.13 vs. Sys.1&7, Sys.14 vs. Sys.2&8). Second, the

decision-level fusion by majority voting performed better than early fusion by concatenation (Sys.13 vs. Sys.19, col.5-6, Table 2). Lastly, decision-level fusion among BERT and Roberta feature based SVM classifiers produced the best feature combination accuracy of 87.96% on the CV data, and 87.50% on the test (Sys.13 vs. 14-17, col.5-6, Table 2).

4.2. Classifiers Combination

Apart from the feature extractor stage, uncertainty could also manifest in architecture design and parameter estimation of back-end AD detection models on limited AD detection training data. To this end, we conducted classifier combination with majority voting over the binary AD detection decisions by a range of BERT (or Roberta) feature based back-end classifiers, as illustrated in Figure 1 (grey, bottom right). All possible feature-classifier combinations involving both the full set 2×5 combination and its various subsets were examined to maximize the diversity and complementarity among systems. The best performing feature-classifier combination configuration involves fine-tuned (with snapshot level model voting of Section 3) BERT or Roberta features trained LDA, GP and MLP classifiers (Sys. 18, col.5-6, Table 2).

The hyperparameters in the back-end classifiers are optimized by greedy search on CV, as mentioned in Section 2.2. More specifically, we used the SVM setting as in the baseline, LDA solved by SVD, GP with kernel $4.0 * RBF(5.0)$, XGB based on tree models with the number of boosting round set as 16, the maximum tree depth for base learners set as 2, learning rate $\eta = 0.4$, minimum loss reduction required to make a partition on the tree $\gamma = 0.1$, and MLP with a single hidden layer of 256-dimension, L2 penalty parameter $\alpha = 1e - 5$, optimized by limited memory BFGS. All other hyper-parameters were kept as default settings.

As shown in Table 2, the first general trend is that classifier combination system consistently surpasses, or at least matches, the performance of any single classifier system across the CV and test sets (Sys.6 vs. Sys.1-5, Sys.12 vs. Sys.7-11, col.5-6, Table 2). Secondly, conducting classifier combination on the top of feature combination produced no further improvements over pure feature combination (Sys.18 vs. Sys.13, col.5-6, Table 2). This suggests further cross-system diversity and complementarity need to be injected into individual back-end classifiers' design and training when feature combination is used.

5. ASR Output Based AD Detection

In this Section we apply the feature combination and classifier combination methods of Section 4 to AD detection using ASR

Table 2: Accuracy of AD detection using different back-end classifiers with or without feature and/or model combination on manual transcripts, or one of three ASR systems (CNN-TDNN, Conformer, or their two-pass decoding combination [39] “CNN-TDNN → Conformer”) produced transcripts. “CNN-TDNN + Conformer” refers to separate AD detection using each ASR system’s specific transcripts before final decision voting combination. “LDA” stands for linear discriminant analysis, “GP” for Gaussian Process, “MLP” for multilayer perceptron, and “XGB” for extreme gradient boosting.

Sys.	Feature	Combination Method	Classifier	Manual WER 0%		CNN-TDNN WER 30.69%	Conformer WER 26.70%	CNN-TDNN → Conformer WER 26.50%	CNN-TDNN + Conformer WER N/A
				CV Acc.(%)	Test Acc.(%)	Test Acc.(%)	Test Acc.(%)	Test Acc.(%)	Test Acc.(%)
1	BERT	N/A	SVM	86.11	85.42	85.42	81.25	81.25	81.25
2			LDA	86.11	79.17	72.92	72.92	81.25	72.92
3			GP	85.19	81.25	79.17	85.42	81.25	83.33
4			XGB	84.26	79.17	79.17	70.83	79.17	79.17
5			MLP	84.26	81.25	79.17	79.17	81.25	83.33
6			Classifier Best Selection	SVM + LDA + GP + XGB	87.04	85.42	79.17	81.25	81.25
7	Roberta	N/A	SVM	80.56	85.42	87.50	81.25	83.33	85.42
8			LDA	82.41	85.42	81.25	75.00	81.25	81.25
9			GP	79.63	87.50	81.25	81.25	83.33	83.33
10			XGB	81.48	85.42	83.33	85.42	77.08	87.50
11			MLP	84.26	89.58	77.08	79.17	79.17	83.33
12			Classifier Best Selection	LDA + XGB + MLP	84.26	91.67	81.25	79.17	83.33
13	Bert + Roberta	Feature decision voting	SVM	87.96	87.50	93.75	85.42	87.50	91.67
14			LDA	86.11	87.50	83.33	72.92	85.42	85.42
15			GP	85.19	83.33	83.33	83.33	85.42	87.50
16			XGB	86.11	81.25	85.42	83.33	87.50	87.50
17			MLP	87.04	85.42	81.25	85.42	81.25	81.25
18			Feature + Classifier.	LDA + GP + MLP	87.96	85.42	79.17	83.33	81.25

system generated transcripts.

5.1. AD Detection Using Single ASR System’s Output

The AD detection performance using various feature and classifier combination configurations for a single ASR system produced transcripts are shown in Table 2 (col.7-8). Several trends can be found. First, consistent with the previous trends on manual speech transcripts using feature combination (col.5-6, Table 2), decision-level feature fusion by majority voting produced improvements using all 5 back-end classifiers on the test data over the comparable baselines using either BERT or Roberta features alone (Sys.13-17 vs. Sys.1-5 & Sys.7-11, col.7-8, Table 2). Second, the BERT+Roberta feature combination approach based SVM classifier (Sys.13), among all the combination strategies being considered in the same table, consistently produced the best AD detection accuracy using the two single ASR systems’ outputs (Sys.13 vs. Sys.1-12 & 14-19, col.7-8, Table 2). In particular, using the CNN-TDNN ASR system produced speech transcripts, this decision-level feature combination based SVM system produced the best published speech recognition outputs based AD detection accuracy of 93.75% on the benchmark ADReSS test set.

5.2. AD Detection Using Multiple ASR Systems’ Outputs

Compared with ASR performance measured on non-aged, healthy speech [41, 42], significantly higher speech recognition error rates are often found on elderly speech data [2, 15, 22, 23, 31, 43–45]. In order to mitigate the impact of possible ASR transcript errors on the downstream AD detection task, ASR system combination approaches are also investigated in this paper to account for the uncertainty over the quality of speech recognition system outputs produced by the hybrid CNN-TDNN and end-to-end Conformer ASR systems considered in this paper. In this regard, we studied the efficacy of two approaches to combine the hybrid CNN-TDNN and end-to-end Conformer ASR systems. One is detection based on their two-pass decoding combined system produced outputs [39] (“CNN-TDNN → Conformer”, Table 2), while the other performs separate AD detection using either of the two ASR system’s output transcripts before final decision voting combination (“CNN-TDNN+Conformer”, Table 2).

From results presented in Table 2 (col.9-10), we first confirmed that BERT plus Roberta feature decision voting system with the SVM classifier is again the best performing system among all (Sys.13 vs. Sys.1-12 & 14-19, col.9-10, Table 2),

Table 3: Comparison of test results (rounded to one decimal place) against the SOTA based on different modality for the ADReSS20 challenge. “Pre.,” “Rec.” and “F1” represents precision, recall and F1 scores regarding AD as the relevant class.

Study	Modality	Test			
		Acc.(%)	Pre.(%)	Rec.(%)	F1(%)
Syed et al. [32]	Text (Manual)	91.7	-	91.7	-
Li et al. [22]	Text (ASR)	88.0	91.0	83.0	87.0
Martinc et al. [46]	Audio + Text (Manual)	93.8	-	-	-
Laguarta et al. [47]	Audio + pre-trained biomarkers ⁶	93.8	-	-	-
Our work (Sys.13, Tab. 2)	Text (Manual)	91.7	91.7	91.7	91.7
	Text (ASR)	93.8	92.0	95.8	93.9

regardless of which ASR combination approach being used. However, using neither of the two ASR system combined outputs led to improvements over using the hybrid CNN-TDNN system transcripts only (Sys. 13, col.9-10 vs. col.7). This suggests further diversity and complementarity may be required in ASR system combination and integration with the downstream AD detection component. The performance of the best performing AD detection system in this paper (sys. 13, Table 2) are further contrasted in Table 3 with those that are obtained on the same ADReSS20 challenge task and report in the recent literature to demonstrate their competitiveness.

6. Conclusion

The development of a state-of-the-art speech-based AD detection system constructed using the ADReSS20 challenge dataset was presented in this paper. AD detection was performed based on either manual or ASR system transcripts. A series of model ensemble approaches were used in fine-tuning of BERT and Roberta pre-trained text encoders, text embedding feature and back-end classifier combination, as well as the integration with front-end ASR system to address the modeling uncertainty issue resulted from the sparsity of AD detection training data. Experiments conducted on the ADReSS20 Challenge dataset suggest consistent performance improvements were obtained using model and feature combination in AD detection system development. Tighter integration between the ASR system and AD detection components will be investigated in future research.

7. Acknowledgements

This research is supported by Hong Kong RGC GRF grant No. 14200021, 14200218, 14200220, TRS T45-407/19N and Innovation & Technology Fund grant No. ITS/254/19.

⁶They transferred knowledge from pre-trained detection system using biomarkers (e.g. cough features) developed with non-public data

8. References

- [1] P. S. Sachdev *et al.*, “Classifying neurocognitive disorders: the dsm-5 approach,” *Nat. Rev. Neurol.*, 2014.
- [2] R. Vipperla *et al.*, “Ageing voices: The effect of changes in voice parameters on asr performance,” *Eurasip J. Audio Speech Music Process.*, 2010.
- [3] V. G. Yadav, “The hunt for a cure for alzheimer’s disease receives a timely boost,” *Sci. Transl. Med.*, 2019.
- [4] D. Roeck *et al.*, “Brief cognitive screening instruments for early detection of alzheimer’s disease: a systematic review,” *Alzheimer’s research & therapy*, 2019.
- [5] S. Luz *et al.*, “Alzheimer’s dementia recognition through spontaneous speech: The adress challenge,” *INTERSPEECH*, 2020.
- [6] S. Luz *et al.*, “Detecting cognitive decline using speech only: The addresso challenge,” *INTERSPEECH*, 2021.
- [7] A. Balagopalan *et al.*, “To bert or not to bert: comparing speech and language-based approaches for alzheimer’s disease detection,” *INTERSPEECH*, 2020.
- [8] M. S. S. Syed *et al.*, “Automated screening for alzheimer’s dementia through spontaneous speech,” in *INTERSPEECH*, 2020.
- [9] E. Edwards *et al.*, “Multiscale system for alzheimer’s dementia recognition through spontaneous speech,” in *INTERSPEECH*, 2020.
- [10] M. Martinc *et al.*, “Tackling the adress challenge: A multimodal approach to the automated recognition of alzheimer’s dementia,” in *INTERSPEECH*, 2020.
- [11] J. Koo *et al.*, “Exploiting multi-modal features from pre-trained networks for alzheimer’s dementia recognition,” *INTERSPEECH*, 2020.
- [12] Z. S. Syed *et al.*, “Tackling the addresso challenge 2021: The muet-rmit system for alzheimer’s dementia recognition from spontaneous speech,” *INTERSPEECH*, 2021.
- [13] M. Rohanian *et al.*, “Alzheimer’s dementia recognition using acoustic, lexical, disfluency and speech pause features robust to noisy inputs,” *INTERSPEECH*, 2021.
- [14] R. Pappagari *et al.*, “Automatic detection and assessment of alzheimer disease using speech and language technologies in low-resource scenarios,” in *INTERSPEECH*, 2021.
- [15] Y. Pan *et al.*, “Using the outputs of different automatic speech recognition paradigms for acoustic-and bert-based alzheimer’s dementia detection through spontaneous speech,” in *INTERSPEECH*, 2021.
- [16] L. B. D. Santos *et al.*, “Enriching complex networks with word embeddings for detecting mild cognitive impairment from speech transcripts,” *arXiv preprint arXiv:1704.08088*, 2017.
- [17] R. B. Ammar *et al.*, “Speech processing for early alzheimer disease diagnosis: machine learning based approach,” in *AICCSA*, 2018.
- [18] Y. Chien *et al.*, “An assessment system for alzheimer’s disease based on speech using a novel feature sequence design and recurrent neural network,” in *SMC*, 2018.
- [19] J. Yuan *et al.*, “Disfluencies and fine-tuning pre-trained language models for detection of alzheimer’s disease,” in *INTERSPEECH*, 2020.
- [20] A. Pompili *et al.*, “The inesc-id multi-modal system for the adress 2020 challenge,” *INTERSPEECH*, 2020.
- [21] Y. Qiao *et al.*, “Alzheimer’s disease detection from spontaneous speech through combining linguistic complexity and (dis) fluency features with pretrained language models,” *INTERSPEECH*, 2021.
- [22] J. Li *et al.*, “A comparative study of acoustic and linguistic features classification for alzheimer’s disease detection,” in *ICASSP*, 2021.
- [23] G. Gosztolya *et al.*, “Identifying mild cognitive impairment and mild alzheimer’s disease based on spontaneous speech using asr and linguistic features,” *CSL*, 2019.
- [24] Y. Zhu *et al.*, “Waybert: Exploiting semantic and non-semantic speech using wav2vec and bert for dementia detection,” *INTERSPEECH*, 2021.
- [25] J. Chen *et al.*, “An attention-based hybrid network for automatic detection of alzheimer’s disease from narrative speech,” in *INTERSPEECH*, 2019.
- [26] T. Searle *et al.*, “Comparing natural language processing techniques for alzheimer’s dementia prediction in spontaneous speech,” *INTERSPEECH*, 2020.
- [27] U. Sarawgi *et al.*, “Multimodal inductive transfer learning for detection of alzheimer’s dementia and its severity,” *INTERSPEECH*, 2020.
- [28] R. Pappagari *et al.*, “Using state of the art speaker recognition and natural language processing technologies to detect alzheimer’s disease and assess its severity,” in *INTERSPEECH*, 2020.
- [29] N. Cummins *et al.*, “A comparison of acoustic and linguistics methodologies for alzheimer’s dementia recognition,” in *INTERSPEECH*, 2020.
- [30] S. de la Fuente Garcia *et al.*, “Artificial intelligence, speech, and language processing approaches to monitoring alzheimer’s disease: a systematic review,” *Journal of Alzheimer’s Disease*, 2020.
- [31] Z. Ye *et al.*, “Development of the cuhk elderly speech recognition system for neurocognitive disorder detection using the dementia-bank corpus,” in *ICASSP*, 2021.
- [32] Z. S. Syed *et al.*, “Automated recognition of alzheimer’s dementia using bag-of-deep-features and model ensembling,” *IEEE Access*, 2021.
- [33] J. Devlin *et al.*, “Bert: Pre-training of deep bidirectional transformers for language understanding,” *arXiv preprint arXiv:1810.04805*, 2018.
- [34] Y. Liu *et al.*, “Roberta: A robustly optimized bert pretraining approach,” *arXiv preprint arXiv:1907.11692*, 2019.
- [35] J. T. Becker *et al.*, “The natural history of alzheimer’s disease: description of study cohort and accuracy of diagnosis,” *Archives of neurology*, 1994.
- [36] B. MacWhinney, “The childes project: Tools for analyzing talk: Volume i: Transcription format and programs, volume ii: The database,” 2000.
- [37] A. Gulati *et al.*, “Conformer: Convolution-augmented transformer for speech recognition,” *INTERSPEECH*, 2020.
- [38] X. Xie *et al.*, “Blhuc: Bayesian learning of hidden unit contributions for deep neural network speaker adaptation,” in *ICASSP*, 2019.
- [39] T. Wang *et al.*, “Conformer based elderly speech recognition system for alzheimer’s disease detection,” *INTERSPEECH*, 2022.
- [40] M. Cui *et al.*, “Two-pass decoding and cross-adaptation based system combination of end-to-end conformer and hybrid tdnn asr systems,” *INTERSPEECH*, 2022.
- [41] C. Lüscher *et al.*, “Rwth asr systems for librispeech: Hybrid vs attention,” *INTERSPEECH*, 2019.
- [42] Z. Tüske *et al.*, “On the limit of english conversational speech recognition,” *INTERSPEECH*, 2021.
- [43] M. Geng *et al.*, “Speaker adaptation using spectro-temporal deep features for dysarthric and elderly speech recognition,” *In submission to IEEE T AUDIO SPEECH*, 2022.
- [44] F. Rudzicz *et al.*, “Speech recognition in alzheimer’s disease with personal assistive robots,” in *SLPAT*, 2014, pp. 20–28.
- [45] L. Zhou *et al.*, “Speech recognition in alzheimer’s disease and in its assessment,” in *INTERSPEECH*, 2016.
- [46] M. Martinc *et al.*, “Temporal integration of text transcripts and acoustic features for alzheimer’s diagnosis based on spontaneous speech,” *Frontiers in Aging Neuroscience*, 2021.
- [47] J. Laguarda *et al.*, “Longitudinal speech biomarkers for automated alzheimer’s detection,” *Frontiers in Computer Science*, 2021.