



Leveraging Cascaded Binary Classification and Multimodal Fusion for Dementia Detection through Spontaneous Speech

Yin-Long Liu^{1,4}, Yuanchao Li³, Rui Feng¹, Liu He², Jia-Xin Chen^{1,4}, Yi-Ming Wang¹, Yu-Ang Chen⁴, Yan-Han Peng⁴, Jia-Hong Yuan^{1,2*}, Zhen-Hua Ling^{1,2,4}

¹National Engineering Research Center of Speech and Language Information Processing, University of Science and Technology of China, Hefei, P. R. China

²Interdisciplinary Research Center for Linguistic Sciences, University of Science and Technology of China, Hefei, P. R. China

³Centre for Speech Technology Research, University of Edinburgh, UK

⁴Department of Electronic Engineering and Information Science, University of Science and Technology of China, Hefei, P. R. China

lyl2001@mail.ustc.edu.cn, {jiahongyuan, zhling}@ustc.edu.cn

Abstract

This paper presents our submission to the PROCESS Challenge 2025, focusing on spontaneous speech analysis for early dementia detection. For the three-class classification task (Healthy Control, Mild Cognitive Impairment, and Dementia), we propose a cascaded binary classification framework that fine-tunes pre-trained language models and incorporates pause encoding to better capture disfluencies. This design streamlines multi-class classification and addresses class imbalance by restructuring the decision process. For the Mini-Mental State Examination score regression task, we develop an enhanced multimodal fusion system that combines diverse acoustic and linguistic features. Separate regression models are trained on individual feature sets, with ensemble learning applied through score averaging. Experimental results on the test set outperform the baselines provided by the organizers in both tasks, demonstrating the robustness and effectiveness of our approach.

Index Terms: Dementia Detection, Cascaded Binary Classification, Pre-trained Language Models, Multimodal Fusion

1. Introduction

Dementia is a neurodegenerative disease that worsens over time and causes irreversible damage to the brain, manifested by a persistent deterioration of an individual's cognitive and functional abilities, including memory, attention, and executive function [1]. Mild Cognitive Impairment (MCI), an early dementia symptom, is typically regarded as a transitional stage between normal aging and dementia, leading to a mild decline in cognitive functions without significantly impairing daily activities. Today, several biochemical methods, such as magnetic resonance imaging and positron emission tomography, have been developed for early dementia evaluation [2, 3]. However, these methods are often expensive, demanding, time-consuming, and require validation by neurologists in manual clinical settings. Therefore, more efficient and reliable automated methods are needed to achieve early dementia detection [4]. While memory impairment is widely recognized as the hallmark symptom of dementia, speech and language impairments are equally prevalent and serve as critical indicators for the early detection of the disease [5, 6]. Compared to traditional clinical methods, speech- and language-based automatic dementia diagnosis has garnered increasing attention due to its non-invasive, cost-effective, and

convenient nature [7–11]. PROCESS (Prediction and Recognition Of Cognitive declinE through Spontaneous Speech) Signal Processing Grand Challenge [12] focuses on early dementia detection through the analysis of spontaneous speech. This challenge consists of two tasks: a classification task and a regression task. Unlike the binary classification tasks in previous challenges [13–15], PROCESS involves a three-class classification task, aiming to predict the diagnostic label of the subjects as Healthy Control (HC), MCI, or Dementia. The inclusion of the MCI category increases the difficulty of the classification task. The regression task is similar to those in previous challenges, aiming to predict the Mini-Mental State Examination (MMSE) score reflecting cognitive decline. In this work, we propose:

- A cascaded binary classification framework that incorporates pause encoding to effectively address multi-class challenges.
- An enhanced multimodal fusion system that seamlessly integrates diverse acoustic and linguistic features for improved performance.

Experimental results, in contrast to the baseline, validate the efficacy of our approach and highlight its potential in multi-class classification and early detection of dementia.

2. Related Work

Current research primarily utilizes two types of features extracted from spontaneous speech for dementia detection: acoustic features from raw audio signals and linguistic features from transcripts [10, 14, 16]. Both features offer unique insights into cognitive decline associated with Dementia. Based on the recent studies [17–20], acoustic features used in dementia detection can be divided into prosody, duration of pauses, emotional embeddings, pre-trained models embeddings features, and more. Key linguistic features include, but are not limited to, lexical richness, syntactic complexity and pre-trained textual embedding features [21–23]. Previous research has shown that, compared to using acoustic features alone, employing linguistic features extracted from transcripts is more effective in distinguishing between Dementia and HC [18, 24]. Advanced Natural Language Processing (NLP) techniques, particularly those involving Pre-trained Language Models (PLMs) such as BERT [25], RoBERTa [26], and ERNIE [27], are capable of capturing complex patterns and contextual information that are crucial for detecting cognitive decline, and have proven effective in dementia detection [28, 29].

Nevertheless, very limited research has explored the three-

* Corresponding author.

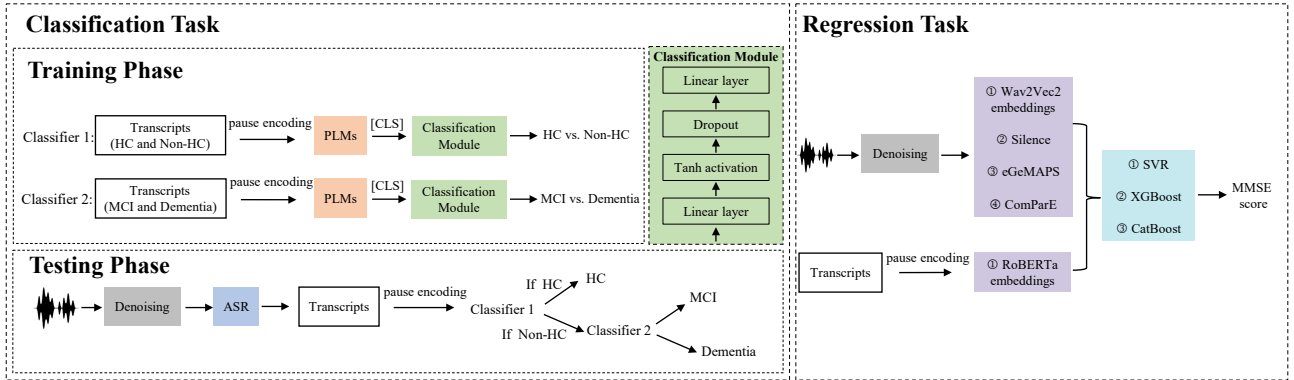


Figure 1: Architecture of methods for classification and regression tasks. PLMs refer to the pre-trained language models BERT, RoBERTa, and ERNIE. Non-HC refers to a label that is either MCI or Dementia. For the training phase, we use manual transcripts, while during the testing phase, ASR transcripts are used. The regression task shown in the figure represents the training phase.

class classification problem in dementia detection, posing challenges in the early diagnosis of dementia, particularly MCI. To address this long-standing issue, we propose a cascaded binary classification approach. Furthermore, we investigate the integration of multimodal features to enhance performance in the regression task.

3. Dataset

The provided dataset contains a total of 157 subjects, later divided into training and validation sets. Each subject has a diagnostic label (82 HC, 59 MCI, 16 Dementia) for the classification task. For the regression task, MMSE scores are available for a subset of 69 subjects. To evaluate language communication and cognitive function, subjects completed three standardized English elicitation tasks commonly used in dementia research: the Cookie Theft Description (CTD), Phonemic Fluency Test (PFT), and Semantic Fluency Test (SFT).

For the CTD task, subjects describe the ‘‘Cookie Theft’’ picture, originally part of the Boston Diagnostic Aphasia Examination and later widely used for dementia detection. This task evaluates cognitive functions, including language comprehension and memory. In the PFT task, subjects list as many words as possible starting with the letter ‘‘P’’ within one minute, excluding proper nouns and country names, assessing verbal fluency and executive language functions. The SFT task requires subjects to name as many animals as possible in one minute, primarily evaluating naming skills and language comprehension to identify potential language impairments.

Each of the 157 subjects has an audio recording and a corresponding manual transcript for each of the elicitation tasks described above. The manual transcripts include annotations for speaker identities, external noises, and durations between utterances. To align the transcripts with the actual spoken content, we remove these annotations. For the test set, only the audio recordings are available, while manual transcripts and diagnostic labels are withheld.

The dataset presents a significant class imbalance, with a predominance of HC subjects and relatively few Dementia cases. While this mirrors real-world clinical distributions, it complicates the training of classification models and motivates our cascaded binary classification approach. Another challenge lies in the substantial overlap of MMSE scores between HC and Dementia groups, despite the expectation of clear separation. For instance, two Dementia subjects scored 29 out of 30, highlighting the subjectivity inherent in clinical assessments. Furthermore, MMSE scores are available for only 69 subjects,

adding complexity to the regression task.

4. Methods

This section outlines our methods for classification and regression tasks, covering data preprocessing, architectural details for both tasks (Figure 1), and the ensemble strategies applied.

4.1. Preprocessing for classification and regression

4.1.1. Denoising

Through manual inspection of the dataset, we identified some audio recordings with loud background noise, beeping sounds at the start and end, and occasional interviewer interventions. To address the first two issues, we applied the MP-SENet model [30], which denoises both magnitude and phase spectra in parallel. Since interviewer interventions can negatively impact dementia detection performance [31], we explored the use of an advanced speaker diarization model [32] to exclude the interviewer’s speech. However, initial experiments revealed frequent misclassifications, with substantial portions of the subject’s speech wrongly identified as the interviewer’s. Further comparisons between manual removal and retaining the interventions showed minimal impact on performance. Consequently, we opted to omit this step in subsequent study.

4.1.2. Fine-tuning ASR model

To enhance transcription accuracy for the test set, we fine-tune the advanced Automatic Speech Recognition (ASR) model, *Whisper-large-v3*¹, using the challenge dataset. The 157 subjects are split into ASR training and validation sets in a 3:1 ratio, stratified by diagnostic labels, incorporating data from all three elicitation tasks for fine-tuning. Given the extended duration of the audio recordings, we first apply forced alignment [33] between each audio file and its corresponding manual transcript. Based on the alignment timestamps, we segment the audio and transcripts into 30-second audio-transcript pairs for fine-tuning.

Fine-tuning is conducted using the *WhisperForConditionalGeneration* and *WhisperProcessor* modules from the *Transformers* Python library. The model checkpoint achieving the highest accuracy on the ASR validation set is selected for final transcription.

4.1.3. Pause encoding

We encode pauses between words in input transcripts to better capture disfluency patterns. Specifically, transcripts are force-

¹ <https://huggingface.co/openai/whisper-large-v3>

aligned with audio recordings using the same forced aligner [33], where pauses are marked as “SIL”. Pauses at the start and end of recordings are removed. Following [29], pauses are categorized as short (<0.5s), medium (0.5–2s), or long (>2s), and encoded using “,”, “.”, and “...”, respectively. Since original punctuation is removed, these symbols exclusively represent pause information.

4.2. Classification task

4.2.1. PLMs

Modern PLMs based on the Transformer architecture [34], such as BERT², RoBERTa³, ERNIE⁴, have proven capable of capturing a wide range of linguistic facts, including lexical knowledge, syntax, and semantics. We fine-tuned these three PLMs for dementia classification. Notably, we excluded speech-related features, as incorporating Pre-trained Speech Models (PSMs) such as Wav2Vec2 [35] or traditional handcrafted features did not yield noticeable performance gains.

4.2.2. Cascaded binary classification

Due to the significant data imbalance among diagnostic classes in the dataset and the inherent complexity of the three-class classification task, we propose a cascaded binary classification strategy to mitigate these challenges. Specifically, we replace the three-class classifier with two independently trained binary classifiers. As shown in Figure 1, during the training phase, we first relabel all MCI and Dementia samples in the training set as Non-HC. We then train a binary classifier to differentiate HC from Non-HC using the entire training set (i.e., HC and Non-HC samples). This approach simplifies training, as the challenge dataset provides an approximate HC:MCI:Dementia ratio of 5:4:1, resulting in a nearly balanced HC:Non-HC distribution of approximately 1:1. Additionally, the distinction between Healthy Controls (HC) and patients (Non-HC) is likely more pronounced than the differences among the three original classes (HC, MCI, and Dementia).

For the second classifier, we train a binary model exclusively on MCI and Dementia samples. While this reduces the number of available training samples (excluding HC), the classification task is potentially less challenging than the original three-class classification, as it focuses solely on distinguishing between MCI and Dementia.

During the testing phase, we first feed the preprocessed samples into an initial classifier to determine whether they belong to the HC or Non-HC. If a sample is classified as HC, it is directly assigned the HC label. Otherwise, it undergoes further processing by a second classifier to refine the classification into either MCI or Dementia. To construct these classifiers, we fine-tune PLMs. Specifically, for each preprocessed transcript, we input it into a PLM, such as BERT, and extract the 1024-dimensional embedding of the [CLS] token from the 24th layer. This embedding serves as input to a simple classification module comprising a linear layer with a Tanh activation function, a dropout layer, and a final linear layer with binary output.

4.3. Regression task

4.3.1. Feature extraction

We extract both acoustic and linguistic features from speech and text, respectively, to leverage complementary information

across modalities. For acoustic features, we extract 1024-dimensional embeddings from the 24th transformer layer of Wav2Vec2, computed as the mean of all token embeddings. Additionally, we design a 10-dimensional silence feature set comprising the number of silence occurrences per second, the ratio of silence to speech duration, and statistical descriptors (maximum, minimum, mean, and standard deviation) of both silence and speech durations. Silence segments are identified using a Recurrent Neural Network (RNN)-based Voice Activity Detection (VAD) model from pyannote⁵.

To capture low-level acoustic features, we employ the OpenSMILE toolkit⁶ to extract the eGeMAPS (88-dimensional) and ComParE (6373-dimensional) feature sets. For linguistic features, we obtain 1024-dimensional embeddings of the [CLS] token from the 24th transformer layer of RoBERTa. Consequently, we construct a total of five distinct feature sets.

4.3.2. Regression models

We use Support Vector Regression (SVR)⁷, XGBoost⁸, and CatBoost⁹, as existing research demonstrates their effectiveness in dementia detection [17].

4.4. Model ensemble

For both classification and regression tasks, we conduct experiments on each individual elicitation task. Each model undergoes 10 training runs with different random seeds to ensure robustness. In the classification task, for both Classifier 1 and Classifier 2, we employ majority voting across the output labels of three PLMs, three elicitation tasks, and 10 random seeds, resulting in a total of $3 \times 3 \times 10 = 90$ models contributing to the final decision.

For the regression task, we first compute the average predictions across the 10 random seeds. Subsequently, we select models with a Root Mean Square Error (RMSE) below 3, drawn from three regression models applied to five feature sets across three elicitation tasks ($3 \times 5 \times 3 = 45$ models in total), and further refine the final MMSE score by averaging their outputs.

5. Experiments

5.1. Experimental setup

For the classification task, we partitioned the 157 subjects into training and validation sets in a 3:1 ratio based on diagnostic labels. To enhance robustness, we conducted four different data splits and evaluated performance using the Macro F1 score. Similarly, for the regression task, 69 subjects were divided into training and validation sets in a 3:1 ratio, maintaining the distribution of MMSE scores. We also performed four different splits and assessed model performance using RMSE.

The fine-tuning of PLMs was conducted over 20 epochs with a batch size of 8 and a learning rate of 2×10^{-5} . We employed the AdamW optimizer and used cross-entropy loss for classification tasks. For regression models, optimal hyperparameters were determined via 4-fold cross-validation combined with a grid search strategy.

All experiments were conducted on NVIDIA RTX 4090 or A100 GPUs. Participants were allowed to submit up to three models for evaluation by the organizers.

⁵<https://github.com/pyannote/pyannote-audio>

⁶<https://github.com/audeering/opensmile-python>

⁷<https://scikit-learn.org/stable/modules/generated/sklearn.svm.SVR.html>

⁸<https://github.com/dmlc/xgboost>

⁹<https://github.com/catboost/catboost>

²<https://huggingface.co/google-bert/bert-large-uncased>

³<https://huggingface.co/FacebookAI/roberta-large>

⁴<https://huggingface.co/nghuyong/ernie-2.0-large-en>

Table 1: Macro F1 score (%) of the classification task.

Different splits	Validation set		Test set
	Cascaded binary	Direct three-class	
Split 1	59.9	55.3	—
Split 2	57.9	57.3	—
Split 3	63.6 ✓	59.9	45.1
Split 4	61.8	61.1	—
	Majority voting (4 splits)✓		57.8
	Majority voting (Retraining(4 splits))✓		58.6
Baseline			55.0

5.2. Results and analyses

5.2.1. ASR model performance

After fine-tuning, the Word Error Rate (WER) of Whisper on the ASR validation set decreased from the original 29.7% to 23.2%, demonstrating the effectiveness of the method.

5.2.2. Classification task

For the classification task, we conducted experiments by fine-tuning PLMs for direct three-class classification, serving as a baseline for comparison with the proposed cascaded binary classification approach. The three submitted models were selected based on the following criteria: (1) the model that achieved the highest validation performance among the four data splits; (2) an ensemble model leveraging majority voting across all models from the four splits; and (3) a model obtained by retraining all models using the entire dataset, followed by majority voting to aggregate their predictions.

Table 1 presents the performance comparison between the proposed cascaded binary classification approach and the direct three-class classification method, both employing model ensemble, on the validation set of each split. Additionally, it reports the performance of the selected models on the test set based on the predefined criteria. From the results, we observe the following key findings:

1) The proposed cascaded binary classification approach consistently outperforms direct three-class classification on the validation set, effectively simplifying multiclass classification and mitigating data imbalance across diagnostic classes. 2) On the test set, our method achieved an F1 score of 58.6%, exceeding the baseline (55%) by 3.6%, further confirming its effectiveness. 3) Although Split 3 showed strong validation performance, it suffered a significant drop on the test set, indicating potential overfitting. However, a multi-split ensemble strategy improved the test F1 score, enhancing model robustness. 4) Finally, retraining on the entire dataset yielded further performance gains, underscoring the advantages of full data utilization for model optimization.

5.2.3. Regression task

For the regression task, we explored fine-tuning RoBERTa and Wav2Vec2 by appending a single-output linear layer to each model and averaging their outputs to predict the MMSE score. This approach was compared against our proposed method, which extracts multimodal features and integrates them with machine learning regression models. The selection criteria for the three submitted models were aligned with those used in the classification task: (1) selecting the model from the split that achieved the highest validation performance among the four splits; (2) averaging the scores from all models across the four splits; and (3) retraining all models using the entire dataset and

Table 2: RMSE of the regression task.

Different splits	Validation set		Test set
	Multimodal	Fine-tuning RoB+W2V	
Split 1	2.88	2.95	—
Split 2	2.56 ✓	2.83	2.88
Split 3	2.66	2.98	—
Split 4	2.92	2.97	—
	Score averaging (4 splits)✓		2.87
	Score averaging (Retraining(4 splits))✓		2.89
Baseline			2.98

subsequently averaging their predicted scores.

Table 2 presents the performance of the proposed multimodal feature fusion method alongside fine-tuned RoBERTa and Wav2Vec2 models with ensemble learning, evaluated on the validation set of each split and the test set for selected models. Key observations include:

1) The proposed multimodal feature fusion method consistently outperforms fine-tuned pre-trained models in terms of RMSE on the validation set. Fine-tuned models often predict MMSE scores near the mean of the training data, likely due to the dataset’s limited size (69 subjects). The complexity and larger parameter size of pre-trained models make it challenging to capture dementia-specific features from such a small dataset. This highlights the advantage of multimodal feature extraction combined with simpler machine learning models for regression tasks. 2) All three submitted models surpassed the baseline (RMSE = 2.98), with the best achieving an RMSE of 2.87, an 0.11 improvement, demonstrating the effectiveness of our proposed approach for MMSE score prediction. 3) Model ensembles integrating diverse acoustic and linguistic features consistently achieved superior performance across all splits, underscoring the complementary nature of these modalities and validating the multimodal fusion framework. 4) Despite these gains, the model exhibited mild overfitting on the validation set, indicating the need for stronger generalization strategies. Additionally, dataset-specific challenges, such as overlapping MMSE scores across diagnostic categories and class imbalance (Section 3), further complicate model optimization.

6. Conclusions

In this paper, we present the results and analyses of our submission to this year’s PROCESS challenge. For the classification task, we introduce a cascaded binary classification strategy to tackle the underexplored three-class problem, with particular emphasis on the early diagnosis of MCI. This approach simplifies multi-class classification and mitigates data imbalance across diagnostic classes. For the regression task, we propose an enhanced multimodal framework that integrates diverse acoustic and linguistic features, leveraging their complementary strengths for improved performance. Experimental results on the test set demonstrate that our methods outperform the provided baseline in both tasks, underscoring the potential of our approach for advancing early dementia detection. Future work will focus on developing more robust and generalizable automated dementia detection models for real-world applications.

7. Acknowledgements

This work was partially supported by the National Social Science Foundation of China (Grant No. 23AYY012), and by the Supercomputing Center of the University of Science and Technology of China.

8. References

- [1] J. C. Beck, D. F. Benson, A. B. Scheibel, J. E. Spar, and L. Z. Rubenstein, "Dementia in the elderly: the silent epidemic," *Annals of Internal Medicine*, vol. 97, no. 2, pp. 231–241, 1982.
- [2] C. R. Jack Jr, M. A. Bernstein, N. C. Fox, P. Thompson, G. Alexander, D. Harvey, B. Borowski, P. J. Britson, J. L. Whitwell, and C. Ward, "The alzheimer's disease neuroimaging initiative (adni): Mri methods," *Journal of Magnetic Resonance Imaging: An Official Journal of the International Society for Magnetic Resonance in Medicine*, vol. 27, no. 4, pp. 685–691, 2008.
- [3] J. Samper-González, N. Burgos, S. Bottani, S. Fontanella, P. Lu, A. Marcoux, A. Routier, J. Guillon, M. Bacci, and J. Wen, "Reproducible evaluation of classification methods in alzheimer's disease: Framework and application to mri and pet data," *NeuroImage*, vol. 183, pp. 504–521, 2018.
- [4] Y. Li, Z. Zhang, J. Han, P. Bell, and C. Lai, "Semi-supervised cognitive state classification from speech with multi-view pseudo-labeling," *ICASSP*, 2025.
- [5] J. L. Cummings and D. F. Benson, "Speech and language alterations in dementia syndromes," in *Brain organization of language and cognitive processes*. Springer, 1989, pp. 107–120.
- [6] A. Yeung, A. Iaboni, E. Rochon, M. Lavoie, C. Santiago, M. Yancheva, J. Novikova, M. Xu, J. Robin, and L. D. Kaufman, "Correlating natural language processing and automated speech analysis with clinician assessment to quantify speech-language changes in mild cognitive impairment and alzheimer's dementia," *Alzheimer's research & therapy*, vol. 13, no. 1, p. 109, 2021.
- [7] T. Talkar, S. Charles, C. Krantsevich, and K. Kawabata, "Detection of cognitive impairment and alzheimer's disease using a speech- and language-based protocol," in *Proc. Interspeech*, 2024, pp. 3025–3029.
- [8] Y.-L. Liu, R. Feng, J.-H. Yuan, and Z.-H. Ling, "Clever hans effect found in automatic detection of alzheimer's disease through speech," in *Proc. Interspeech*, 2024, pp. 2435–2439.
- [9] Y. Li, C. Lai, D. Lala, K. Inoue, and T. Kawahara, "Alzheimer's dementia detection through spontaneous dialogue with proactive robotic listeners," in *2022 17th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE, 2022, pp. 875–879.
- [10] Y. Pan, B. Mirheidari, J. M. Harris, J. C. Thompson, M. Jones, and J. S. Snowden, "Using the outputs of different automatic speech recognition paradigms for acoustic-and BERT-based Alzheimer's dementia detection through spontaneous speech," in *Proc. Interspeech*, 2021, pp. 3810–3814.
- [11] Y.-L. Liu, R. Feng, Y.-X. Lu, J.-X. Chen, Y. Ai, J.-H. Yuan, and Z.-H. Ling, "Can automated speech recognition errors provide valuable clues for alzheimer's disease detection?" in *Proc. ICASSP*, 2025, pp. 1–5.
- [12] F. Tao, B. Mirheidari, M. Pahar, S. Young, Y. Xiao, H. Elghazaly, F. Peters, C. Illingworth, D. Braun, and R. O'Malley, "Early dementia detection using multiple spontaneous speech prompts: The process challenge," *arXiv preprint arXiv:2412.15230*, 2024.
- [13] S. Luz, F. Haider, S. de la Fuente Garcia, D. Fromm, and B. MacWhinney, "Alzheimer's dementia recognition through spontaneous speech: The ADReSS challenge," in *Proc. Interspeech*, 2020, pp. 2172–2176.
- [14] S. Luz, F. Haider, S. de la Fuente, D. Fromm, and B. MacWhinney, "Detecting cognitive decline using speech only: The ADReSSo challenge," *arXiv preprint arXiv:2104.09356*, 2021.
- [15] S. Luz, F. Haider, D. Fromm, I. Lazarou, I. Kompatsiaris, and B. MacWhinney, "Multilingual Alzheimer's dementia recognition through spontaneous speech: a signal processing grand challenge," in *Proc. ICASSP*, 2023, pp. 1–2.
- [16] F. Braun, S. P. Bayerl, F. Hönig, H. Lehfeld, T. Hillemacher, T. Bocklet, and K. Riedhammer, "Infusing acoustic pause context into text-based dementia assessment," in *Proc. Interspeech*, 2024, pp. 1980–1984.
- [17] K. Mei, X. Ding, Y. Liu, Z. Guo, F. Xu, X. Li, T. Naren, J. Yuan, and Z. Ling, "The USTC system for ADReSS-M challenge," in *Proc. ICASSP*, 2023, pp. 1–2.
- [18] M. S. S. Syed, Z. S. Syed, M. Lech, and E. Pirogova, "Automated screening for Alzheimer's dementia through spontaneous speech," in *Proc. Interspeech*, 2020, pp. 2222–2226.
- [19] Z. S. Syed, M. S. S. Syed, M. Lech, and E. Pirogova, "Tackling the addresso challenge 2021: The muet-rmit system for alzheimer's dementia recognition from spontaneous speech," in *Interspeech*, 2021, pp. 3815–3819.
- [20] M. Rohanian, J. Hough, and M. Purver, "Alzheimer's dementia recognition using acoustic, lexical, disfluency and speech pause features robust to noisy inputs," *arXiv preprint arXiv:2106.15684*, 2021.
- [21] M. Martinc and S. Pollak, "Tackling the address challenge: A multimodal approach to the automated recognition of alzheimer's dementia," in *Interspeech*, 2020, pp. 2157–2161.
- [22] Y. Qiao, X. Yin, D. Wiechmann, and E. Kerz, "Alzheimer's disease detection from spontaneous speech through combining linguistic complexity and (dis) fluency features with pretrained language models," *arXiv preprint arXiv:2106.08689*, 2021.
- [23] L. B. d. Santos, E. A. Corrêa Jr, O. N. Oliveira Jr, D. R. Amancio, L. L. Mansur, and S. M. Alufio, "Enriching complex networks with word embeddings for detecting mild cognitive impairment from speech transcripts," *arXiv preprint arXiv:1704.08088*, 2017.
- [24] J. Li, J. Yu, Z. Ye, S. Wong, M. Mak, B. Mak, X. Liu, and H. Meng, "A comparative study of acoustic and linguistic features classification for Alzheimer's disease detection," in *Proc. ICASSP*, 2021, pp. 6423–6427.
- [25] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "BERT: Pre-training of deep bidirectional transformers for language understanding," *arXiv preprint arXiv:1810.04805*, 2018.
- [26] Y. Liu, M. Ott, N. Goyal, J. Du, M. Joshi, and D. Chen, "RoBERTa: A robustly optimized BERT pretraining approach," *arXiv preprint arXiv:1907.11692*, 2019.
- [27] Y. Sun, S. Wang, Y. Li, S. Feng, H. Tian, H. Wu, and H. Wang, "Ernie 2.0: A continual pre-training framework for language understanding," in *Proc. AAAI*, 2020, pp. 8968–8975.
- [28] Y. Wang, T. Wang, Z. Ye, L. Meng, S. Hu, X. Wu, X. Liu, and H. Meng, "Exploring linguistic feature and model combination for speech recognition based automatic ad detection," in *Proc. Interspeech*, 2022, pp. 3328–3332.
- [29] J. Yuan, Y. Bian, X. Cai, J. Huang, Z. Ye, and K. Church, "Disfluencies and fine-tuning pre-trained language models for detection of Alzheimer's disease," in *Proc. Interspeech*, 2020, pp. 2162–2166.
- [30] Y.-X. Lu, Y. Ai, and Z.-H. Ling, "Mp-senet: A speech enhancement model with parallel denoising of magnitude and phase spectra," in *Proc. Interspeech*, 2023, pp. 3834–3838.
- [31] Y.-L. Liu, R. Feng, J.-H. Yuan, and Z.-H. Ling, "Leveraging prompt learning and pause encoding for alzheimer's disease detection," in *Proc. ISCSLP*, 2024, pp. 486–490.
- [32] A. Plaquet and H. Bredin, "Powerset multi-class cross entropy loss for neural speaker diarization," in *Proc. Interspeech*, 2023, pp. 3222–3226.
- [33] J. Yuan and M. Liberman, "Speaker identification on the scotus corpus," *Journal of the Acoustical Society of America*, vol. 123, no. 5, p. 3878, 2008.
- [34] A. Vaswani, "Attention is all you need," *Advances in Neural Information Processing Systems*, 2017.
- [35] A. Baevski, Y. Zhou, A. Mohamed, and M. Auli, "wav2vec 2.0: A framework for self-supervised learning of speech representations," in *Proc. NeurIPS*, 2020, pp. 12 449–12 460.