

Data-driven Approach to Differentiating between Depression and Dementia from Noisy Speech and Language Data

Malikeh Ehghaghi^{1,2}, Frank Rudzicz^{1,2,3,4,5}, Jekaterina Novikova¹

¹Winterlight Labs, Toronto, ON

²Department of Computer Science, University of Toronto, ON

³Vector Institute for Artificial Intelligence, Toronto, ON

⁴Li Ka Shing Knowledge Institute, St Michael’s Hospital, Toronto, ON

⁵Surgical Safety Technologies Inc., Toronto, ON

{malikeh, jekaterina}@winterlightlabs.com, {frank}@spoclub.com

Abstract

A significant number of studies apply acoustic and linguistic characteristics of human speech as prominent markers of dementia and depression. However, studies on discriminating depression from dementia are rare. Co-morbid depression is frequent in dementia and these clinical conditions share many overlapping symptoms, but the ability to distinguish between depression and dementia is essential as depression is often curable. In this work, we investigate the ability of clustering approaches in distinguishing between depression and dementia from human speech. We introduce a novel aggregated dataset, which combines narrative speech data from multiple conditions, i.e., Alzheimer’s disease, mild cognitive impairment, healthy control, and depression. We compare linear and non-linear clustering approaches and show that non-linear clustering techniques distinguish better between distinct disease clusters. Our interpretability analysis shows that the main differentiating symptoms between dementia and depression are acoustic abnormality, repetitiveness (or circularity) of speech, word finding difficulty, coherence impairment, and differences in lexical complexity and richness.

1 Introduction

Depressive disorder and dementia are clinical conditions that both impose a substantial cost globally in terms of mortality and morbidity and have a significant negative impact on social and economic productivity (Jaeschke et al., 2021). Distinguishing between these conditions has proven to be a challenging task (Murray, 2010) as they frequently co-occur and have many overlapping symptoms such as apathy (Lee and Lyketsos, 2003), changes in sleep patterns (Thorpe, 2009), and concentration issues (Korczyń and Halperin, 2009). However, de-

pression is generally curable by either psychotherapy or medication, while dementia is a neurodegenerative disease, which is caused by irreversible deterioration of the nervous system. It is hence crucial to differentiate between these two conditions (Fraser et al., 2016b).

Previous studies demonstrated that machine learning methods and speech analysis are useful in detecting dementia from depression (Fraser et al., 2016b; Murray, 2010). However, the machine learning methods used in prior studies suffer from three main limitations:

Firstly, the datasets applied in prior literature only comprise Alzheimer’s disease (AD), healthy control (HC), and depression (Depr) samples of senior participants with similar demographic distributions and recording environments (Fraser et al., 2016b; Murray, 2010). In real world settings, the datasets are very noisy due to variations in the data collection procedures. Additionally, dementia is not necessarily of the AD type in all cases, and other types of dementia like mild cognitive impairment (MCI) can be included.

Secondly, to the best of our knowledge, previous studies have only used classification approaches to detect AD from HC (Pulido et al., 2020; Balagopalan et al., 2021; Balagopalan and Novikova, 2021), Depr from HC (Wu et al., 2022), or AD from Depr (Fraser et al., 2016b) using speech. This might not be an ideal simulation of the real world diagnosis procedure. In clinical diagnosis, the first step is to detect the symptoms and explore the pattern changes in patient records before diagnosing the disease (Regier et al., 2013), while in classification, we first map the samples to the disease labels and then, apply interpretability methods to explore the differentiating features between the classes (Gordon, 1999).

Lastly, prior studies demonstrated that acoustic

and linguistic features extracted from spontaneous speech provide valuable indicators of both mental disorders such as depression (Low et al., 2020) and cognitive impairment like AD or MCI (Fraser et al., 2016a; Boschi et al., 2017). However, they did not derive a strong conclusion about the main distinguishing speech-based symptoms in classifying dementia from depression (Fraser et al., 2016b).

To address the first limitation, we generate a novel aggregated dataset, which combines several speech datasets comprising AD, MCI, HC, and Depr labels with a variety of data collection procedures. To address the second and third limitations, we introduce a novel approach, which applies clustering techniques to inspect what data-driven feature categories (symptoms) are the main differentiators between AD, MCI, Depr, and HC samples. We then use the distinguishing symptoms as a feature selection technique to classify AD, MCI, and Depr. Our key findings indicate that 1) the non-linear clustering approaches outperform the linear techniques in terms of separability level of distinct disease clusters; 2) acoustic abnormalities, variations in lexical complexity and richness, repetitiveness (or circularity) of speech, word finding difficulty, and coherence impairment are the main differentiating symptoms to distinguish between different types of dementia (e.g., AD and MCI), and Depr; 3) data-driven differentiators are able to substantially improve performance of classification across diseases.

2 Related Work

There has been a substantial number of studies on detecting either dementia (e.g., MCI or AD) or depression from spontaneous speech. However, little has been done to distinguish dementia from depression using discourse patterns.

To discriminate dementia from depression, Fraser et al. (2016b) applied speech data from the Pitt corpus in the DementiaBank database (Becker et al., 1994), elicited from elderly participants through picture description task, with ‘Cookie Theft’ (Goodglass et al., 2001) used as a picture. The samples were labeled as either AD or HC based on a personal history and a neuropsychological assessment battery (Iverson et al., 2008). A subset of the samples were labeled as depressed or non-depressed based on the established threshold on Hamilton Depression Rating Scale (HAM-D) test scores (Bagby et al., 2004). To explore the distin-

guishing discourse patterns between AD and Depr, Murray (2010) collected a speech dataset of elderly participants (with Depr, AD, or HC labels) who completed a picture description task, with Norman Rockwell’s painting ‘The Soldier’ used as a picture. Samples with Depr were diagnosed based on DSM-IV criteria (Frances et al., 1995) and samples with AD met NINCDS-ADRDA criteria (Tierney et al., 1988) for probable AD. The datasets used in these studies didn’t include other types of dementia such as MCI, and all of their samples followed the same data collection procedure, while we create an aggregated dataset, which consists of AD, MCI, HC, and Depr samples from different speech datasets with various data collection procedures.

Murray (2010) examined whether elderly individuals with depression can be distinguished from those at early stages of AD through distinct patterns in narrative speech. Based on their findings, individuals with AD generated less informative speech compared to the depressed patients in their picture descriptions, while there were no significant differences in the informativeness of the narratives between HC and Depr samples. Furthermore, quantitative and syntactic measures of discourse did not differ across the three groups. However, Murray (2010) did not attempt to make predictions using the data.

Fraser et al. (2016b) investigated if the automated AD screening tools misclassify cognitively healthy participants with Depr as AD when using narrative speech. They also used linguistic and acoustic features to classify non-depressed AD subjects from those with comorbid depression from speech elicited through picture description task. In their study, they compared logistic regression (LR) with support vector machines (SVM) classification models. Their performance in distinguishing between depressed and non-depressed AD samples was moderate (accuracy = 0.658) due to a wide range of overlapping symptoms. In addition, they only applied classification approaches and they didn’t derive the most informative features discriminating between AD patients with and without depression. In the present work, we apply clustering approaches to cluster the diseases based on the similarities in the discourse patterns, and apply interpretability techniques to explore the distinguishing feature categories (symptoms) between distinct diagnosis labels (i.e., HC, AD, MCI, and Depr). We use the differentiating symptoms as a

feature selection technique to classify the diseases.

3 Methods

3.1 Dataset

In this paper, we generated an aggregated superset of the datasets listed in Table 1 that contains speech recordings of English-speaking participants describing pictures. All the audio recordings were manually transcribed by trained transcriptionists, using the CHAT protocol and annotations (MacWhinney, 2014).

| Dataset | AD | MCI | Depr | HC |
|------------------------------------|------|-----|------|------|
| DementiaBank (Becker et al., 1994) | 178 | 138 | 0 | 229 |
| Healthy Aging | 0 | 214 | 0 | 211 |
| ADReSS (Luz et al., 2020) | 54 | 0 | 0 | 54 |
| DEPAC+ (Tasnim et al., 2022) | 0 | 0 | 222 | 532 |
| AD Clinical Trial | 1616 | 0 | 0 | 0 |
| Aggregated dataset | 1848 | 352 | 222 | 1026 |

Table 1: Speech datasets used. For each dataset, the number of samples with each diagnosis label is reported in the following columns.

DementiaBank (Becker et al., 1994) and **ADReSS** (Luz et al., 2020) are the datasets of pathological speech elicited from participants through picture description task, with ‘Cookie Theft’ (Goodglass et al., 2001) used as a picture. The recordings are labeled as AD, MCI, and HC.

Healthy Aging is the dataset of speech elicited from community volunteers through picture description task, with ‘Family in the Kitchen’, ‘Man in the Living Room’, ‘Food Market’, ‘Picnic’, ‘Grandmother’s Birthday’, and ‘Romantic Dinner’ proprietary images. The recordings are labeled as possible HC and MCI. Soft labels are based on the established threshold on Montreal Cognitive Assessment (Nasreddine et al., 2005) screening tool.

DEPAC+ is the extended version of the **DEPAC** (Tasnim et al., 2022) dataset, with more samples collected using the same data collection procedure. This is a dataset of narrative speech elicited from participants through picture description task, with ‘Family in the Kitchen’ and ‘Man Falling’ images. The recordings are labeled as HC and Depr. Soft labels are based on the established threshold on Patient Health Questionnaire-9 (PHQ-9) (Kroenke et al., 2001) test scores¹.

AD Clinical Trial is a dataset of speech recordings from the baseline and screening visits of a clin-

¹The participants with a PHQ-9 score ≤ 9 were labeled as HC, and the remaining samples with a PHQ-9 score ≥ 10 met criteria for symptoms of depression.

ical trial elicited from participants through picture description task, with ‘Family in the kitchen’, ‘Man in the Living Room’, ‘Grandmother’s Birthday’, ‘Romantic Dinner’, and ‘Cookie Theft’ (Goodglass et al., 2001) images. All the recordings are labeled as AD according to the the National Institute on Aging/Alzheimer’s Association criteria (Frisoni et al., 2011).

All images other than ‘Cookie Theft’ (Goodglass et al., 2001) were designed to match the ‘Cookie theft’ picture in style and the amount of information content units according to picture design principles described by Patel and Connaghan (2014).

3.2 Feature Extraction

We extracted 220 acoustic features from audio, and 325 linguistic features from the associated transcripts. These features were classified into the following categories (the full list is in Appendix A):

Acoustic: This category includes spectral and voicing-related features (e.g., Mel-Frequency Cepstral Coefficients (MFCC) (Rudzicz et al., 2012), Fundamental frequency (F_0), or statistical functionals of Zero-Crossing Rate (ZCR) (Kulkarni, 2018)) describing the acoustic properties of the sound wave.

Syntactic Complexity: This category comprises variables like the frequencies of various production rules from the constituency parsing tree of the transcripts (Chae and Nenkova, 2009), or Lu’s syntactic complexity features (Lu, 2010) enumerating the rate of usage of different syntactic structures.

Discourse Mapping: This category consists of features such as utterance distances, or speech-graph features (Mota et al., 2012) like graph density (Mirheidari et al., 2018) to calculate the repetitiveness or circularity of speech.

Lexical Complexity and Richness: This category accounts for the variables like frequency of words, or measures of vocabulary diversity such as type-token ratio (Richards, 1987) describing the lexical complexity and vocabulary richness of the transcripts.

Information Content Units: This category includes variables such as the number of objects, subjects, locations, and actions used to measure the number of items correctly named in the picture description task previously found to be associated with memory impairment (Croisile et al., 1996).

Sentiment: This category contains features such as valence, arousal, and dominance scores (War-

riner et al., 2013) for all words and word types describing the sentiment of the words used.

Word Finding Difficulty: This category consists of features including speech rate, duration of words, and number of filled (e.g., um, uh) and un-filled pauses as signs of word finding difficulty, which result in less fluid or fluent speech.

Coherence (Global and Local): This category includes variables like average, minimum, and maximum cosine distances (Mirheidari et al., 2018) between subsequent utterances (local coherence) or between utterances and key words (global coherence) using word2vec (Church, 2017) representation of the utterances to calculate their semantic similarity.

4 Proposed Novel Approach: Data-Driven Approach to Detecting Differentiating Speech-based Symptoms between Dementia and Depression

4.1 Dimensionality Reduction and Clustering

We first applied dimensionality reduction techniques to the preprocessed features (see Appendix B). To explore linear dimension reduction approaches, we applied Principal Component Analysis (PCA) (Wold et al., 1987) as well as Linear Discriminant Analysis (LDA) (Izenman, 2013). For non-linear dimensionality reduction techniques, we used Uniform Manifold Approximation and Projection (UMAP) (McInnes et al., 2018) and T-distributed Stochastic Neighbour Embedding (t-SNE) (Van der Maaten and Hinton, 2008) (See details of implementation and hyperparameter setting in Appendix C).

Next, we clustered the low-dimensional data points by K-Means clustering (Mysiak, 2020) to group them in an unsupervised way into distinguishable clusters. Clusters were meant to represent groups associated with data labels - HC, AD, MCI, and Depr.

4.1.1 Performance Metrics

The performance of the clustering methods was measured based on the following metrics:

1. *Optimal number of disease clusters* determined by the elbow method (Yuan and Yang, 2019) after training K-Means clustering on the feature embeddings resulted from dimension reduction. The ideal case is to derive 4 distinct disease clusters in line with the 4 di-

agnosis labels in the aggregated dataset (i.e., HC, AD, MCI, and Depr).

2. *Silhouette score* (Rousseeuw, 1987) was used to measure the level of cluster separability. Its value ranges from -1 to 1. '1' means clusters are well apart from each other and clearly distinguished. '0' means that the distance between clusters is not significant. '-1' means clusters are assigned in the wrong way (Bhardwaj, 2020). The results were recorded for K=4 (the number of labels in the dataset), where K is the number of clusters generated by K-Means clustering.

4.2 Analysis of the Differentiating Feature Categories between the Disease Clusters

Analysis of the differentiating feature categories across the disease clusters consists of 3 main steps: LIME-based explanation of the low-dimensional embeddings, analysis of feature contributions to the non-linear components, and feature selection using the differentiating feature categories in classification of AD vs MCI vs Depr.

4.2.1 Local Explanation of the Non-linear Embeddings by LIME

We applied a LIME-for-t-SNE² interpretability method developed by Bibal et al., 2020 to find the main differentiating feature categories between AD, Depr, MCI, and HC diagnosis labels. This approach adapts Local Interpretable Model-agnostic Explanations (LIME) (Ribeiro et al., 2016) to locally explain t-SNE components.

4.2.2 Analysis of Feature Contributions to the Non-linear Components

In this experiment, we investigated what feature categories are the main differentiating factors between the distinguishable disease clusters derived by K-Means clustering. As the first step, we randomly selected 10 HC samples from each cluster and applied Lime-for-t-SNE model to explain the local trends in their neighborhood. We also picked 10 Depr and 10³ AD data points from the associated disease clusters and followed the same procedure to locally explain the low-dimensional components.

²Publicly available at <https://github.com/vu-minh/mlteam-lime-for-tsne>

³We selected 10 samples from each disease cluster, since each group must contain at least 5 samples for both Kruskal-Wallis H-Test (Lomuscio, 2021) and Mann-Whitney U-Test (Bedre, 2021) explained in 4.2.2.

| Dimension reduction method | Is the optimal number of clusters (K) equal to 4? | Silhouette score (K = 4) |
|----------------------------|---------------------------------------------------|--------------------------|
| PCA (linear) | - | 0.2010 |
| LDA (linear) | - | 0.4125 |
| t-SNE (non-linear) | x | 0.4723 |
| UMAP (non-linear) | x | 0.5580 |

Table 2: Summary of the performance of all dimensionality reduction techniques. The second column checks if the optimal number of clusters is equal with the total number of labels (e.g., HC, MCI, AD, and Depr) in the aggregated dataset. ‘K’ refers to the number of clusters in K-Means clustering applied on the embeddings in the low-dimensional space.

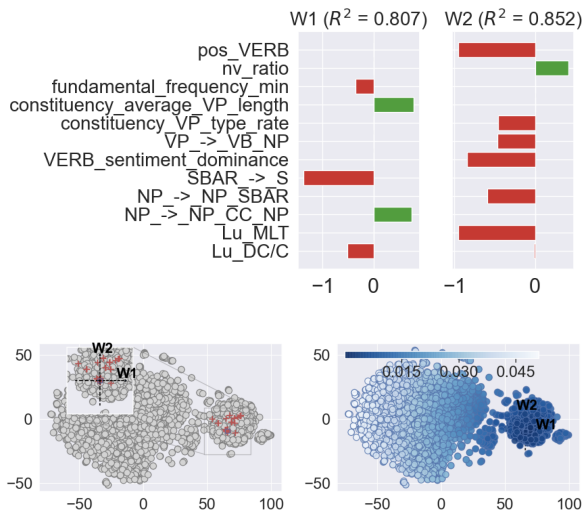


Figure 1: Explanation of the local trends in the t-SNE embeddings for a selected Depr instance. The figure at the top indicates the weights of the highly-contributed features explaining each local dimension (R^2 score indicates how well the local trends are linearly explained per each axis.) The blue transparency in the scatter plot represents the errors of the linear model applied locally on the original instance. The figure at the bottom left is a zoom on the zone of interest for local explanation, with projected samples in red (Bibal et al., 2020)

Figure 1 depicts an example of the local explanation of t-SNE embeddings for a selected Depr instance. For each candidate sample, we generated a vector of length 9 indicating the total number of highly-contributed features explaining either quasi-horizontal (e.g., $W1$ in Figure 1) or quasi-vertical (e.g., $W2$ in Figure 1) trends per each feature category including acoustic, syntactic complexity, discourse mapping, lexical complexity and richness, information content units, sentiment, word finding difficulty, coherence (global and local), and utterance cohesion.

Overall Group Comparison (Kruskal-Wallis H-Test): After calculating the feature frequency vectors of the selected samples, we applied overall

group comparison per each feature category to test the overall difference between the feature frequencies across the disease groups. For this purpose, we used Kruskal-Wallis H-test (Kruskal and Wallis, 1952) using the `scipy.stats.kruskal` library in python.

Pairwise Group Comparison (Mann-Whitney U-Test): As a post-hoc comparison method, we then applied pairwise Mann-Whitney U-test (Mann and Whitney, 1947) using the `scipy.stats.mannwhitneyu` python library to determine the distributions of which feature categories are significantly different between each pair of the selected disease groups.

4.2.3 Classification of AD vs Depr vs MCI

After analyzing the feature contributions to the non-linear components, we used the main differentiating feature categories as a feature selection technique to investigate whether they improve the classification performance of AD vs Depr vs MCI. For this purpose, we separately trained Multi-layer Perceptron classifier (MLPClassifier) on the following feature sets:

1. F : All the hand-crafted acoustic and linguistic features
2. F_d : Only the feature categories that are shown to be the main differentiators between AD and Depr based on Mann-Whitney test
3. $F - F_d$: All the hand-crafted features excluding the main distinguishing feature categories

We implemented MLPClassifier by `neural_network.MLPClassifier` package of Scikit-learn (Pedregosa et al., 2011) with all the hyperparameters set to their default parameter settings. We trained the models using grouped 10-fold cross validation to avoid overlapping subjects between the train and test folds and evaluated the performance of the models in terms of macro average accuracy, precision, recall, and F1 scores across the 10 folds.

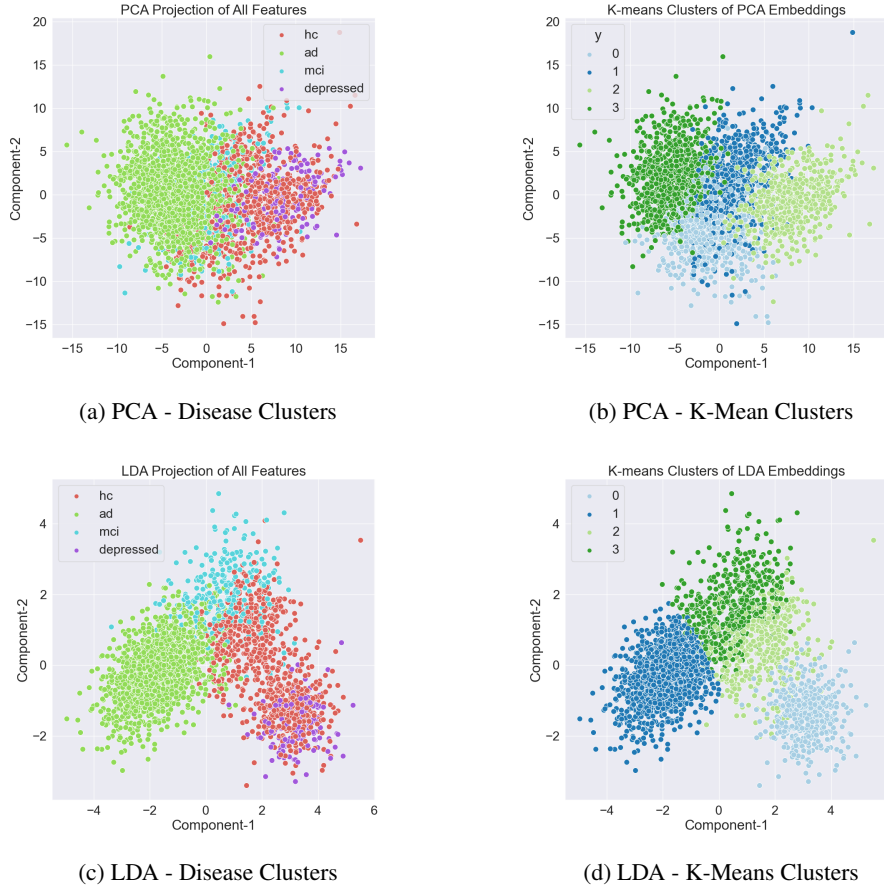


Figure 2: Pairwise scatter plots of the linear dimensionality reduction techniques (Component-1 vs Component-2). Left figures: 2-D representation of the samples colored based on their diagnosis labels. Right figures: 2-D representation of the samples colored based on the data-driven clusters resulted from K-Means clustering for $K=4$.

5 Results and Discussion

5.1 Comparison of Linear and Non-linear Dimensionality Reduction Approaches

Table 2 compares how the linear approaches (e.g., PCA, and LDA) perform versus the non-linear techniques (e.g., t-SNE, and UMAP) in distinguishing between distinct diagnosis labels (i.e., AD, MCI, HC, and Depr). Their performance is compared according to their optimal number of K-Means clusters, and Silhouette score. The second column of Table 2 represents whether the optimal number of data-driven disease clusters in K-Means clustering is equal to the total number of diagnosis labels in the aggregated dataset, which is our ideal case.

Between the linear techniques, the Silhouette score obtained by LDA is about twice in value compared to PCA. This can be due to the fact that LDA (Izenman, 2013) is a supervised dimensionality reduction technique which focuses on maximizing the class separability by projecting the data points

on a new linear axis, while PCA (Wold et al., 1987) tries to find the directions of maximal variance. Based on Figure 2c and 2d, the clusters of different diseases, as well as the K-Means clusters in LDA, are more visually distinguishable when compared to PCA (See Figure 2a and 2b). It is also interesting to note how the clusters are placed in LDA plots. MCI comes between AD and HC samples, while depressed data points are positioned on the right end of the figure. This visualization creates a spectrum from AD to MCI, to healthy samples and also, well-separated depressed data points from the rest of the samples.

Interestingly, the optimal number of K-Means clusters in t-SNE is exactly equal with 4 (the total number of disease labels in our data set), which is our ideal case. In addition, its Silhouette score is higher than both PCA and LDA methods. Figure 3a illustrates how well the disease clusters are separated in this model.

In Table 2, we observe that UMAP demonstrates

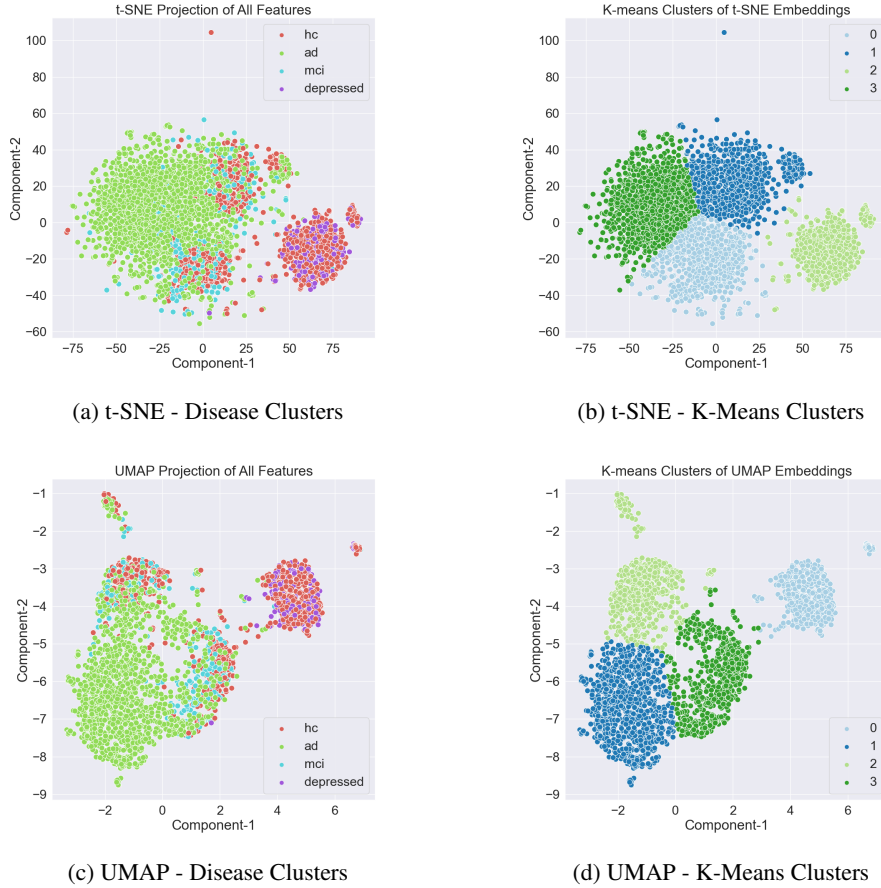


Figure 3: Pairwise scatter plots of the non-linear dimensionality reduction techniques (Component-1 vs Component-2). Left figures: 2-D representation of the samples colored based on their diagnosis labels. Right figures: 2-D representation of the samples colored based on the data-driven clusters resulted from K-Means clustering for $K=4$.

the best performance among all clustering techniques according to its optimal number of K-Means clusters and Silhouette score. Its optimal number of clusters determined by elbow method is exactly the same as the original number of diagnosis labels. In addition, its Silhouette score is higher than other approaches meaning that the level of separability of the data-driven disease clusters is higher in UMAP. The associated cluster visualizations for UMAP are also depicted in Figure 3c. We see depressed samples are well-separated from AD, and MCI, although AD and MCI themselves are not easily distinguishable.

In summary, linear dimensionality reduction techniques like PCA and LDA transform the data to a low-dimensional space as a linear combination of the original variables, while non-linear techniques are applied when the original high-dimensional data contains non-linear relationships (Sumithra and Surendran, 2015). Consequently, our findings suggest that the linearity assumption might be in-

correct for our aggregated dataset and hence, this can be another reason why the non-linear dimensionality reduction techniques outperformed the linear ones.

5.2 Analysis of the Differentiating Feature Categories between the Disease Clusters

As it is illustrated in Figure 3b, K-Means clustering derived four distinct disease clusters in a data-driven way using t-SNE embeddings. Cluster 2 corresponds to the right-most cluster in Figure 3a, which is a mixture of Depr and HC samples. Cluster 3 associates with the AD green clump of points on the left-most side of Figure 3a and clusters 0 and 1 match with the two zones in the middle comprising a combined set of AD, MCI, and HC data points. We randomly selected 10 HC samples from three distinct clusters 0, 1, and 2. We also picked 10 random Depr points from cluster 0 and 10 random AD points from cluster 3. For each instance, we applied LIME-for-t-SNE to explain its local

| Compare | Acoustic | Syntactic Complexity | Discourse Mapping | Lexical Complexity and Richness | Information Content Units | Sentiment | Word Finding Difficulty | Coherence (Global and Local) | Utterance Cohesion |
|---------------|----------|----------------------|-------------------|---------------------------------|---------------------------|-----------|-------------------------|------------------------------|--------------------|
| AD vs HC | x | - | x | x | x | - | x | x | - |
| Depr vs HC | x | - | - | x | - | x | - | x | - |
| HC Variations | x | - | - | x | - | x | - | x | - |
| AD vs Depr | x | - | x | x | - | - | x | x | - |

Table 3: Pairwise Mann-Whitney U-Test on frequency vectors of disease groups. For each pair of disease groups, the feature categories with p-value < 0.05 are marked as differentiating symptoms.

| Feature Set | Precision | Recall | F1 Score | Accuracy |
|-------------|-----------------------------------|-----------------------------------|-----------------------------------|-----------------------------------|
| F | 0.88 ± 0.04 | 0.86 ± 0.04 | 0.87 ± 0.04 | 0.90 ± 0.02 |
| F_d | 0.90 ± 0.03 | 0.88 ± 0.03 | 0.89 ± 0.03 | 0.92 ± 0.02 |
| $F - F_d$ | 0.74 ± 0.06 | 0.69 ± 0.07 | 0.71 ± 0.06 | 0.82 ± 0.03 |

Table 4: Performance of AD vs MCI vs Depr classification using different feature sets. Here, F denotes all hand-crafted acoustic and linguistic features. F_d denotes differentiating feature categories between AD and Depr. $F - F_d$ denotes all features excluding differentiating feature categories.

neighbourhood and derive its frequency vector of feature categories (See Section 4.2). Overall group comparison using Kruskal-Wallis H-Test on the frequency vectors represents that the feature categories including acoustic, lexical complexity and richness, and coherence are significantly different (with p-value < 0.05) across the disease groups including AD, Depr, and different variations of HC.

As post-hoc group comparison, we applied pairwise Mann-Whitney U-test on each pair of disease groups to assess what feature categories are the main differentiating symptoms across the disease clusters. As it is shown in Table 3, acoustic, lexical complexity and richness, sentiment, and coherence are significantly different across different variations of HC. These differences show variations within the group of healthy samples that can root in the data origin, and data collection procedures.

Our results indicate that some samples labeled as Depr are similar to HC samples across all the feature categories. This can be due to the distribution of PHQ-9 scores in DEPAC+ dataset with the majority of samples with scores in the range of 5 to 14 from mild to moderate levels of depression severity. Minor levels of depression does not meet the full criteria of major depressive disorder and the symptoms of minor forms of depression are less severe compared to major depressive disorder (Shin et al., 2021). This increases the risk of confusing modest rates of depression with control samples (Cummins et al., 2015).

Acoustic, discourse mapping (repetitiveness or circularity of speech), lexical complexity and rich-

ness, word finding difficulty, and coherence are found to be the main differentiating symptoms between AD and Depr disease clusters. To investigate the effectiveness of our results, we used these feature categories as a feature selection method.

5.3 Change in Classification Performance

We reported the performance of classification of AD vs MCI vs Depr in Table 4. According to paired sample t-test, the expected value of the accuracy, precision, recall, and F1 scores across 10 folds are significantly different between each pair of F , F_d , and $F - F_d$ feature sets, with p-value < 0.05 . Compared to when using all the features, feature selection using only the differentiating feature categories significantly improved the classification performance in terms of all metrics. Also, excluding the differentiating feature categories significantly worsened the performance of the model in classifying the diseases. These observations support that our proposed method shows a promising avenue toward detecting the data-driven symptoms that can successfully differentiate between Depr, AD, and MCI diseases.

6 Conclusion

In this work, we generate a novel aggregated dataset composed of a number of speech corpora including a combination of different clinical conditions (e.g., AD, MCI, HC, and Depr). We extract a hand-crafted set of acoustic and linguistic features derived from speech data, which are used as model predictors for discriminating between the diagnosis labels and we categorize these features under data-driven feature categories in line with the clinical symptoms of these diseases. We cluster the samples into distinguishable disease clusters and examine what speech symptoms are the main differentiating factors between the diseases. Based on our findings, non-linear clustering approaches outperform the linear ones in terms of distinguishing between distinct disease clusters. Our results

signify that acoustic abnormality, repetitiveness, or circularity of speech, word finding difficulty, coherence, and differences in lexical complexity and richness are the main differentiating symptoms between different types of dementia (e.g., MCI and AD), and depression.

References

- R Michael Bagby, Andrew G Ryder, Deborah R Schuller, and Margarita B Marshall. 2004. The hamilton depression rating scale: has the gold standard become a lead weight? *American Journal of Psychiatry*, 161(12):2163–2177.
- Aparna Balagopalan, Benjamin Eyre, Jessica Robin, Frank Rudzicz, and Jekaterina Novikova. 2021. Comparing pre-trained and feature-based models for prediction of alzheimer’s disease based on speech. *Frontiers in aging neuroscience*, 13:635945.
- Aparna Balagopalan and Jekaterina Novikova. 2021. Comparing Acoustic-based Approaches for Alzheimer’s Disease Detection. *INTERSPEECH 2021*.
- 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.
- Renesh Bedre. 2021. [Mann-whitney u test \(wilcoxon rank sum test\) in python \[pandas and scipy\]](#).
- Ashutosh Bhardwaj. 2020. [Silhouette coefficient](#).
- Adrien Bibal, Viet Minh Vu, Géraldin Nanfack, and Benoît Frénay. 2020. Explaining t-sne embeddings locally by adapting lime. In *ESANN*, pages 393–398.
- 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.
- Étienne Brunet et al. 1978. *Le vocabulaire de Jean Giraudoux structure et évolution*. Slatkine.
- Marc Brysbaert and Boris New. 2009. Moving beyond kučera and francis: A critical evaluation of current word frequency norms and the introduction of a new and improved word frequency measure for american english. *Behavior research methods*, 41(4):977–990.
- Simone Centellegher. 2020. [How to compute PCA loadings and the loading matrix with scikit-learn](#).
- Jieun Chae and Ani Nenkova. 2009. Predicting the fluency of text with shallow structural features: Case studies of machine translation and human-written text. *Association for Computational Linguistics*.
- Kenneth Ward Church. 2017. Word2Vec. *Natural Language Engineering*, 23(1):155–162.
- Bernard Croisile, Bernadette Ska, Marie-Josée Brabant, Annick Duchene, Yves Lepage, Gilbert Aimard, and Marc Trillet. 1996. Comparative study of oral and written picture description in patients with Alzheimer’s disease. *Brain and language*, 53(1):1–19.
- Nicholas Cummins, Stefan Scherer, Jarek Krajewski, Sebastian Schnieder, Julien Epps, and Thomas F Quatieri. 2015. A review of depression and suicide risk assessment using speech analysis. *Speech communication*, 71:10–49.
- Alexis Dinno. 2009. Implementing Horn’s parallel analysis for principal component analysis and factor analysis. *The Stata Journal*, 9(2):291–298.
- Allen Frances, Michael B First, and Harold Alan Pincus. 1995. *DSM-IV guidebook*. American Psychiatric Association.
- Kathleen C Fraser, Jed A Meltzer, and Frank Rudzicz. 2016a. Linguistic features identify Alzheimer’s disease in narrative speech. *Journal of Alzheimer’s Disease*, 49(2):407–422.
- Kathleen C Fraser, Frank Rudzicz, and Graeme Hirst. 2016b. Detecting late-life depression in alzheimer’s disease through analysis of speech and language. In *Proceedings of the Third Workshop on Computational Linguistics and Clinical Psychology*, pages 1–11.
- Giovanni B Frisoni, Bengt Winblad, and John T O’Brien. 2011. Revised NIA-AA criteria for the diagnosis of alzheimer’s disease: a step forward but not yet ready for widespread clinical use. *International psychogeriatrics*, 23(8):1191–1196.
- Harold Goodglass, Edith Kaplan, and Sandra Weintraub. 2001. *BDAE: The Boston Diagnostic Aphasia Examination*. Lippincott Williams & Wilkins Philadelphia, PA.
- Allan David Gordon. 1999. *Classification*. CRC Press.
- Antony Honoré et al. 1979. Some simple measures of richness of vocabulary. *Association for literary and linguistic computing bulletin*, 7(2):172–177.
- Grant L Iverson, Brian L Brooks, Travis White, and Robert A Stern. 2008. Neuropsychological assessment battery: Introduction and advanced interpretation. *The neuropsychology handbook*, pages 279–343.
- Alan Julian Izenman. 2013. Linear discriminant analysis. In *Modern multivariate statistical techniques*, pages 237–280. Springer.
- Kara Jaeschke, Fahmy Hanna, Suhailah Ali, Neerja Chowdhary, Tarun Dua, and Fiona Charlson. 2021. Global estimates of service coverage for severe mental disorders: findings from the WHO Mental Health Atlas 2017–addendum. *Global Mental Health*, 8.

- Amos D Korczyn and Ilan Halperin. 2009. Depression and dementia. *Journal of the neurological sciences*, 283(1-2):139–142.
- Kurt Kroenke, Robert L Spitzer, and Janet BW Williams. 2001. The PHQ-9: validity of a brief depression severity measure. *Journal of general internal medicine*, 16(9):606–613.
- William H Kruskal and W Allen Wallis. 1952. Use of ranks in one-criterion variance analysis. *Journal of the American statistical Association*, 47(260):583–621.
- Nilesh Kulkarni. 2018. Use of complexity based features in diagnosis of mild Alzheimer disease using EEG signals. *International Journal of Information Technology*, 10(1):59–64.
- Hochang B Lee and Constantine G Lyketsos. 2003. Depression in alzheimer’s disease: heterogeneity and related issues. *Biological psychiatry*, 54(3):353–362.
- Samantha Lomuscio. 2021. [Getting started with the kruskal-wallis test.](#)
- Daniel M Low, Kate H Bentley, and Satrajit S Ghosh. 2020. Automated assessment of psychiatric disorders using speech: A systematic review. *Laryngoscope Investigative Otolaryngology*, 5(1):96–116.
- Xiaofei Lu. 2010. Automatic analysis of syntactic complexity in second language writing. *International journal of corpus linguistics*, 15(4):474–496.
- Saturnino Luz, Fasih Haider, Sofia de la Fuente, Davida Fromm, and Brian MacWhinney. 2020. Alzheimer’s dementia recognition through spontaneous speech: the ADReSS challenge. *arXiv preprint arXiv:2004.06833*.
- Brian MacWhinney. 2014. *The CHILDES project: Tools for analyzing talk, Volume II: The database.* Psychology Press.
- Henry B Mann and Donald R Whitney. 1947. On a test of whether one of two random variables is stochastically larger than the other. *The annals of mathematical statistics*, pages 50–60.
- Leland McInnes, John Healy, and James Melville. 2018. Umap: Uniform manifold approximation and projection for dimension reduction. *arXiv preprint arXiv:1802.03426*.
- Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013. Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems*, pages 3111–3119.
- Bahman Mirheidari, Daniel Blackburn, Traci Walker, Annalena Venneri, Markus Reuber, and Heidi Christensen. 2018. Detecting signs of dementia using word vector representations. In *Interspeech*, pages 1893–1897.
- Natalia B Mota, Nivaldo AP Vasconcelos, Nathalia Lemos, Ana C Pieretti, Osame Kinouchi, Guillermo A Cecchi, Mauro Copelli, and Sidarta Ribeiro. 2012. Speech graphs provide a quantitative measure of thought disorder in psychosis. *PLoS one*, 7(4):e34928.
- Laura L Murray. 2010. Distinguishing clinical depression from early alzheimer’s disease in elderly people: can narrative analysis help? *Aphasiology*, 24(6-8):928–939.
- Kamil Mysiak. 2020. [Explaining k-means clustering.](#)
- Ziad S Nasreddine, Natalie A Phillips, Valérie Bédirian, Simon Charbonneau, Victor Whitehead, Isabelle Collin, Jeffrey L Cummings, and Howard Chertkow. 2005. The Montreal Cognitive Assessment, MoCA: a brief screening tool for mild cognitive impairment. *Journal of the American Geriatrics Society*, 53(4):695–699.
- Rupal Patel and Kathryn Connaghan. 2014. Park play: A picture description task for assessing childhood motor speech disorders. *International Journal of Speech-Language Pathology*, 16(4):337–343.
- F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. 2011. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830.
- Jeffrey Pennington, Richard Socher, and Christopher D Manning. 2014. Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, pages 1532–1543.
- Benjamin Pope, Thomas Blass, Aron W Siegman, and Jack Raheer. 1970. Anxiety and depression in speech. *Journal of Consulting and Clinical Psychology*, 35(1p1):128.
- María Luisa Barragán Pulido, Jesús Bernardino Alonso Hernández, Miguel Ángel Ferrer Ballester, Carlos Manuel Travieso González, Jiří Mekyska, and Zdeněk Smékal. 2020. Alzheimer’s disease and automatic speech analysis: a review. *Expert systems with applications*, 150:113213.
- Darrel A Regier, Emily A Kuhl, and David J Kupfer. 2013. The dsm-5: Classification and criteria changes. *World psychiatry*, 12(2):92–98.
- Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. 2016. " why should i trust you?" explaining the predictions of any classifier. In *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*, pages 1135–1144.
- Brian Richards. 1987. Type/token ratios: What do they really tell us? *Journal of child language*, 14(2):201–209.

- Peter J Rousseeuw. 1987. Silhouettes: a graphical aid to the interpretation and validation of cluster analysis. *Journal of computational and applied mathematics*, 20:53–65.
- Frank Rudzicz, Graeme Hirst, and Pascal van Lieshout. 2012. Vocal tract representation in the recognition of cerebral palsied speech. *Journal of speech, language, and hearing research : JSLHR*.
- Daun Shin, Won Ik Cho, C Hyung Keun Park, Sang Jin Rhee, Min Ji Kim, Hyunju Lee, Nam Soo Kim, and Yong Min Ahn. 2021. Detection of minor and major depression through voice as a biomarker using machine learning. *Journal of Clinical Medicine*, 10(14):3046.
- Hans Stadthagen-Gonzalez and Colin J Davis. 2006. The Bristol norms for age of acquisition, imageability, and familiarity. *Behavior research methods*, 38(4):598–605.
- V Sumithra and Subu Surendran. 2015. A review of various linear and non linear dimensionality reduction techniques. *Int. J. Comput. Sci. Inf. Technol*, 6(3):2354–2360.
- Mashrura Tasnim, Malikeh Ehghaghi, Brian Diep, and Jekaterina Novikova. 2022. Depac: a corpus for depression and anxiety detection from speech. In *Proceedings of the Eighth Workshop on Computational Linguistics and Clinical Psychology*, pages 1–16.
- Lilian Thorpe. 2009. Depression vs. dementia: how do we assess. *The Canadian Review of Alzheimer’s Disease and Other Dementias*, 12(3):17–21.
- Mary C Tierney, Rory H Fisher, Anthony J Lewis, Maria L Zorzitto, W Gary Snow, David W Reid, and Paula Nieuwstraten. 1988. The nincds-adrda work group criteria for the clinical diagnosis of probable alzheimer’s disease: A clinicopathologic study of 57 cases. *Neurology*, 38(3):359–359.
- Laurens Van der Maaten and Geoffrey Hinton. 2008. Visualizing data using t-SNE. *Journal of machine learning research*, 9(11).
- Amy Beth Warriner, Victor Kuperman, and Marc Brysbaert. 2013. Norms of valence, arousal, and dominance for 13,915 english lemmas. *Behavior research methods*, 45(4):1191–1207.
- Martin Wattenberg, Fernanda Viégas, and Ian Johnson. 2016. [How to use t-SNE effectively](#). *Distill*.
- Haydée F Wertzner, Solange Schreiber, and Luciana Amaro. 2005. Analysis of fundamental frequency, jitter, shimmer and vocal intensity in children with phonological disorders. *Revista Brasileira de Otorrinolaringologia*, 71:582–588.
- Svante Wold, Kim Esbensen, and Paul Geladi. 1987. Principal component analysis. *Chemometrics and intelligent laboratory systems*, 2(1-3):37–52.
- Pingping Wu, Ruihao Wang, Han Lin, Fanlong Zhang, Juan Tu, and Miao Sun. 2022. Automatic depression recognition by intelligent speech signal processing: A systematic survey. *CAAI Transactions on Intelligence Technology*.
- Victor H Yngve. 1960. A model and an hypothesis for language structure. *Proceedings of the American philosophical society*, 104(5):444–466.
- Chunhui Yuan and Haitao Yang. 2019. Research on k-value selection method of k-means clustering algorithm. *J*, 2(2):226–235.

A List of the features

Detailed description of the linguistic and acoustic variables in our conventional feature set is represented respectively in Table 5 and Table 6.

B Data Preprocessing

B.1 Standardization

In the data preprocessing step, the features with constant values were removed and then, the feature values were standardized by removing the mean and scaling to unit variance. The standard score of a sample x was calculated as:

$$y = \frac{x - \mu}{\sigma} \quad (1)$$

here μ and σ are the mean and standard deviation of the sample x in all training samples.

B.2 Feature selection

To remove multicollinearity, one of each pair of the features with Pearson correlation higher than 0.9 was removed.

C Implementation and Hyperparameter Setting of the Dimensionality Reduction Models

C.1 Linear Approaches

PCA: PCA was implemented by the `sklearn.decomposition.PCA` package in Scikit-learn (Pedregosa et al., 2011) and its number of components was set to the optimal number of Principal Components (PCs) calculated by Horn’s parallel analysis (Dinno, 2009), which was equal to 46. After sorting the PCs based on their explained variance ratio, the feature loadings (Centellegher, 2020) were calculated to measure the correlation between the features and the low-dimensional components. According to the distribution of the feature loadings, features with absolute value of loadings ≥ 0.4 were selected as the highly-correlated features in each PC. The number of components with the largest values of explained variance ratio and at least one highly-correlated feature was chosen as the optimal number of components. As a result, the tuned number of components was equal with 8. This approach selects the components which explain the most variance in data and include features which are highly-correlated with PCs on a linear scale.

LDA: LDA was implemented by the `LinearDiscriminantAnalysis` package of Scikit-learn (Pedregosa et al., 2011) with its default parameter settings. The number of components was set equal to 3 (the maximum allowed value), which is the number of classes-1, to achieve the highest total explained variance ratio. The classes represent the diagnosis labels in our study including HC, AD, Depr, and MCI.

C.2 Non-linear Approaches

t-SNE: t-SNE (Van der Maaten and Hinton, 2008) was implemented by the `sklearn.manifold.TSNE` package of Scikit-learn (Pedregosa et al., 2011). Perplexity was tuned by grid search to obtain the highest Silhouette score (See Section 4.1.1) in K-Means clustering trained on the t-SNE embeddings. The rest of the hyper-parameters were left unchanged with their default values. We used perplexity=30 to preserve both local and global structure of the given data to an adequate level (Wattenberg et al., 2016), in line with the recommended range of perplexity values by Van der Maaten and Hinton (2008). The number of components in t-SNE was manually tuned to 2, which was the best performing one based on the Silhouette score metric (See Section Section 4.1.1).

UMAP: This algorithm was implemented using the original UMAP⁴ library. Among different combinations of parameter settings, grid search indicated that number of components=2, the number of neighbours=50, and minimum distance=0.1 obtained the highest Silhouette score (See Section 4.1.1) in K-Means clustering trained on the UMAP embeddings. The remaining parameters were set to their default values.

⁴<https://umap-learn.readthedocs.io/en/latest/>

Linguistic Features

| Feature Category | #Features | Brief Description |
|---------------------------------|-----------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Syntactic complexity | 143 | <p>Constituency-parsing based features: Scores based on the parse tree (Chae and Nenkova, 2009) (e.g., the height of the tree, the statistical functions of Yngve depth (a measure of embeddedness) (Yngve, 1960), and the frequencies of various production rules (Chae and Nenkova, 2009)).</p> <p>Lu’s syntactic complexity features: Metrics of syntactic complexity suggested by Lu (2010) such as the length of sentences, T-units, and clauses, etc.</p> <p>Utterance length: Statistical functionals of utterance length.</p> |
| Lexical complexity and richness | 103 | <p>Grammatical constituents: The constituents of the parse tree represented in a collection of context-free grammar variables.</p> <p>Vocabulary richness: Type-token ratios; brunet (Brunet et al., 1978); Honore’s statistic (Honoré et al., 1979).</p> <p>Lexical norm-based: Average norms across all words, verbs only, and nouns only for imageability, age of acquisition, familiarity (Stadthagen-Gonzalez and Davis, 2006) and frequency (Brysbaert and New, 2009).</p> |
| Discourse mapping | 18 | <p>Utterance distances quantifying the utterance similarity via distance metrics and speech-graph (Mota et al., 2012) features based on the graph representation of the transcripts.</p> |
| Global coherence | 15 | <p>Statistical functionals of cosine distance between GloVe (Pennington et al., 2014) word embeddings of each utterance and its nearest content unit centroid utterances.</p> |
| Local coherence | 15 | <p>Statistical functionals of the similarity between Word2Vec (Mikolov et al., 2013) embeddings of the successive utterances.</p> |
| Word finding difficulty | 11 | <p>Pauses and fillers: Variables like hesitation, speech rate, word duration, and number of filled and unfilled pauses as markers of difficulty in finding words resulting in less fluent speech (Pope et al., 1970).</p> <p>Invalid words: The proportion of words not in the English dictionary (NID).</p> |
| Information units | 10 | <p>The number of information content units including objects, subjects, locations, and actions applied to quantify the number of items correctly named through the picture description task.</p> |
| Sentiment | 9 | <p>Valence, arousal, and dominance scores for all words and word types describing the sentiment of the words used (Warriner et al., 2013).</p> |
| Utterance cohesion | 1 | <p>Proportion of the number of switches in verb tense across utterances.</p> |

Table 5: List of all hand-curated linguistic features derived from transcripts. The number of features in each feature category is indicated in the second column (titled ‘#Features’).

Spectral and Energy Related Features

| Feature | #Features | Brief Description |
|-------------------------------------------------|-----------|----------------------------------------------------------------------------------------|
| Mel-Frequency Cepstral Coefficients (MFCC) 0-12 | 168 | Statistical functionals of 42 MFCC coefficients. |
| Intensity | 8 | Statistical functionals of the perceived loudness in <i>dB</i> (auditory model based). |
| Zero-Crossing Rate (ZCR) | 4 | Statistical functionals of zero crossing rate across all the voiced frames. |

Voicing Related Features

| | | |
|---------------------------------|----|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Harmonic-to-Noise Ratio (HNR) | 12 | Statistical functionals of the degree of acoustic periodicity in dB using both auto-correlation and cross-correlation methods. |
| Jitter and Shimmer | 11 | Jitter indicates the variability or perturbation of fundamental frequency, while shimmer refers to the same perturbation, but it is related to the amplitude of sound wave, or intensity of vocal emission (Wertzner et al., 2005). |
| Pauses and Fillers | 8 | Number and duration of short, medium, and long pauses, fillers (um, uh), mean pause duration, and pause-to-speech ratio. |
| Fundamental Frequency (F_0) | 6 | Statistical functionals of the fundamental frequency in Hz. |
| Durational features | 2 | Total sample and speech duration in the audio record. |
| Phonation Rate | 1 | Number of voiced samples over the total number of samples. |

Table 6: List of all hand-curated acoustic features derived from audio records. The number of features in each feature category is indicated in the second column (titled '#Features').