

# Factors Affecting the Performance of Automated Speaker Verification in Alzheimer’s Disease Clinical Trials

Malikeh Ehghaghi, Marija Stanojevic, Ali Akram, Jekaterina Novikova

Winterlight Labs Inc., Toronto, Canada

{malikeh, marija, aliakram, jekaterina}@winterlightlabs.com

## Abstract

Detecting duplicate patient participation in clinical trials is a major challenge because repeated patients can undermine the credibility and accuracy of the trial’s findings and result in significant health and financial risks. Developing accurate automated speaker verification (ASV) models is crucial to verify the identity of enrolled individuals and remove duplicates, but the size and quality of data influence ASV performance. However, there has been limited investigation into the factors that can affect ASV capabilities in clinical environments. In this paper, we bridge the gap by conducting analysis of how participant demographic characteristics, audio quality criteria, and severity level of Alzheimer’s disease (AD) impact the performance of ASV utilizing a dataset of speech recordings from 659 participants with varying levels of AD, obtained through multiple speech tasks. Our results indicate that ASV performance: 1) is slightly better on male speakers than on female speakers; 2) degrades for individuals who are above 70 years old; 3) is comparatively better for non-native English speakers than for native English speakers; 4) is negatively affected by clinician interference, noisy background, and unclear participant speech; 5) tends to decrease with an increase in the severity level of AD. Our study finds that voice biometrics raise fairness concerns as certain subgroups exhibit different ASV performances owing to their inherent voice characteristics. Moreover, the performance of ASV is influenced by the quality of speech recordings, which underscores the importance of improving the data collection settings in clinical trials.

## 1 Introduction

Healthcare systems are increasingly relying on automatic speaker verification (ASV) models to ensure secure and accurate identification of patients and healthcare providers, with the aim of preventing fraud, safeguarding patient privacy, and ensur-

ing the accuracy of medical records (Upadhyay et al., 2022; Arasteh et al., 2022).

Conducting large-scale clinical trials, involving numerous patients, doctors, clinics, and even different countries can pose significant challenges in identifying instances of duplicate participation, which occurs when a single individual joins the same study more than once, either at different sites or time points, leading to skewed results and undermining the validity of study findings (Irum and Salman, 2019). Shiovitz et al. (2013) discovered that as much as 7.78% of patients involved in a clinical study were duplicated across different sites.

In some cases, individuals participate in multiple clinical trials concurrently in order to earn more money. When a trial enrolls an adequate number of substandard participants, it risks not meeting the primary endpoints and ultimately causing a multimillion-dollar study to fail. Pinho et al. (2021) examined the financial effect of duplicate participants on the pharmaceutical companies conducting a set of short-term study programs across psychiatric disorders including Schizophrenia, Major Depressive Disorder, and Bipolar Depression. Based on their results, enrolling ineligible subjects in the selected studies results in a loss of around \$29,680,000 for the sponsor pharmaceutical company. In addition, duplicate participation results in higher placebo rates and compromised data integrity. These findings highlight the importance of addressing the duplicate participant problem and underscore the need for reliable and accurate ASV methods in healthcare systems to verify whether an unknown voice belongs to a known enrolled individual (Upadhyay et al., 2022; Arasteh et al., 2022).

Cognitive impairment has been linked to a decline in vocabulary richness, syntactic complexity, and speech fluency, according to previous research (Thomas et al., 2005; Roark et al., 2011; Guinn and Habash, 2012; Meilán et al., 2012). Therefore,

it is important to investigate whether the abnormal speech patterns exhibited by individuals with cognitive impairment can affect ASV performance. Despite this concern, there is a paucity of research examining the relationship between cognitive impairment and ASV in the existing literature. This research gap motivated us to address this issue by examining the effect of Alzheimer’s disease (AD) severity level on ASV performance.

Furthermore, external factors such as participants’ demographic information (Si et al., 2021), recording environment, or data collection procedure (Woo et al., 2006; Wan, 2017) may also have an impact on ASV performance, but their impact is not well-studied in the healthcare industry. An extensive analysis of these external factors could provide valuable insights into the accuracy and reliability of ASV models and identify potential sources of bias due to differences in inherent voice characteristics among subgroups (Si et al., 2021).

The purpose of this study is to investigate the effectiveness of ASV models in identifying duplicate patient participation in large-scale clinical trials, and to explore the factors that influence ASV performance in such settings. To this end, we utilize a longitudinal clinical dataset of English speech recordings obtained through multiple speech tasks from 659 participants with varying levels of AD. We employ the TitaNet model, an end-to-end deep learning text-independent ASV model pre-trained on a large volume of speech recordings of English speakers. ASV models can be classified into two groups: text-dependent (TD) and text-independent (TI). TI ASV models allow for more flexibility in the enrollment and verification phases without constraints on the speech content. When pre-trained on extensive audio datasets, TI models demonstrate a comparable level of accuracy to TD models. We evaluate the performance of TitaNet on our dataset in a zero-shot setting, achieving a 3.1% equal error rate (EER). In addition, we analyze the impact of various external factors on ASV performance, including participant demographic characteristics (i.e., age, and gender), audio quality criteria (i.e., clinician interference, background noise, participant accent, and participant clarity), as well as AD severity level. This study aims to provide valuable insights into the factors that can affect the performance of ASV models in clinical trial environments, with the goal of improving the accuracy, fairness, and reliability.

Our findings indicate that voice biometrics may present fairness issues, as certain subgroups demonstrate differing speaker verification performances due to their inherent voice characteristics. In addition, the quality of speech recordings can impact ASV performance, highlighting the importance of monitoring and enhancing data collection and recording settings during clinical trials.

## 2 Related Work

Speaker verification technology has been increasingly utilized in various domains, including healthcare systems. Several studies have analyzed the feasibility and effectiveness of speaker verification models in healthcare settings (Hao and Hei; Weng et al., 2014). However, the external factors that can affect the performance of ASV models has not been extensively studied through research in the healthcare field.

**Race and Gender Effect:** Si et al. (2021) utilized three state-of-the-art ASV models including the Xvector-TDNN (Snyder et al., 2018), ECAPA-TDNN (Desplanques et al., 2020), and DTW (Dutta, 2008) models to explore demographic effects on speaker verification. For this purpose, they used a subset of the mPower study (Bot et al., 2016), a Parkinson’s disease mobile dataset, comparing a diverse group of 300 speakers by race and gender. Their results demonstrated that the Latinx subgroup indicates the worst ASV performance among the four major races in the dataset (i.e., White, Black, Latinx, and Asian). Based on their findings, gender represents minor differences in ASV performance between male-only and female-only subgroups. We did similar gender-level and accent-level analyses on patients with Alzheimer’s disease to detect the potential sources of bias in ASV due to the inherent voice characteristics of distinct genders or English accents.

**Age Effect:** Kelly and Harte (2011) analyzed the effect of long-term ageing on ASV performance. They utilized a conventional GMM-UBM system (Irum and Salman, 2019) on a longitudinal voice dataset of a cohort of 13 adult speakers, whose recordings were collected over a time span of 30-40 years. According to their results, short-term aging (less than 5 years) does not have a significant impact on verification performance, compared to normal inter-session variations. However, for longer periods, aging has a negative effect on veri-

fication accuracy. Moreover, the researchers found that the rate of verification score decline is more rapid for speakers aged 60 years and above. However, they only evaluated their models on small cohorts and inter-speaker differences across different age groups were not further analyzed, while in the present work, we evaluate the ASV performance across different age groups over 55 years old and incorporate a larger clinical dataset with 659 speakers in total, while we controlled for AD effect. [Taylor et al. \(2020\)](#) also demonstrated that some speech and vocal characteristics (e.g., the spectral center of gravity, spectral skewness, or spectral kurtosis) undergo alterations with aging, and these changes can vary between men and women. These findings suggest that age is an effective factor in speaker’s voice characteristics and this underscores the importance of assessing age effect on our ASV model to ensure the fairness of the model across different age groups.

**Noise Effect:** [Wan \(2017\)](#) applied LibriSpeech corpus ([Panayotov et al., 2015](#)) of English novel reading speech with varying lengths and tested ASV performance across different types and levels of background noise (e.g., babble, car, office and airplane noise) with a great mismatch between training and testing speech. Based on their findings, performance varies across different types of noise and the number of errors grow with a decrease in the sound-to-noise-ratio value. However, other metrics of audio quality were not considered in their study and their models were only trained and tested on healthy speech recordings. In the present study, we assess the effect of other audio quality aspects, such as participant clarity, clinician interference, and background noise, on ASV performance using a dataset of speakers with varying severity levels of AD, which was collected in clinical environment.

**Speech Pathology Effect:** [Arasteh et al. \(2022\)](#) have investigated the vulnerability of pathological speech to re-identification in ASV systems. In a large-scale study, they explored the effects of different speech pathologies on ASV using a real-world pathological speech corpus of more than 2,000 test subjects with various speech and voice disorders. Their results indicated that some types of speech pathology, particularly dysphonia, regardless of speech intelligibility, are more vulnerable to a breach of privacy compared to healthy speech. They do not analyze the effect of AD on ASV per-

formance, even though speech and language impairment are prevalent issues in moderate to severe stages of AD that may potentially affect the ASV performance. This motivates us to evaluate ASV performance across varying severity levels of AD.

## 3 Methods

### 3.1 Datasets

The Alzheimer’s Disease Clinical Trial (ADCT) dataset comprises speech recordings of English-speaking patients with a clinical diagnosis of mild to moderate AD who participated in a clinical trial. This is a proprietary dataset, which was collected every 12 weeks for a 48-week treatment period. It includes recordings of participants performing a set of self-administered speech tasks, including picture description ([Goodglass et al., 2001](#); [Becker et al., 1994](#)), phonemic verbal fluency ([Borkowski et al., 1967](#)), and semantic verbal fluency ([Tombaugh et al., 1999](#)) tasks.

#### 3.1.1 Demographic Information

Demographic data were collected about the participants at the beginning of the study. This data includes the age, and gender of the individuals upon consenting. The data collection study was approved by the ethical committee.

#### 3.1.2 Transcription and Quality Assessment

All the audio recordings were manually transcribed by 49 trained transcriptionists based on the CHAT protocol and annotations ([MacWhinney, 2014](#)). The transcriptionists utilized an online tool that granted them access to the recordings and enabled them to transcribe the audio content, segment the files into utterances, and perform quality assessment. In addition, the transcriptions manually rated the quality of the recordings according to different quality criteria. The values range from 0 to 3 for each quality criterion. Value higher than 0 indicates that the audio recording has minor to major issues under that quality criterion. The quality criteria consist of background noise, clinician interference, participant accent, and participant clarity. The background noise criterion indicates whether there is noise in the background from the environment. Clinician interference indicates whether the clinician (or another speaker) interferes with the speech task. The participant accent criterion indicates whether the participant is a native or near-native speaker (values of 0) or has a detectable non-native accent (value higher than 0). Participant

clarity indicates whether the participant’s voice is hard to hear or understand.

### 3.1.3 Clinical Assessment

Participants were assessed on the severity level of AD using the Mini-Mental State Examination (MMSE) (Hennekes et al., 2016) rating scale. MMSE is a brief cognitive function assessment, which consists of 30 questions that can be completed in less than 10 minutes. The questions are divided into seven categories, with each subscore examining a particular aspect of cognition: Orientation in time (score range 0-5), orientation in place (score range 0-5), registration (score range 0-3), attention and concentration (score range 0-5), recall (score range 0-3), language (score range 0-8), and drawing (score range 0-1). MMSE total score ranges from 0 to 30, with higher scores indicating better cognitive function and lower scores indicating more severe cognitive impairment. In this study, the participants were categorized into four levels of AD severity based on MMSE criteria (Wimo et al., 2013): Healthy Control (HC) (MMSE score > 26 points), Mild AD (MMSE score 21-26 points), Moderate AD (MMSE score 15-20 points), and Severe AD (MMSE score < 15 points).

### 3.1.4 Dataset Composition

The ADCT data comprises 7084 audio recordings from 659 speakers with  $10.7 \pm 7.0$  samples on average per each speaker. The average duration of total audio and speech-only audio are equal to 69.31 and 37.30 seconds, respectively.

In the dataset, 43.4% of the speakers are male and 56.6% of the speakers are female. The age range of the subjects spans from 55 to 80 years old. Age distribution of the subjects is represented in Figure 1, with an average value equal to  $69.7 \pm 6.7$ .

Figure 2 indicates the distribution of MMSE scores in the ADCT dataset, showing that the majority of the samples consist of mild to severe levels of AD with scores in the range of 15 to 26 points. The average MMSE score is equal to  $17.3 \pm 4.4$ . It should also be noted that the severity level of AD may vary over time for some of the speakers.

## 3.2 Models

In this study, we utilized the TitaNet model (Koluguri et al., 2022), which is a state-of-the-art end-to-end TI ASV model from the Nvidia NeMo toolkit<sup>1</sup> that had been pre-trained on an

<sup>1</sup><https://github.com/NVIDIA/NeMo>

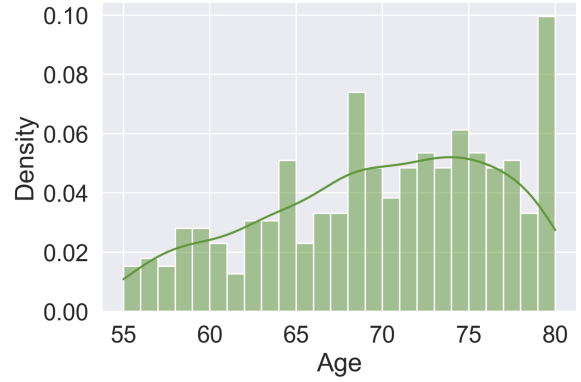


Figure 1: Age distribution of the speakers in ADCT dataset.

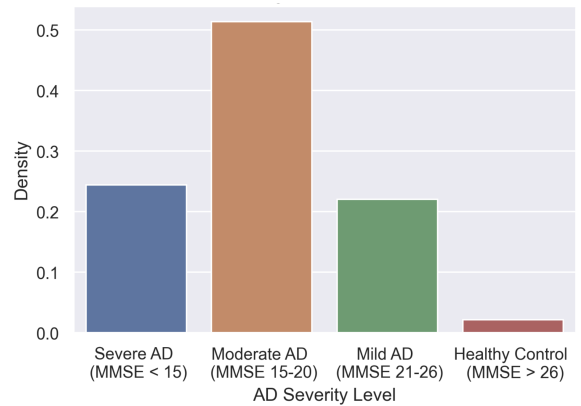


Figure 2: Distribution of AD severity levels in ADCT dataset.

extensive collection of English speech data, from various publicly-available resources. The TitaNet model is a neural network model that adopts an encoder-decoder architecture to extract speaker embeddings from speech. The model architecture is inspired by ContextNet (Han et al., 2020) model, which comprises 1D depth-wise separable convolutions followed by squeeze and excitation (SE) layers combined with channel attention pooling to convert utterances of varying lengths into a fixed-length embedding. The model contains 25.3M parameters and it is pre-trained on the VoxCeleb1 Dev (Nagrani et al., 2017), VoxCeleb2 Dev (Chung et al., 2018), Fisher (Cieri et al., 2004), Switchboard-Cellular1, Switchboard-Cellular2 (Godfrey and Holliman, 1993), and LibriSpeech (Panayotov et al., 2015) datasets.

We applied the model to the ADCT dataset in a zero-shot setting and used it for further analysis of the effect of external factors.

### 3.3 Evaluation

To analyze the effect of external factors on ASV performance, we separately evaluated the performance of the TitaNet model across different subsets of ADCT data including genders, age groups, audio quality levels, and AD severity levels. Initially, we produced embeddings for all audio files within each group. Subsequently, we aggregated these embeddings to create a set of tuples comprising positive and negative pairs.

**Positive tuples** refer to pairs of embeddings that belong to the same speaker. The total number of positive tuples in each group is given by  $\sum_{i=1}^m \binom{n_i}{2}$ , which is calculated by summing up all pairs of  $n_i$  speech recordings for the  $i^{\text{th}}$  speaker, where  $m$  is the total number of speakers in the same group.

**Negative tuples** refer to pairs of embeddings from different speakers within the same group. The total number of negative pairs is calculated as  $\sum_{i=1}^m n_i * (N - n_i)/2$ , where  $N$  is the total number of audio files in the group, and  $n_i$  is the number of speech recordings for speaker  $i$ . The sum is divided by 2 to avoid counting each pair twice.

After generating all the positive and negative pair tuples, we proceeded to compute the cosine similarity between the pairs of vector embeddings within each tuple. Subsequently, we adjusted a threshold value  $\theta$  for each group through manual tuning until the true positive rate equalled the true negative rate, which enabled us to evaluate the performance of the ASV model using the equal error rate (EER) metric. If the cosine similarity value exceeded the threshold, we considered the corresponding tuple as belonging to the same speaker. Conversely, if the cosine similarity value was below the threshold, we deemed the two embeddings to represent different speakers.

## 4 Results and Discussion

In order to have a baseline level of ASV performance, we evaluated the TitaNet model on all the speech recordings of the ADCT dataset and obtained a 3.1% EER. To further analyze the impact of participant demographic characteristics, audio quality and AD severity level, we then recalculated EER and compared the ASV performance across different subgroups.

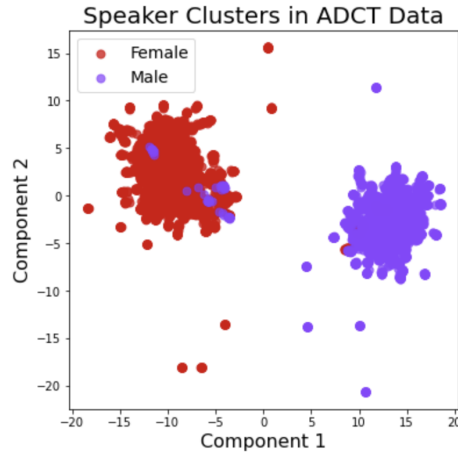


Figure 3: Visualization of speaker clusters using TitaNet embeddings of all audio recordings in ADCT dataset, created using the PaCMAP (Wang et al., 2021) dimensionality reduction method, where each color represents the gender of the speaker.

### 4.1 Effect of Participant Demographic Characteristics on ASV Performance

#### 4.1.1 Gender Effect

We first analyzed the effect of speaker genders on the performance of ASV. As shown in Figure 3, two visually distinguishable speaker clusters appeared in the visualization of speech embeddings of all speakers in our dataset. Colouring the data points based on their gender indicates that each cluster is representative of a specific gender. The left cluster mostly comprises female speakers and the majority of the right cluster consists of male speakers. For further analysis of the ASV performance, we separately evaluated the model performance on male and female speech samples within each cluster. To control for the confounding factors, we randomly downsized the size of the female subgroup to the number of speakers in the male subgroup and also, matched the average age and MMSE score between the two subgroups. The results (Table 1) show that according to the EER metric, the ASV model performs better on the total dataset comprising diverse genders compared to when it is applied to the male-only or female-only speakers (Table 1). The ASV performance for the male subgroup (EER = 4.98%) is slightly better than that for the female subgroup (5.13% EER), although the difference is not substantial. This is in line with prior literature (Hanilçi and Ertaş, 2013) demonstrating that male speakers exhibit higher speaker recognition accuracy compared to female speakers regardless of the database

Gender	Tuned Thr.	EER(%)	#Spkrs	#Smpls	Avg #Smpls per Spkr	Avg Age	Avg MMSE Score
Female	0.75	5.13	170	2735	16.09±3.94	69.53±6.72	17.33±4.37
Male	0.75	4.98	170	2671	15.72±4.02	69.41±6.96	17.45±4.45
All	0.74	<b>3.10</b>	659	7084	10.70±7.00	69.55±6.75	17.32±4.44

Table 1: Benchmarking TitaNet ASV model across different genders on the ADCT English dataset. ‘#Spkrs’ denotes the number of speakers per each gender. ‘#Smpls’ denotes the number of audio recordings per each gender. ‘Tuned Thr.’ denotes tuned threshold.

Age Group	Tuned Thr.	EER(%)	#Spkrs	#Smpls	Avg #Smpls per Spkr	Gender	Avg MMSE Score
Age <= 70	0.73	3.62	197	3235	16.42±3.86	Male+Female	17.09±4.72
Age > 70	0.74	4.20	195	3022	15.50±4.07	Male+Female	17.57±4.11
All	0.74	<b>3.10</b>	659	7084	10.7±7.0	Male+Female	17.32±4.44

Table 2: Benchmarking TitaNet ASV model across different age groups on the ADCT English dataset. ‘#Spkrs’ denotes the number of speakers per each age group. ‘#Smpls’ denotes the number of audio recordings per each age group. ‘Tuned Thr.’ denotes tuned threshold.

and classifier used. The results also align with Si et al. (2021) indicating that there is little difference in gender in terms of the performance of ASV models in general.

#### 4.1.2 Age Effect

We then evaluated the age effect on ASV performance. For this purpose, we categorized the speakers into two age subgroups with Age <= 70 and Age > 70 according to the speaker ages at the study enrollment date. The threshold was set to 70 because it is equal to the approximate median and mean value of the age distribution of the speakers in the ADCT dataset (Section 3.1.4). To control for the confounding factors, we designed our experiments to ensure that the age-based subgroups had a comparable number of speakers and average MMSE scores and included a combination of male and female speakers.

Based on the results indicated in Table 2, EER for participants under 70 is 0.58% lower than the older age group. These results can be explained by Taylor et al. (2020) revealing that specific attributes of speech and voice characteristics (e.g., fricative spectral moments, semitone standard deviation, etc.) vary according to age.

## 4.2 Effect of Audio Quality on ASV Performance

We examined how different audio quality factors affect the performance of speaker verification in clinical environments. We divided ADCT samples

into two subgroups based on their quality rating for each criterion: ‘No Issue’ for samples with a rating of 0 and ‘Minor to Major Issue’ for those with a rating higher than 0 (maximum is 3). In order to mitigate the influence of possible confounding factors, we structured our experiments in such a way that the quality-based subgroups had a similar number of speakers and average MMSE score, and included both male and female speakers. Table 3 shows the comparison of EER values between each pair of subgroups per audio quality criterion. Our results indicate that subgroups of audio samples with no background noise, and high participant clarity yielded lower EER than subgroups with varying levels of background noise and poor participant clarity. Therefore, better control of the data collection setting, along with the use of high-clarity audio recordings with minimal background noise, would be recommended in order to improve ASV performance in clinical trials. Our findings are in line with Eskimez et al. (2018), who demonstrated that incorporating a DNN-based speech enhancement technique as a front-end noise reduction module can enhance the ASV performance when applied to noisy speech data obtained from real customers.

Our results also suggest that clinician interference can negatively impact ASV performance with 0.48% increase in EER. Therefore, it is recommended that clinicians refrain from interrupting participants during speech tasks in recording sessions to prevent any decline in performance.

We also evaluated the speakers’ accents as a qual-

Audio Quality Criterion	Tuned Thr.	EER(%)	#Spkrs	#Smpls	Avg #Smpls per Spkr	Gender	Avg Age	Avg MMSE Score
Background Noise - No Issue	0.75	<b>2.90</b>	125	426	3.40±1.45	M + F	69.60±6.72	16.94±5.83
Background Noise - Minor to Major Issue	0.74	3.54	125	511	4.08±2.08	M + F	69.21±6.45	16.78±5.58
Participant Clarity - No Issue	0.75	<b>2.85</b>	112	481	4.29±1.70	M + F	69.80±6.38	16.81±5.62
Participant Clarity - Minor to Major Issue	0.74	3.41	112	432	3.85±1.83	M + F	69.23±6.81	16.04±5.54
Clinician Interference - No Issue	0.75	<b>2.90</b>	103	659	4.30±2.08	M + F	69.40±6.77	17.65±5.61
Clinician Interf. - Minor to Major Issue	0.73	3.38	103	399	3.87±1.86	M + F	69.43±6.84	14.77±5.22
Participant Accent - Native	0.74	2.97	188	901	4.79±2.82	M + F	68.63±6.89	17.22±5.01
Participant Accent - Non-Native	0.74	<b>2.01</b>	188	594	3.16±1.54	M + F	70.45±6.32	17.19±4.56
All	0.74	3.10	659	7084	10.7±7.0	M + F	69.55±6.75	17.32±4.44

Table 3: Benchmarking TitaNet speaker verification model across different levels of audio quality on the ADCT dataset. ‘#Spkrs’ denotes the number of speakers per each quality subgroup. ‘#Smpls’ denotes the number of audio recordings per each quality subgroup. ‘Tuned Thr.’ denotes tuned threshold. For each quality criterion, ‘No Issue’ indicates samples with a rating = 0 and ‘Minor to Major Issue’ indicates samples with a rating > 0 (maximum is 3). ‘M’ denotes male speakers and ‘F’ denotes female speakers. Bold font denotes the subgroup yielding best ASV performance in each quality criterion.

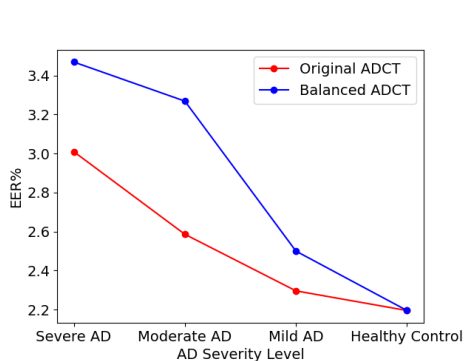


Figure 4: Comparison of the performance of the TitaNet ASV model across different AD severity levels based on EER% metric. Original ADCT refers to the dataset with the original number of speakers per severity level. Balanced ADCT refers to the dataset with each group downsized to the number of speakers in the smallest group, which is the HC group.

ity indicator, and our results show that ASV performs better on non-native speakers (2.01% EER) than on native speakers (2.97% EER). This suggests that ASV performance can be better in trials that involve participants who speak with diverse non-native English accents as a way to identify unique speech characteristics for each individual.

### 4.3 Effect of the Severity Level of Alzheimer’s Disease on ASV Performance

We examined how the AD severity level impacts the performance of the ASV model. For this purpose, we performed a separate recalculation of EER for subgroups of audio samples consisting of Severe AD (Number of speakers = 218), Moderate AD (Number of speakers = 436), Mild AD (Number of speakers = 244), and HC (Number of speakers = 25), while retaining the original number of speakers. To establish a fair comparison, we then bal-

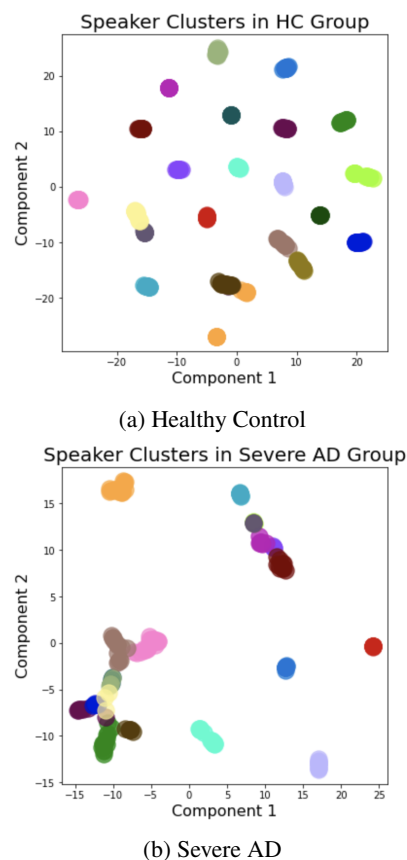


Figure 5: Speaker cluster visualizations of HC and Severe AD groups based on TitaNet embeddings for ADCT dataset, created using the PaCMAP (Wang et al., 2021) dimensionality reduction method, with each colour representing a distinct speaker.

anced the number of speakers across the subgroups by downsizing each to the smallest subgroup size of 25 speakers, which was the size of the HC subgroup. In both scenarios, higher AD severity levels lowered speaker verification performance, by about 1% to 1.5% of EER (Figure 4). Also, EER was lower within the groups where a higher num-

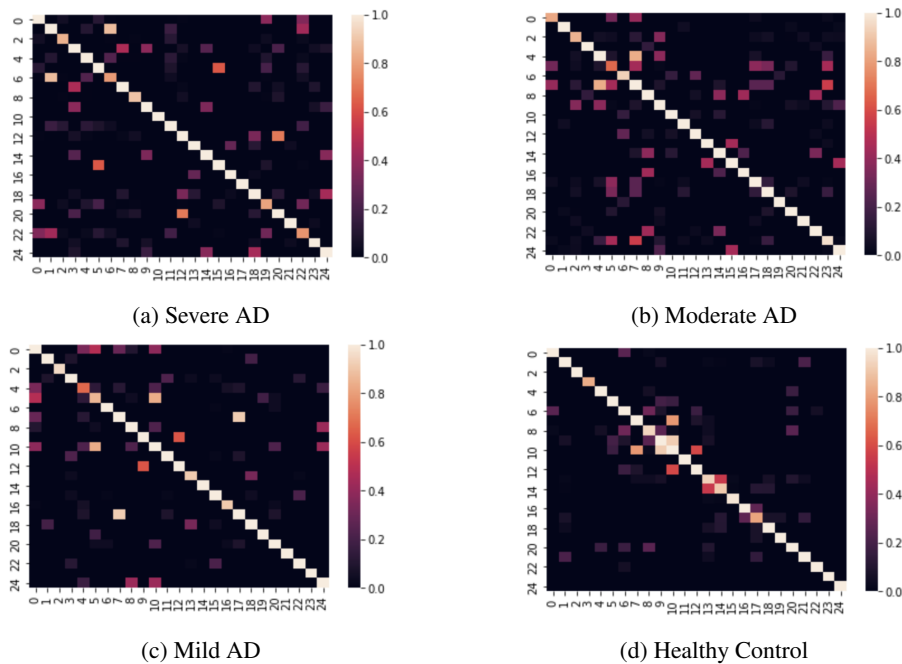


Figure 6: Similarity heatmap visualizations for ADCT dataset using TitaNet embeddings for different AD severity levels each downsized to the number of speakers in the smallest group, which is equal to 25. The lighter colours correspond to a higher level of similarity between the associated speaker tuples.

ber of speakers were included. Figure 5a and 5b display the PaCMAP (Wang et al., 2021) visualization of TitaNet speaker embeddings for the HC and Severe AD groups in the balanced ADCT dataset, where each colour representing a unique speaker. In both the HC and Severe AD groups, samples from the same speakers are clustered close to each other, while in the HC group, the samples from different speakers are more distinguishable compared to the Severe AD group. Moreover, there is a higher level of similarity between negative speaker tuples within Moderate and Severe AD subgroups in comparison to Mild AD and HC subgroups, as indicated in the similarity heatmaps for different AD severity levels (Figure 6). Overall, our findings suggest that the unique voice characteristics associated with varying levels of AD severity (Boschi et al., 2017) can be entangled with the identity of the speaker and may introduce a potential source of bias in the ASV models.

## 5 Conclusion

Large-scale clinical trials require accurate verification of participants, as duplicate participation may lead to substandard data quality and significant financial and health risks. Therefore, developing accurate ASV models for verifying participant identity is essential in these settings. External fac-

tors such as participant profile or audio quality can cause errors and biases in ASV performance during the trials, but limited research has been conducted in this area. In the present work, we utilize a longitudinal speech dataset of participants with varying levels of AD severity and investigate the impact of external factors, such as different participant demographic characteristics, audio quality criteria, and AD severity levels, on the performance of an end-to-end TI ASV model. Our findings show that variations in ASV performance can be attributed to the inherent voice characteristics of different subgroups (e.g., different ages, genders, accents, or AD severity levels) that are likely to be confused with the identity of the speaker. Hence, it is critical to reassess this technology to mitigate the risk of potential biases toward certain subgroups. Based on our results, poor audio quality with unclear speech, noisy background, and clinician interference also negatively impacts ASV performance. This highlights the importance of quality assurance for the speech recordings during the trials. In future work, we aim to automate the audio quality assessment process by leveraging existing automated methods such as perceptual evaluation of speech quality (PESQ) (Rix et al., 2001) or short-time objective intelligibility (STOI) (Taal et al., 2010), which would reduce the human effort required for this task.



## References

- Soroosh Tayebi Arasteh, Tobias Weise, Maria Schuster, Elmar Noth, Andreas Maier, and Seung Hee Yang. 2022. The effect of speech pathology on automatic speaker verification—a large-scale study. [arXiv preprint arXiv:2204.06450](#).
- James T Becker, François Boiler, Oscar L Lopez, Judith Saxton, and Karen L McGonigle. 1994. The natural history of alzheimer’s disease: description of study cohort and accuracy of diagnosis. [Archives of neurology](#), 51(6):585–594.
- John G Borkowski, Arthur L Benton, and Otfried Spreen. 1967. Word fluency and brain damage. [Neuropsychologia](#), 5(2):135–140.
- Veronica Boschi, Eleonora Catricala, Monica Consonni, Cristiano Chesi, Andrea Moro, and Stefano F Cappa. 2017. Connected speech in neurodegenerative language disorders: a review. [Frontiers in psychology](#), 8:269.
- Brian M Bot, Christine Suver, Elias Chaibub Neto, Michael Kellen, Arno Klein, Christopher Bare, Megan Doerr, Abhishek Pratap, John Wilbanks, E Dorsey, et al. 2016. The mpower study, parkinson disease mobile data collected using researchkit. [Scientific data](#), 3(1):1–9.
- Joon Son Chung, Arsha Nagrani, and Andrew Senior. 2018. Voxceleb2: Deep speaker recognition. [arXiv preprint arXiv:1806.05622](#).
- Christopher Cieri, David Miller, and Kevin Walker. 2004. The fisher corpus: A resource for the next generations of speech-to-text. In [LREC](#), volume 4, pages 69–71.
- Brecht Desplanques, Jenthe Thienpondt, and Kris Demuynck. 2020. Ecapa-tdnn: Emphasized channel attention, propagation and aggregation in tdnn based speaker verification. [arXiv preprint arXiv:2005.07143](#).
- Tridibesh Dutta. 2008. Dynamic time warping based approach to text-dependent speaker identification using spectrograms. In [2008 Congress on Image and Signal Processing](#), volume 2, pages 354–360. IEEE.
- Sefik Emre Eskimez, Peter Soufleris, Zhiyao Duan, and Wendi Heintzelman. 2018. Front-end speech enhancement for commercial speaker verification systems. [Speech Communication](#), 99:101–113.
- John Godfrey and Edward Holliman. 1993. Switchboard-1 release 2 ldc97s62. [Linguistic Data Consortium](#), page 34.
- Harold Goodglass, Edith Kaplan, and Sandra Weintraub. 2001. [BDAE: The Boston Diagnostic Aphasia Examination](#). Lippincott Williams & Wilkins Philadelphia, PA.
- Curry I Guinn and Anthony Habash. 2012. Language analysis of speakers with dementia of the alzheimer’s type. In [2012 AAAI Fall Symposium Series](#).
- Wei Han, Zhengdong Zhang, Yu Zhang, Jiahui Yu, Chung-Cheng Chiu, James Qin, Anmol Gulati, Ruoming Pang, and Yonghui Wu. 2020. Contextnet: Improving convolutional neural networks for automatic speech recognition with global context. [arXiv preprint arXiv:2005.03191](#).
- Cemal Hanilçi and Figen Ertaş. 2013. Investigation of the effect of data duration and speaker gender on text-independent speaker recognition. [Computers & Electrical Engineering](#), 39(2):441–452.
- Bin Hao and Xiali Hei. Voice print liveness detection for medical devices.
- Carsten Hennes, Catherine Reed, Yun-Fei Chen, Grazia Dell’Agnello, and Jeremie Lebec. 2016. Describing the sequence of cognitive decline in alzheimer’s disease patients: results from an observational study. [Journal of Alzheimer’s Disease](#), 52(3):1065–1080.
- Amna Irum and Ahmad Salman. 2019. Speaker verification using deep neural networks: A. [International Journal of Machine Learning and Computing](#), 9(1).
- Finnian Kelly and Naomi Harte. 2011. Effects of long-term ageing on speaker verification. In [Biometrics and ID Management: COST 2101 European Workshop, BioID 2011, Brandenburg \(Havel\), Germany, March 8-10, 2011. Proceedings 3](#), pages 113–124. Springer.
- Nithin Rao Koluguri, Taejin Park, and Boris Ginsburg. 2022. Titanet: Neural model for speaker representation with 1d depth-wise separable convolutions and global context. In [ICASSP 2022-2022 IEEE International Conference on Acoustics, Speech and Signal Processing \(ICASSP\)](#), pages 8102–8106. IEEE.
- Brian MacWhinney. 2014. [The CHILDES project: Tools for analyzing talk, Volume II: The database](#). Psychology Press.
- Juan JG Meilán, Francisco Martínez-Sánchez, Juan Carro, José A Sánchez, and Enrique Pérez. 2012. Acoustic markers associated with impairment in language processing in alzheimer’s disease. [The Spanish journal of psychology](#), 15(2):487–494.
- Arsha Nagrani, Joon Son Chung, and Andrew Senior. 2017. Voxceleb: a large-scale speaker identification dataset. [arXiv preprint arXiv:1706.08612](#).
- Vassil Panayotov, Guoguo Chen, Daniel Povey, and Sanjeev Khudanpur. 2015. Librispeech: an asr corpus based on public domain audio books. In [2015 IEEE international conference on acoustics, speech and signal processing \(ICASSP\)](#), pages 5206–5210. IEEE.

- Andre Pinho, Kerri Weingard, and Mitchell D Efron. 2021. Improving safety and preventing failure in clinical trials by detecting and preventing duplicate and professional research subjects: the case for use of a research subject database registry. Journal of Clinical Trials, 11(4):1–6.
- Antony W Rix, John G Beerends, Michael P Hollier, and Andries P Hekstra. 2001. Perceptual evaluation of speech quality (pesq)-a new method for speech quality assessment of telephone networks and codecs. In 2001 IEEE international conference on acoustics, speech, and signal processing. Proceedings (Cat. No. 01CH37221), volume 2, pages 749–752. IEEE.
- Brian Roark, Margaret Mitchell, John-Paul Hosom, Kristy Hollingshead, and Jeffrey Kaye. 2011. Spoken language derived measures for detecting mild cognitive impairment. IEEE transactions on audio, speech, and language processing, 19(7):2081–2090.
- Thomas M Shiovitz, Charles S Wilcox, Lilit Gevorgyan, and Adnan Shawkat. 2013. Cns sites cooperate to detect duplicate subjects with a clinical trial subject registry. Innovations in Clinical Neuroscience, 10(2):17.
- Sophie Si, Zhengxiong Li, and Wenyao Xu. 2021. Exploring demographic effects on speaker verification. In 2021 IEEE Conference on Communications and Network Security (CNS), pages 1–2. IEEE.
- David Snyder, Daniel Garcia-Romero, Gregory Sell, Daniel Povey, and Sanjeev Khudanpur. 2018. X-vectors: Robust dnn embeddings for speaker recognition. In 2018 IEEE international conference on acoustics, speech and signal processing (ICASSP), pages 5329–5333. IEEE.
- Cees H Taal, Richard C Hendriks, Richard Heusdens, and Jesper Jensen. 2010. A short-time objective intelligibility measure for time-frequency weighted noisy speech. In 2010 IEEE international conference on acoustics, speech and signal processing, pages 4214–4217. IEEE.
- Sammi Taylor, Christopher Dromey, Shawn L Nissen, Kristine Tanner, Dennis Eggett, and Kim Corbin-Lewis. 2020. Age-related changes in speech and voice: spectral and cepstral measures. Journal of Speech, Language, and Hearing Research, 63(3):647–660.
- Calvin Thomas, Vlado Keselj, Nick Cercone, Kenneth Rockwood, and Elissa Asp. 2005. Automatic detection and rating of dementia of alzheimer type through lexical analysis of spontaneous speech. In IEEE International conference mechatronics and automation, 2005, volume 3, pages 1569–1574. IEEE.
- Tom N Tombaugh, Jean Kozak, and Laura Rees. 1999. Normative data stratified by age and education for two measures of verbal fluency: Fas and animal naming. Archives of clinical neuropsychology, 14(2):167–177.
- Shrikant Upadhyay, Mohit Kumar, Ashwani Kumar, Ramesh Karnati, Gouse Baig Mahommad, Sara A Althubiti, Fayadh Alenezi, and Kemal Polat. 2022. Feature extraction approach for speaker verification to support healthcare system using blockchain security for data privacy. Computational and Mathematical Methods in Medicine, 2022.
- Qianhui Wan. 2017. Speaker Verification Systems under Various Noise and SNR Conditions. Ph.D. thesis, Université d’Ottawa/University of Ottawa.
- Yingfan Wang, Haiyang Huang, Cynthia Rudin, and Yaron Shaposhnik. 2021. Understanding how dimension reduction tools work: An empirical approach to deciphering t-sne, umap, trimap, and pacmap for data visualization. Journal of Machine Learning Research, 22(201):1–73.
- Fuliang Weng, Taufiq Hasan, and Zhe Feng. 2014. Speaker verification in a health monitoring system. US Patent 8,818,810.
- Anders Wimo, Catherine C Reed, Richard Dodel, Mark Belger, Roy W Jones, Michael Happich, Josep M Argimon, Giuseppe Bruno, Diego Novick, Bruno Vellas, et al. 2013. The geras study: a prospective observational study of costs and resource use in community dwellers with alzheimer’s disease in three european countries—study design and baseline findings. Journal of Alzheimer’s Disease, 36(2):385–399.
- Ram H Woo, Alex Park, and Timothy J Hazen. 2006. The mit mobile device speaker verification corpus: data collection and preliminary experiments. In 2006 IEEE Odyssey-The Speaker and Language Recognition Workshop, pages 1–6. IEEE.