

Not All Errors Are Equal: Investigation of Speech Recognition Errors in Alzheimer’s Disease Detection

Jiawen Kang, Junan Li, Jinchao Li, Xixin Wu, Helen Meng

The Chinese University of Hong Kong, Hong Kong SAR

{jwkang, jli, jcli, wuxx, hmmeng}@se.cuhk.edu.hk

Abstract

Automatic Speech Recognition (ASR) plays an important role in speech-based automatic detection of Alzheimer’s disease (AD). However, recognition errors could propagate downstream, potentially impacting the detection decisions. Recent studies have revealed a non-linear relationship between word error rates (WER) and AD detection performance, where ASR transcriptions with notable errors could still yield AD detection accuracy equivalent to that based on manual transcriptions. This work presents a series of analyses to explore the effect of ASR transcription errors in BERT-based AD detection systems. Our investigation reveals that not all ASR errors contribute equally to detection performance. Certain words, such as stopwords, despite constituting a large proportion of errors, are shown to play a limited role in distinguishing AD. In contrast, the keywords related to diagnosis tasks exhibit significantly greater importance relative to other words. These findings provide insights into the interplay between ASR errors and the downstream detection model.

Index Terms: Alzheimer’s disease, speech recognition, pre-trained language model, spoken language processing

1. Introduction

Alzheimer’s disease (AD) is a neurodegenerative disorder characterized by progressive cognitive impairment, including deterioration in memory, attention, and executive function. Due to the irreversible progression of AD pathology [1], early detection and diagnosis play a pivotal role in facilitating timely intervention and management, conventionally relying on in-person clinical assessments [2, 3]. With recent progress in spoken language technology, speech-based automatic AD detection has emerged as a promising area due to its potential for more cost-effective and scalable AD screening [4, 5].

Automatic detection of Alzheimer’s Disease (AD) from speech has attracted increasing attention in recent years. Early research attempted to identify AD by utilizing handcrafted acoustic and linguistic features [6, 7, 8, 9]. For example, Winer and Frankenberg et al. [8, 9] demonstrated that parts-of-speech (POS), word categories, and pause features are highly related to AD. More recently, deep embedding features from pre-trained models have been widely explored for AD detection, leveraging either speech-based models [10, 11, 12] or text-based models [13, 14, 15]. Among these models, pre-trained language models (PLMs) such as BERT [16] and RoBERTa [17], have shown remarkable performance by leveraging their strong ability to capture rich linguistic and semantic patterns. These PLMs can be used either as feature extractors [13, 18, 19] or directly fine-tuned on the AD detection task [14]. Moreover, researchers have made further enhancements to PLMs [20, 21, 22], pushing the state-of-the-art detection accuracy to new heights.

While PLMs have achieved impressive performance in AD detection, the practical applications of these approaches largely

rely on Automatic Speech Recognition (ASR) front-ends to transcribe speech into text, as transcription errors could lead to significant bias in PLM embeddings and alter the decisions of downstream AD detection models. Many works have investigated tailored ASR systems for AD or other disordered speech. Early works explored incorporating language models [23] and Maximum A Posteriori (MAP) adaptation [24] for ASR in AD detection. Recent efforts have investigated novel approaches such as data augmentation [25, 26, 27], domain adaptation [28] and neural architectural search [29]. Nevertheless, due to the difficulty in collecting in-domain data, there is still a substantial accuracy gap between ASR for disordered speech and normal speech.

On the other hand, many studies have shown a non-linear relationship between ASR Word Error Rate (WER) and AD detection accuracies. For instance, ASR transcription with notable errors (WER ~30%) could yield equivalent AD detection accuracy as manual transcription does [19, 30], while ASR models with lower WER might not necessarily lead to better AD detection results [31, 29]. Few studies have investigated the in-depth relationship between ASR errors and AD detection. [23] might be the first work revealing that the performances of ASR and AD detection systems are weakly correlated. Balagopalan et al. [32] demonstrated that deletion error is effective in the AD detection model using a handcrafted feature set. Recently, Li et al. [31] showed that ASR errors could serve as helpful cues for fine-tuned PLMs to identify speech with AD. These findings suggest that certain ASR error words might be trivial and harmless to the downstream AD detection task, while others could have a strong effect.

This work aims to explore the impact of ASR errors on Alzheimer’s Disease (AD) detection using a BERT-based model. We first implement a subject ASR and AD detection systems which achieves the same accuracy when evaluated with ASR and manual transcriptions. We then conducted a thorough analysis based on the ASR errors. It reveals that not all ASR errors are equally detrimental to downstream AD detection. Certain words, such as stopwords, constitute a large proportion of errors but are less significant concerning the picture description tasks performed in the ADReSS data used in this study. In contrast, task-related keywords occupy only 9% of all ASR errors but have proven to play a pivotal role compared to other words. These findings provide insights into the interplay between ASR errors and the downstream detection model, benefiting the development of ASR-robust AD detection systems.

2. Approach

2.1. Data

The data used in this work comes from the Alzheimer’s Dementia Recognition Through Spontaneous Speech (ADReSS) Challenge 2020 corpus [33]. This challenge selects a sub-task of Pitt Corpus in the DementiaBank database [34], which requires all

the participants to describe the Cookie Theft picture [35]. The ADReSS corpus consists of 156 different English speakers’ audio samples with corresponding transcripts. Among them, 78 of the speakers are healthy controls (35 male, 43 female) while the rest are with AD (35 male, 43 female). The corpus is divided into a standard train (108 speakers, about 2 hours) and test (48 speakers, about 1 hour) sets with balanced distributions of age, gender, and disease conditions.

2.2. Automatic Speech Recognition Models

We employed a tailor-made ASR system based on the Time Delay Neural Network (TDNN). This model was trained with 1000-hour Librispeech corpus and adapted on 59-hour Pitt corpus [34]. Additionally, data augmentation, speaker adaptation, and a Transformer language model were adopted to enhance the performance on the in-domain speech. As a result, the employed system has shown state-of-the-art level performance on the ADReSS corpus. More details can be found in [25].

2.3. AD Detection Models

We formulate AD detection as a binary classification problem, i.e. classifying participants as healthy controls or individuals with AD. The development of the detection system follows work [19], which has shown superior detection performance on the ADReSS corpus. This system utilizes the BERT model [16] to generate embeddings from the automatically transcribed text. Specifically, we take the embedding corresponding to the [CLS] token for each speaker throughout the entire transcript. This process yields a feature vector with a dimensionality of 768, which serves as the input representation for the following analysis. To avoid the over-fitting problem, we use the Principal Component Analysis (PCA) method to reduce the dimension of the embeddings to 108, which is the size of the training set. The compressed embeddings are then fed into a Support Vector Machine (SVM) with linear kernel and regularization parameter $C = 1$ for classification. Both PCA and SVM models are fitted with manual transcription. As for evaluation, we use accuracy scores as the main criterion, where precision, recall, and F1 scores are also used for analysis.

3. AD Detection with ASR Transcriptions

The non-linear dependency between ASR WER and AD detection accuracy has been observed in multiple studies [19, 29, 30, 31]. To look into this observation, we first set up a subject system using the above-mentioned settings. Table 1 shows the classification results of the subject detection system. Note that the detection model was fitted with manual transcription, and ASR transcriptions were only used in the test stage. We can see that even though the ASR system has a WER of 33.9%, it still achieves an equivalent detection accuracy of 88% compared to manual transcriptions.

Table 1: *AD detection performance using Manual and ASR transcription. “Trans.” stands for transcriptions.*

Trans.	WER	Accuracy	Precision	Recall	F1	AUC
Manual	—	0.88	0.91	0.83	0.87	0.88
ASR	33.9%	0.88	0.82	0.96	0.88	0.88

Table 2 further breaks down prediction results into confusion matrices. It reflects that although two types of transcriptions achieve the same accuracy, prediction with ASR transcriptions tends *more false positive errors*. Compared to manual transcription, using ASR transcriptions results in 7 additional

Table 2: *Confusion matrix of AD detection results. “TP, TN, FP, FN” stands for true positive, true negative, false positive and false negative, respectively.*

Transcriptions	TP	TN	FP	FN	Total
Manual	20	22	2	4	48
ASR	23	19	5	1	48

positive predictions (3 cases changed from true negatives to false positives, 4 cases changed from false negatives to true positives), and only one more negative prediction (a case changed from false negative to true positive). This suggests that transcriptions with ASR errors might be more likely to be classified as originating from AD patients. However, there are 40/48 ASR-robust cases where manual and ASR transcriptions yield the same classification results, despite the high word error rate of 30%. This raises questions about how these errors influence downstream decision-making.




Figure 1: *The top 20 error word distribution of the ASR system.*

4. Analysis of speech recognition errors

Based on the detection system in Section 3, this section characterizes the composition of ASR errors and investigates why a high word error rate did not significantly degrade AD detection performance. We will start with error distribution and then examine certain types of errors using the test set transcriptions.

4.1. Basic analysis

The classification results show that there are 40 out of 48 cases (ASR-robust cases) where ASR and manual transcriptions yielded identical predictions. To look into details, we first draw the distribution of the top 20 ASR errors. Given that these errors did not lead to prediction changes, we assume that frequent errors play a relatively marginal role in AD detection. As shown in Figure 1, we can first observe these errors follow a long-tail distribution, with a small proportion of words constituting a large proportion of errors. Moreover, we noticed that most of these top word errors are stopwords, which have been considered to carry less semantic weight and are not directly related to the picture description task used to diagnose AD patients.

In addition to overall errors, we conducted a thorough case study to observe the individual patterns of ASR errors. Particularly, we generated alignment maps for each participant to visualize the compositions of ASR transcriptions and their discrepancy with manual transcription. As showcased in Figure 2, the alignment map compares manual and ASR transcriptions word-by-word, where correct words are marked as squares and errors are marked as crosses. Following the findings in the previous section, we highlight the low-semantic stopwords¹ with blue colors. In comparison, we additionally highlight the relatively high-semantic task-related keywords with red colors. As the diagnosis of AD in the ADReSS corpus is based on a picture description task, we hypothesize that content words related to the elements in the target picture may play a role in distinguishing

¹Following NLTK [36] (v3.8.1) Stopwords List




Figure 2: Samples of alignment map for participants *s170* (a) and *s179* (b). Squares represent correctly transcribed words, while crosses ('x') indicate ASR errors. Blue color for stopwords, red for keywords, and gray for other words.

Table 3: List of the task-related keywords

Task-related Keywords
boy, girl, two, children, cookie, jar, brother, sister, steal, cupboard, door, finger, women, lady, dry, washing, dish, water, overflow, floor, kitchen, window, curtain, mother, plate, stand, sink, running, fall, outside, tree, house, look, open, hand, glass

AD patients from healthy controls. Table 3 lists 39 task-related keywords used in this experiment, drawing on study [31].

Going through the alignment maps of ASR-robust cases, we have the following findings. First, stopwords and keywords covered most of the words in transcriptions, thus investigations of these words considered a majority of content. Second, healthy participants generally could mention two times of keywords than participants with AD, which was reflected by more red markers in their alignment maps. Additionally, aligning with the error distribution, a large proportion of ASR errors are on stopwords, which means *most keywords are well preserved in ASR transcriptions*. These observations suggest that ASR errors are dominated by low-semantic stopwords, while keywords related to AD diagnosis are relatively well recognized. This may explain why these errors have a limited impact on the final AD classification. The following part will present our detailed investigation of stopwords and keywords.

4.2. Analysis of stopwords

A series of experiments was conducted to examine the role of stopwords in AD detection. First, Table 4 presents a detailed breakdown of stopwords and non-stopwords in ASR errors from the ASR-robust cases. It shows that a small proportion of stopword types (24%) constitute the majority of errors (60%), while the remaining 76% types of words constitute only 40% of errors. This aligns with the observation in Figure 1 and 2. Given these errors did not affect AD classification, we hypothesize that when using BERT features, the misrecognized stopwords error may not significantly propagate to the downstream detection decision. Meanwhile, we note that the non-stopword errors also show an insignificant impact on classification results, suggesting that words other than stopwords could also play a marginal role in distinguishing AD. We will discuss this further in the later part of this paper.

Words ablation experiments To verify the impact of stopwords on downstream classification, we conduct a word ablation experiment to probe the variation of BERT embeddings when stopwords are misrecognized. Starting from manual transcriptions, we randomly removed or substituted 10% to 100% of stopwords in 10% increments and observed the correspond-

Table 4: Word error rates (WER) and counts (ratio %) of stopwords and non-stopwords in manual transcription (Counts) and ASR errors (Errors Counts). “tokens” represents the count of words and “types” represents the count of distinct words.

	WER	Counts		Errors Counts	
		tokens	types	tokens	types
Stopwords	32.4%	2574(58%)	95(18%)	834(60%)	87(24%)
Non-stopwords	29.8%	1897(42%)	447(82%)	566(40%)	281(76%)

ing BERT embeddings. For substitution, we replaced random stopwords with an equal number of non-stopwords randomly sampled from all manual transcriptions. We then projected the high-dimensional BERT embeddings of the edited transcriptions onto a 2-dimensional subspace orthogonal to the original SVM decision hyperplane. The two sides of the decision boundary were filled with white (healthy) and gray (AD) colors to demonstrate the relationship between BERT embeddings and SVM decisions. Figure 3 shows the resulting visualizations from an example participant. When gradually removing or substituting stopwords (Figure 3 (a) & (c)), BERT embeddings show a shift towards the decision boundary. This suggests that stopwords may have an effect on AD detection. However, the edited transcription still falls on the negative (recognized as healthy) side until all stopwords are removed/substituted. In contrast, when gradually removing/substituting the same number of non-stopwords (Figure 3 (b) & (d)), the embeddings shift across the decision boundary with a more direct trend. Moreover, in this case, substitutions result in a larger embedding variance compared to removal. We attribute this to the randomness in the substitution operation, as random words were sampled for replacement.

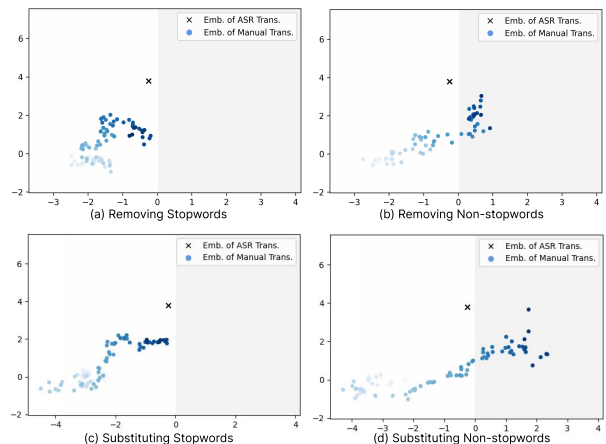


Figure 3: BERT embedding variations as stopwords are incrementally removed/substituted from manual transcriptions (participant *s172*). Light to dark blue: fewer to more stopwords removed/substituted. 'x': ASR transcript. White/gray: health/AD decision regions.

Furthermore, to reduce randomness and quantify this experiment across participants, we calculate the *hyperplane offset* for each participant and each removal/substitution. This offset represents a signed distance from BERT embedding to the SVM hyperplane.

$$d = (w^T x + b) / \|w\| \quad (1)$$

The above equation presents calculation details where w and b are the norm vector and interception of SVM, x is the input vector, and $\|\cdot\|$ stands for L2 norm. The resulting d reflects




Figure 4: Average hyperplane offset for transcription embeddings as stopwords (a) and keywords (b) are incrementally removed or substituted. Edit ratio represents the percentage of words removed/substituted.

how embeddings move toward or against the decision boundary. We averaged d across participants to observe the average embedding movement when transcriptions are edited. In Figure 4 (a), we present the resulting offset curves, using blue and red background colors to indicate healthy and AD decision regions, respectively. Similar to the embedding visualization, when removing or substituting stopwords (solid lines), the average offset values gradually approach 0 from negative values but do not cross the decision boundary. Suggesting that the edit of stopwords leads to embeddings shift towards the AD decision region, but not enough to change the classification outcome. In contrast, editing non-stopwords easily leads to set values reaching 0, indicating a likely change in the classification result. This difference implies distinct roles of stopwords and non-stopwords in the representations, with non-stopwords carrying more semantic content that influences the AD detection task.

4.3. Analysis of task-related keywords

In comparison to stopwords, this part investigates keywords related to elements in the Cookie Thief picture used for AD diagnosis. Table 5 breakdown keyword and non-keyword errors of ASR-robust cases. We highlight that the WER for keywords is only 14.3%, which is much lower than the overall WER of 33.9%. The high accuracy in keywords is understandable, as task-related keywords are frequently mentioned words in the training set, with only 60^2 distinct keywords in total. Compared to stopwords, most keywords have more syllables, making them easier to recognize. Furthermore, keyword errors account for just 9% of all ASR errors. This aligns with the observations using the alignment maps in Section 4.1.

Words ablation experiments We conducted word ablation experiments using the same settings as in the stopword investigation, replacing stopwords with keywords. The visualizations of BERT embeddings are exemplified in Figure 5. It is not surprising that the results of editing keywords and stopwords are opposite. Removing or substituting keywords in transcriptions causes BERT embeddings to shift across the SVM decision boundary, whereas similar manipulations of non-keywords do not produce this effect. Figure 4 (b) depicted the hyperplane offset experiment. There is also a clear divergence when editing keywords and non-keywords. Opposite to stopwords, removing or substituting keywords shows more significant changes in embedding offset. This divergence indicates *keywords may play a distinctly meaningful role compared to other words*. Some of the 88% non-keywords may also be non-stopwords, which could explain the insignificant non-stopwords in Table 4. No-

²Including inflections of words in Table 3

Table 5: Word error rates (WER) and counts (ratio %) of keywords and non-keywords in manual transcription (Counts) and ASR errors (Errors Counts). “tokens” represents the count of words and “types” represents the count of distinct words.

	WER	Counts		Errors Counts	
		tokens	types	tokens	types
Keywords	14.3%	909(20%)	60(11%)	130(9%)	46(12%)
Non-keywords	35.6%	3562(80%)	482(89%)	1270(91%)	322(88%)

tably, editing keywords didn’t result in embedding shifts as effectively as expected, i.e., the averaged offset started to be positive only when 70% of the keywords were removed. We interpret this as an indication that there exists information above the word level that can be helpful to AD classification. E.g., sentence length and the interplay between words within the transcriptions.

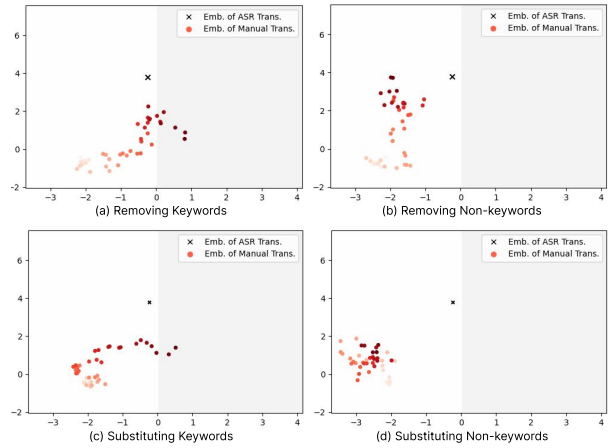


Figure 5: BERT embedding variations as keywords are incrementally removed/substituted from manual transcriptions (participant s172). Light to dark blue: fewer to more keywords removed/substituted. ‘x’: ASR transcript. White/gray: health/AD decision regions.

5. Conclusions

We performed a series of analyses to investigate the effect of ASR transcription errors in BERT-based Alzheimer’s Disease (AD) detection. Specifically, we explored why transcriptions with notable word error rates could yield detection accuracy equivalent to that of manual transcriptions. We have shown that not all ASR errors are equally detrimental to downstream AD detection. Certain word types, particularly stopwords, constitute a large proportion of errors while carrying limited information for AD discrimination than others. The task-related keywords were demonstrated to be significant to AD classification while occupying only 9% of all word errors. These findings provide insights into understanding the influence of ASR errors on downstream AD detection and facilitate the development of robust AD detection systems. While this work focused on ASR at the word level, future work could explore their influence on other linguistic structures such as syntactic dependencies and discourse coherence.

6. Acknowledgements

This work is supported by the HKSARG Research Grants Council’s Theme-based Research Grant Scheme (Project No. T45- 407/19N) and the CUHK Stanley Ho Big Data Decision Research Centre.

7. References

- [1] C. Lynch, “World alzheimer report 2019: Attitudes to dementia, a global survey: Public health: Engaging people in adrd research,” *Alzheimer’s & Dementia*, 2020.
- [2] Z. S. Nasreddine, N. A. Phillips, V. Bédirian, S. Charbonneau, V. Whitehead, I. Collin, J. L. Cummings, and H. Chertkow, “The montreal cognitive assessment, MoCA: a brief screening tool for mild cognitive impairment,” *Journal of the American Geriatrics Society*, vol. 53, no. 4, pp. 695–699, 2005.
- [3] J. Reilly, M. Losh, U. Bellugi, and B. Wulfeck, ““frog, where are you?” narratives in children with specific language impairment, early focal brain injury, and williams syndrome,” *Brain and language*, vol. 88, no. 2, pp. 229–247, 2004.
- [4] I. Martínez-Nicolás, T. E. Llorente, F. Martínez-Sánchez, and J. J. G. Meilán, “Ten years of research on automatic voice and speech analysis of people with alzheimer’s disease and mild cognitive impairment: a systematic review article,” *Frontiers in Psychology*, vol. 12, p. 620251, 2021.
- [5] H. Meng, B. Mak, M.-W. Mak, H. Fung, X. Gong, T. Kwok, X. Liu, V. Mok, P. Wong, J. Woo *et al.*, “Integrated and enhanced pipeline system to support spoken language analytics for screening neurocognitive disorders,” *Interspeech*, 2023.
- [6] T. Alhanai, R. Au, and J. Glass, “Spoken language biomarkers for detecting cognitive impairment,” 2017.
- [7] K. C. Fraser, J. A. Meltzer, and F. Rudzicz, “Linguistic features identify alzheimer’s disease in narrative speech,” *Journal of Alzheimer’s Disease*, 2016.
- [8] J. Weiner, C. Frankenberg, J. Schröder, and T. Schultz, “Speech reveals future risk of developing dementia: Predictive dementia screening from biographic interviews,” in *ASRU*. IEEE, 2019.
- [9] C. Frankenberg, J. Weiner, M. Knebel, A. Abulimiti, P. Toro, C. J. Herold, T. Schultz, and J. Schröder, “Verbal fluency in normal aging and cognitive decline: Results of a longitudinal study,” *Computer Speech & Language*, 2021.
- [10] J. Koo, J. H. Lee, J. Pyo, Y. Jo, and K. Lee, “Exploiting multi-modal features from pre-trained networks for alzheimer’s dementia recognition,” in *INTERSPEECH*, 2020.
- [11] A. Balagopalan, B. Eyre, J. Robin, F. Rudzicz, and J. Novikova, “Comparing pre-trained and feature-based models for prediction of alzheimer’s disease based on speech,” *Frontiers in aging neuroscience*, 2021.
- [12] Z. S. Syed, M. S. S. Syed, M. Lech, and E. Pirogova, “Automated recognition of alzheimer’s dementia using bag-of-deep-features and model ensembling,” *IEEE Access*, 2021.
- [13] A. Balagopalan, B. Eyre, F. Rudzicz, and J. Novikova, “To bert or not to bert: comparing speech and language-based approaches for alzheimer’s disease detection,” *arXiv preprint arXiv:2008.01551*, 2020.
- [14] 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 *INTERSPEECH*, 2020.
- [15] M. Martinc, F. Haider, S. Pollak, and S. Luz, “Temporal integration of text transcripts and acoustic features for alzheimer’s diagnosis based on spontaneous speech,” *Frontiers in Aging Neuroscience*, 2021.
- [16] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, “BERT: Pre-training of deep bidirectional transformers for language understanding,” 2019.
- [17] Y. Liu, M. Ott, N. Goyal, J. Du, M. Joshi, D. Chen, O. Levy, M. Lewis, L. Zettlemoyer, and V. Stoyanov, “Roberta: A robustly optimized bert pretraining approach,” 2019.
- [18] M. S. S. Syed, Z. S. Syed, M. Lech, and E. Pirogova, “Automated screening for alzheimer’s dementia through spontaneous speech,” in *Interspeech*, vol. 2020, 2020, pp. 2222–6.
- [19] 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 *ICASSP*. IEEE, 2021, pp. 6423–6427.
- [20] J. Li, K. Song, J. Li, B. Zheng, D. Li, X. Wu, X. Liu, and H. Meng, “Leveraging pretrained representations with task-related keywords for alzheimer’s disease detection,” 2023.
- [21] 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,” *INTER-SPEECH*, 2022.
- [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,” in *INTERSPEECH*, 2021.
- [23] L. Zhou, K. C. Fraser, and F. Rudzicz, “Speech recognition in alzheimer’s disease and in its assessment,” in *Interspeech*, vol. 2016, 2016, pp. 1948–1952.
- [24] B. Mirheidari, D. Blackburn, T. Walker, M. Reuber, and H. Christensen, “Dementia detection using automatic analysis of conversations,” *Computer Speech & Language*, vol. 53, pp. 65–79, 2019.
- [25] Z. Ye, S. Hu, J. Li, X. Xie, M. Geng, J. Yu, J. Xu, B. Xue, S. Liu, X. Liu *et al.*, “Development of the cuhk elderly speech recognition system for neurocognitive disorder detection using the dementia-bank corpus,” in *ICASSP*. IEEE, 2021, pp. 6433–6437.
- [26] M. Geng, X. Xie, S. Liu, J. Yu, S. Hu, X. Liu, and H. Meng, “Investigation of data augmentation techniques for disordered speech recognition,” *arXiv preprint arXiv:2201.05562*, 2022.
- [27] Z. Jin, M. Geng, J. Deng, T. Wang, S. Hu, G. Li, and X. Liu, “Personalized adversarial data augmentation for dysarthric and elderly speech recognition,” *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 2023.
- [28] J. Deng, F. R. Gutierrez, S. Hu, M. Geng, X. Xie, Z. Ye, S. Liu, J. Yu, X. Liu, and H. Meng, “Bayesian parametric and architectural domain adaptation of lf-mmi trained tdnn for elderly and dysarthric speech recognition,” in *Interspeech*, 2021, pp. 4818–4822.
- [29] T. Wang, J. Deng, M. Geng, Z. Ye, S. Hu, Y. Wang, M. Cui, Z. Jin, X. Liu, and H. Meng, “Conformer based elderly speech recognition system for alzheimer’s disease detection,” *arXiv preprint arXiv:2206.13232*, 2022.
- [30] C. Li, T. Cohen, and S. Pakhomov, “The far side of failure: Investigating the impact of speech recognition errors on subsequent dementia classification,” *arXiv preprint arXiv:2211.07430*, 2022.
- [31] C. Li, W. Xu, T. Cohen, and S. Pakhomov, “Useful blunders: Can automated speech recognition errors improve downstream dementia classification?” *Journal of Biomedical Informatics*, vol. 150, p. 104598, 2024.
- [32] A. Balagopalan, K. Shkaruta, and J. Novikova, “Impact of asr on alzheimer’s disease detection: All errors are equal, but deletions are more equal than others,” in *W-NUT*, 2020, pp. 159–164.
- [33] S. Luz, F. Haider, S. de la Fuente, D. Fromm, and B. MacWhinney, “Alzheimer’s Dementia Recognition Through Spontaneous Speech: The ADReSS Challenge,” *INTERSPEECH*, 2020.
- [34] J. T. Becker, F. Boiler, O. L. Lopez, J. Saxton, and K. L. McGonigle, “The natural history of alzheimer’s disease: description of study cohort and accuracy of diagnosis,” *Archives of neurology*, 1994.
- [35] E. Kaplan, H. Goodglass, and S. Weintraub, *Boston Naming Test*. Philadelphia, PA: Lea & Febiger, 1983.
- [36] S. Bird and E. Loper, “NLTK: The natural language toolkit,” in *Proceedings of the ACL Interactive Poster and Demonstration Sessions*. Barcelona, Spain: Association for Computational Linguistics, Jul. 2004, pp. 214–217. [Online]. Available: <https://aclanthology.org/P04-3031>