

Explainable Identification of Dementia From Transcripts Using Transformer Networks

Loukas Ilias  and Dimitris Askounis

Abstract—Alzheimer’s disease (AD) is the main cause of dementia which is accompanied by loss of memory and may lead to severe consequences in peoples’ everyday life if not diagnosed on time. Very few works have exploited transformer-based networks and despite the high accuracy achieved, little work has been done in terms of model interpretability. In addition, although Mini-Mental State Exam (MMSE) scores are inextricably linked with the identification of dementia, research works face the task of dementia identification and the task of the prediction of MMSE scores as two separate tasks. In order to address these limitations, we employ several transformer-based models, with BERT achieving the highest accuracy accounting for 87.50%. Concurrently, we propose an interpretable method to detect AD patients based on siamese networks reaching accuracy up to 83.75%. Next, we introduce two multi-task learning models, where the main task refers to the identification of dementia (binary classification), while the auxiliary one corresponds to the identification of the severity of dementia (multiclass classification). Our model obtains accuracy equal to 86.25% on the detection of AD patients in the multi-task learning setting. Finally, we present some new methods to identify the linguistic patterns used by AD patients and non-AD ones, including text statistics, vocabulary uniqueness, word usage, correlations via a detailed linguistic analysis, and explainability techniques (LIME). Findings indicate significant differences in language between AD and non-AD patients.

Index Terms—Alzheimer’s disease, dementia, BERT, multi-task learning, LIME.

I. INTRODUCTION

ALZHEIMER’S disease (AD) constitutes a neurodegenerative disease characterized by a progressive cognitive decline and is the leading cause of dementia. Signs of dementia include amongst others: problems with short-term memory, keeping track of a purse or wallet, paying bills, planning and preparing meals, remembering appointments, or travelling out of the neighborhood [1]. Because of the fact that Alzheimer’s dementia gets worse over time, it is important to be diagnosed early. For this reason, several research works have been introduced targeting at diagnosing dementia, which use imaging

Manuscript received 24 September 2021; revised 19 March 2022; accepted 26 April 2022. Date of publication 5 May 2022; date of current version 9 August 2022. (Corresponding author: Loukas Ilias.)

The authors are with the Decision Support Systems Laboratory, School of Electrical and Computer Engineering, National Technical University of Athens, 15780 Athens, Greece (e-mail: liliask@epu.ntua.gr; askous@epu.ntua.gr).

Digital Object Identifier 10.1109/JBHI.2022.3172479

techniques [2], CSF biomarkers [3], [4], or EEG signals [5]. Due to the fact that dementia affects speech to a high degree, recently the research has moved towards dementia identification from spontaneous speech, where several shared tasks [6], [7] have been developed in order to distinguish AD from non-AD patients.

Several research works have been conducted with regard to the identification of AD patients using speech and transcripts. The majority of them have employed feature extraction techniques [8]–[12], in order to train traditional Machine Learning (ML) algorithms, such as Logistic Regression, k-NN, Random Forest, etc. However, feature extraction constitutes a time-consuming procedure achieving poor classification results and often demands some level of domain expertise. Recently, researchers introduce deep learning architectures [13], [14], such as CNNs and BiLSTMs, so as to improve the classification results. Despite the success of transformer-based models in several domains, their potential has not been investigated to a high degree in the task of dementia identification from transcripts, where research works [15] having proposed them, use their outputs as features to train shallow machine learning algorithms. Concurrently, all research works except one [16], train machine learning models, in order to distinguish AD patients from non-AD patients, without taking into account the severity of dementia via Mini-Mental State Exam (MMSE) scores. Motivated by this limitation, we propose two multi-task learning models minimizing the loss of both dementia identification and its severity.

At the same time, to the best of our knowledge, the research works that have proposed deep learning models based on transformer networks have focused their interest only on improving the classification results obtained by CNNs, BiLSTMs etc. instead of exploring possible explainability techniques. Specifically, due to the fact that deep learning models are considered black boxes, it is important to propose ways of making them interpretable, since it is imperative for a clinician to be informed why the specific deep neural network classified a person as AD patient or not. To the best of our knowledge, only one work [17] has experimented with interpreting its proposed deep learning model (CNN-LSTM model) in the field of dementia detection using transcripts. In order to tackle this limitation, our contribution is twofold. First, we propose an interpretable neural network architecture. Next, we extend prior work and employ LIME [18], a model agnostic framework for interpretability, aiming to explain the predictions made by our best performing model. Concurrently, we propose an in-depth analysis of the language patterns used between AD and non-AD patients aiming

to shed more light on the main differences observed in the vocabulary that may distinguish people suffering from dementia from healthy people.

Our main contributions can be summarized as follows:

- We employ several transformer-based models, pretrained in biomedical and general corpora, and compare their performances.
- We propose an interpretable method based on the siamese neural networks along with a co-attention mechanism, so as to detect AD patients.
- We introduce two models in a multi-task learning framework, where the one task is the identification of dementia and the second one is the detection of MMSE score (severity of dementia). We model the MMSE detection task as a multiclass classification task instead of a regression task.
- We perform a thorough linguistic analysis regarding the differences in language between control and dementia groups.
- We employ LIME, in order to explain the predictions of our best performing model.

II. RELATED WORK

A. Feature-Based

The authors in [19], [20] introduced approaches based on multimodal data (both linguistic and acoustic features) to detect AD patients (binary classification task) and predict MMSE score (regression task). More specifically, the authors in [19] exploited dimensionality reduction techniques followed by machine learning classifiers and stated that Logistic Regression (LR) with language features was their best performing model in terms of classifying AD and non-AD patients. With regards to estimating the MMSE score, they claimed that a Random Forest classifier with language features achieves the lowest RMSE and R^2 scores. The combination of linguistic and acoustic features did not perform well on both tasks. In [20], the authors trained both shallow and deep learning models (LSTM and CNN) on a feature set consisting of acoustic features (*i*-vectors, *x*-vectors) and text features (word vectors, BERT embeddings, LIWC features, and CLAN features) to detect AD patients. They found that the top-performing classification models were the Support Vector Machine (SVM) and Random Forest classifiers trained on BERT embeddings, which both achieved an accuracy of 85.4% on the test set. Regarding the regression task, they claimed that the gradient boosting regression model using BERT embeddings outperformed all the other introduced architectures. Authors in [15] trained shallow machine learning algorithms (Logistic Regression and Support Vector Machine for detecting AD patients, and Support Vector Machines based regression and Partial Least Squares Regressor for predicting the MMSE scores) using embeddings extracted by transformer-based models, namely BERT, RoBERTa, DistilBERT, DistilRoBERTa, and BioMed-RoBERTa-base. A similar approach was conducted by [21], where the authors extracted embeddings for each word of the transcript using transformer-based networks, exploited four types of pooling functions for generating a transcript-level

representation, and trained a Logistic Regression classifier. Research work [22] merged acoustic (*x*-vectors) and linguistic features and trained a Support Vector Machine Classifier. In terms of the language features, (i) a Global Maximum pooling, (ii) a bidirectional LSTM-RNNs provided with an attention module, and (iii) the second model augmented with part-of-speech (POS) embeddings were trained on the top of a pretrained BERT model. Nasreen *et al.* [11] extracted two feature sets, namely disfluency and interactional features, and performed an in-depth statistical analysis in an attempt to investigate the differences between AD and non-AD subjects in terms of these features. Findings show that these two groups of people present significant differences. Then, they exploited shallow machine learning algorithms using the aforementioned feature sets to distinguish AD from non-AD patients and obtained an accuracy of 0.90 when providing both feature sets as input to the SVM classifier.

B. Deep Learning

Research works [23], [24] employed a hierarchical attention neural network to detect AD patients. More specifically, the authors in [23] evaluated their proposed model in both manual and automatic transcripts and found that a hierarchical neural network achieves an improvement in F1-score in comparison to other deep learning models. In [24], the authors tried to interpret the decisions made by the proposed model by visualizing words and sentences and performing statistical analyses. However, they were not able to explain why their model pays attention to some specific words more than others. Moreover, an explainable approach was introduced by [17]. Specifically, after proposing three deep learning architectures based on CNNs and RNNs, the authors applied visualization techniques and showed which linguistic characteristics are indicative of dementia, i.e., short answers, repeated requests for clarification, and interjections at the start of each utterance. Authors in [25] proposed a multi-task learning framework (Sinc-CLA), so as to predict age and MMSE scores (both considered as regression tasks) and used only speech as input for their proposed network. Concurrently, they introduced shallow networks with input *i*-vectors and *x*-vectors both in single and multi-task learning frameworks. They claimed that using *x*-vectors in a multi-task learning framework yields the best results in terms of the estimation of both age and MMSE scores. Ref. [26] introduced both feature-based and transformer-based methods. Regarding transformer-based models, they fine-tuned the BERT model to detect AD patients achieving better evaluation results than the ones achieved via the feature-based methods. For estimating the MMSE score they proposed only feature-based approaches. Research work [16] is the most similar to ours. The authors proposed transformer-based models using text, audio, and images (they converted audio to images using Mel Frequency Cepstral Coefficient). Regarding text, they employed BERT and Longformer. They claimed that models using only text data outperformed all the other proposed ones. The fusion of text and audio did not achieve better results. They introduced also a multi-task learning architecture using only text as input, in order to predict the MMSE score (regression task) and detect AD patients (binary classification

task). Results showed limited improvements in classification and a negative impact in regression. We extend this research work by employing more transformer-based networks with an efficient training strategy, proposing a new interpretable method to detect AD patients based on siamese networks, introducing two models in a multi-task learning framework by regarding the MMSE prediction task as a multiclass classification task and employing explainability techniques. On the other hand, research works [27] & [28] introduced deep learning models including CNNs and LSTM neural networks with feed-forward highway layers respectively. In [27] results suggested that the utterances of the interviewer boost the classification performance. A similar methodology with [28] was proposed by [29], where the authors exploited both BERT and LSTMs with gating mechanism and showed that LSTM with gating mechanism outperforms BERT model with gating mechanism. They stated that this difference may be attributable to the fact that BERT is very large in comparison to the LSTM models. Researchers in [30] introduced four approaches for detecting AD patients. Specifically, they trained a hierarchical neural network with an attention mechanism on linguistic features. Concurrently, they proposed a Siamese Neural Network and a Convolutional Neural Network using audio waveforms. Finally, they extracted features from audio segments and trained an SVM classifier. Results showed that the combination of audio features, CNNs, and hierarchical neural network achieved the best classification results.

C. Related Work Review Findings

From the aforementioned research works, it is evident that despite the negative consequences dementia has in people's everyday life, little work has been done so far towards its identification. More specifically, most researchers introduce feature extraction approaches from audio and transcripts and train ML algorithms, such as SVM, LR, etc. Because of the fact that feature extraction constitutes a time-consuming procedure and does not generalize well to new AD patients, researchers have started exploiting deep learning methods, such as CNNs and LSTMs, which obtain low performances. However, despite the fact that pretrained transformer models achieve new state-of-the-art results in several domains, including the biomedical one, their potential has been mainly used as embeddings for training shallow ML algorithms, such as SVM or LR. Concurrently, little has been done regarding the interpretability of the proposed deep learning models as well as the main differences observed in the language between AD patients and non-AD patients.

Our work is different from the research works mentioned above, since we: (a) propose several pretrained transformer-based models and compare their performances, (b) introduce the idea of siamese neural networks along with a co-attention mechanism towards the task of dementia classification, (c) convert the MMSE regression task into a multiclass classification one and explore if it helps dementia identification, (d) perform a detailed linguistic analysis to find the linguistic patterns that distinguish

TABLE I
MEAN AND STANDARD DEVIATION OF THE MMSE SCORES FOR THE TWO MAIN GROUPS (AD AND NON-AD PATIENTS)

	MMSE	
	mean	standard deviation
AD	17.79	5.48
non-AD	29.01	1.17

AD patients from non-AD ones, and (e) exploit LIME for explaining the predictions made by our best performing model.

III. DATASET

We use the ADReSS Challenge Dataset [6] for conducting our experiments. In contrast to other datasets, this dataset is matched for gender and age, so as to minimize the risk of bias in the prediction tasks. Moreover, it has been selected in such a way so as to mitigate biases often overlooked in evaluations of AD detection methods, including repeated occurrences of speech from the same participant (common in longitudinal datasets) and variations in audio quality. It consists of speech recordings along with their associative transcripts and includes 78 non-AD and 78 AD subjects. In addition, the dataset includes the MMSE scores for each subject except one. We report the mean and standard deviation of the MMSE scores for the two main groups, i.e., AD patients and non-AD ones, in Table I. Each participant (PAR) has been assigned by the interviewer (INV) to describe the Cookie Theft picture from the Boston Diagnostic Aphasia Exam [31]. Due to the fact that the transcripts are annotated using the CHAT coding system [32], we use the python library PyLangAcq [33] for having access to the dataset. We use data (utterances) only from PAR and conduct our experiments at the transcript-level. The ADReSS Challenge dataset has been divided into a train and a test set. The train set consists of 54 AD patients and 54 non-AD ones, while the test set consists of 24 AD patients and 24 non-AD ones.

IV. PROBLEM STATEMENT

In this section, the problem statement used in this paper is presented. More specifically, it can be divided into two problems, namely the Single-Task Learning (STL) Problem and the Multi-Task Learning (MTL) Problem, which are presented in detail in Sections IV-A and IV-B respectively.

A. Single-Task Learning Problem

Let a dataset $\mathcal{S}_{n \times 2} = \begin{bmatrix} s_1, label_1 \\ s_2, label_2 \\ \vdots \\ s_n, label_n \end{bmatrix}$ consist of a set of transcrip-

tions belonging to the dementia group, $d \subset \mathcal{S}$, and a set of transcriptions belonging to the control group, $c \subset \mathcal{S}$. Furthermore, $label_i \in \{0, 1\}$, $1 \leq i \leq n$, where 0 denotes that $s_i \in c$, while 1 denotes that $s_i \in d$. The task is to identify if a transcription $s_i \in \mathcal{S}$, belongs to a person suffering from dementia, i.e., $s_i \in d$, or not, i.e., $s_i \in c$.

B. Multi-Task Learning Problem

Let a dataset $\mathcal{S}_{n \times 3} = \begin{bmatrix} s_1, label_1, mmse_1 \\ s_2, label_2, mmse_2 \\ \vdots \\ s_n, label_n, mmse_n \end{bmatrix}$ consist of a set of

transcriptions belonging to the dementia group, $d \in \mathcal{S}$, and a set of transcriptions belonging to the control group, $c \in \mathcal{S}$. Furthermore, $label_i \in \{0, 1\}$, $1 \leq i \leq n$, where 0 denotes that $s_i \in c$, while 1 denotes that $s_i \in d$. Moreover, $mmse_i$ indicates the MMSE scores. The tasks here are to identify (i) if a transcription $s_i \in \mathcal{S}$, belongs to a person suffering from dementia, i.e., $s_i \in d$, or not, i.e., $s_i \in c$, as well as (ii) to identify the MMSE scores of each person.

V. PREDICTIVE MODELS

In this section, we describe the models used for detecting AD patients. Specifically, Section V-A refers to the models employed in the single-task learning setting, whereas in Section V-B we refer to the models used for jointly learning to identify AD patients and detect the severity of dementia.

A. Single-Task Learning

1) *Transformer-Based Models*: We exploit the following transformer-based networks in our experiments: **BERT** [34], **BioBERT** [35], **BioClinicalBERT** [36], **ConvBERT** [37], **RoBERTa** [38], **ALBERT** [39], and **XLNet** [40].

Regarding our experiments, we pass each transcription through each pretrained model mentioned above. The output of each model is passed through a Global Average Pooling layer followed by two dense layers. The first dense layer consists of 128 units with a ReLU activation function and the second one has one unit with a sigmoid activation function to give the final output.

2) *Transformer-Based Models With Co-Attention Mechanism*: In this section, we present an interpretable method to differentiate AD from non-AD patients. First, we split each transcription s in the dataset into two statements of equal length (s_1 & s_2). In this way, we have to categorize a pair of statements (s_1 & s_2) into dementia or control group. To do this, we pass s_1 and s_2 through the transformer-based models mentioned in Section V-A1, i.e., BERT, BioBERT, BioClinicalBERT, ConvBERT, RoBERTa, ALBERT, and XLNet. These models can be considered as siamese in our experiments, since we make them share the same weights. Then, we implement a co-attention mechanism introduced by [41] and adopted in several studies, including [42], [43], over the two embeddings of the two statements (outputs of the transformer-based models), in order to render the entire architecture interpretable.

Formally, let $x_1^1, x_2^1, x_3^1, \dots, x_N^1$ and $x_1^2, x_2^2, x_3^2, \dots, x_T^2$ be the tokens of s_1 and s_2 respectively. These tokens are passed to the transformer-based models as described via the equations below:

$$C = model(x_1^1, x_2^1, x_3^1, \dots, x_N^1), C \in \mathbb{R}^{d \times N} \quad (1)$$

$$S = model(x_1^2, x_2^2, x_3^2, \dots, x_T^2), S \in \mathbb{R}^{d \times T} \quad (2)$$

where *model* is one of the following: BERT, BioBERT, BioClinicalBERT, ConvBERT, RoBERTa, ALBERT, and XLNet. We have omitted the first dimension, which corresponds to the batch size. Following the methodology proposed by [41], given the output of the model receiving the tokens of s_1 ($C \in \mathbb{R}^{d \times N}$) and the output of the model receiving the tokens of s_2 ($S \in \mathbb{R}^{d \times T}$), where d denotes the hidden size of the model, the affinity matrix $F \in \mathbb{R}^{N \times T}$ is calculated using the equation $F = \tanh(C^T W_l S)$, where $W_l \in \mathbb{R}^{d \times d}$ is a matrix of learnable parameters. Next, this affinity matrix is considered as a feature and we learn to predict the attention maps for both statements via the following, $H^s = \tanh(W_s S + (W_c C) F)$ and $H^c = \tanh(W_c C + (W_s S) F^T)$, where $W_s, W_c \in \mathbb{R}^{k \times d}$ are matrices of learnable parameters. The attention probabilities for each word in both statements are calculated through the softmax function as follows, $a^s = softmax(w_{hs}^T H^s)$, $a^c = softmax(w_{hc}^T H^c)$, where $a_s \in \mathbb{R}^{1 \times T}$ and $a_c \in \mathbb{R}^{1 \times N}$. $W_{hs}, W_{hc} \in \mathbb{R}^{k \times 1}$ are the weight parameters. Based on the above attention weights, the attention vectors for each statement are obtained by calculating the weighted sum of the features from each statement. Formally, $\hat{s} = \sum_{i=1}^N a_i^s s^i$, $\hat{c} = \sum_{j=1}^T a_j^c c^j$, where $\hat{s} \in \mathbb{R}^{1 \times d}$ and $\hat{c} \in \mathbb{R}^{1 \times d}$. Finally, these two vectors are concatenated, i.e., $p = [\hat{s}, \hat{c}]$, where $p \in \mathbb{R}^{1 \times 2d}$ and we pass the vector p to a dense layer with 128 units and a ReLU activation function followed by a dense layer consisting of one unit with a sigmoid activation function.

B. Multi-Task Learning

In this section we propose two architectures based on multi-task learning [44] and adopt the methodology followed by [45] & [46]. To be more precise, we employ a multi-task learning framework consisting of a primary and an auxiliary task. The identification of dementia constitutes the primary task, while the prediction of the MMSE score constitutes the auxiliary one. Our main objective is to explore whether the MMSE score helps in classifying groups into dementia or control. The introduced architectures are trained on the two tasks and updated at the same time with a joint loss:

$$L = (1 - \alpha) L_{dementia} + \alpha L_{MMSE} \quad (3)$$

where $L_{dementia}$ and L_{MMSE} are the losses of dementia identification and MMSE prediction tasks respectively. α is a hyper-parameter that controls the importance we place on each task. We mention below the MTL architectures developed.

a) *MTL-BERT (Multiclass)*: We pass each transcription through a BERT model (which constitutes our best performing STL model). The output of the BERT model is passed through two separate dense layers, so as to identify dementia and predict the MMSE score. For identifying dementia, we use a dense layer with 2 units and a softmax activation function and minimize the cross-entropy loss function. Regarding the estimation of the MMSE score, in contrast with previous research works, we convert the MMSE regression task into a multiclass classification task. More specifically, according to [28], we can create 4 groups of cognitive severity: **healthy** (MMSE score ≥ 25), **mild dementia** (MMSE score of 21–24), **moderate dementia**

(MMSE score of 10–20), and **severe dementia** (MMSE score ≤ 9). Thus, for classifying transcriptions into one of these 4 groups, we use a dense layer of 4 units with a softmax activation function and minimize the cross-entropy loss function.

b) MTL-BERT-DE (Multiclass): Similarly to [46], we pass each transcription into a BERT model. The output of the BERT model is passed through two separate BERT encoders, i.e., double encoders, which are followed by dense layers so as to identify dementia and classify MMSE score into one of the four classes mentioned above. For identifying dementia, we use a dense layer with 2 units and a softmax activation function and minimize the cross-entropy loss function. For classifying the MMSE score, we use a dense layer with 4 units and a softmax activation function and minimize the cross-entropy loss function.

VI. EXPERIMENTS

All experiments are conducted on a single Tesla P100-PCIE-16 GB GPU.

A. Single-Task Learning

Comparison with state-of-the-art approaches: We compare our introduced models with the following research works, since these research works propose single-task learning models and test their proposed approaches on the ADReSS Challenge test set: (1) Text [15], (2) LSTM with Gating (Acoustic + Lexical + Dis) [28], (3) Fusion Maj. (3-best) [30], (4) Logistic Regression (NLP) [20], (5) fastText, bi + trigram [27], (6) Attempt 5 [21], and (7) Fusion of system [22].

Experimental Setup: Firstly, we divide the train set provided by the Challenge into a train and a validation set (65%-35%). Next, we train the proposed architectures five times and test them using the test set provided by the Challenge. Specifically, we freeze the weights of each pretrained model (BERT, BioBERT, BioClinicalBERT, ConvBERT, RoBERTa, ALBERT, and XLNet) and update the weights of the rest layers. In this way, these pretrained models act as fixed feature extractors. We train the proposed architectures using Adam optimizer with a learning rate of $1e-4$. We apply *EarlyStopping* and stop training, if the validation loss has stopped decreasing for 9 consecutive epochs. We also apply *ReduceLRonPlateau*, where we reduce the learning rate by a factor of 0.2, if the validation loss has stopped decreasing for 3 consecutive epochs. When this training procedure stops, we unfreeze the weights of the pretrained models and train the entire deep learning architectures using Adam optimizer with a learning rate of $1e-5$. We apply *EarlyStopping* with a patience of 3 based on the validation loss. In terms of models with a co-attention mechanism, we start training the proposed architectures using Adam optimizer with a learning rate of $1e-3$ and follow the same methodology. We also apply dropout after the co-attention mechanism with a rate of 0.4. For **BERT**, we have used the base-uncased model, for **BioBERT** we have used BioBERT v1.1 (+PubMed), for **ConvBERT** we have used the base model, for **RoBERTa** we have employed the base model, for **ALBERT** we have used the base-v1 model, and

for **XLNet** we have used the base model. For these pretrained models, we have used the Transformers library [47].¹

Evaluation Metrics: We evaluate our results using Accuracy, Precision, Recall, F1-score, and Specificity. All these metrics have been calculated using the dementia class as the positive one.

B. Multi-Task Learning

Comparison with state-of-the-art approaches: For the primary task (AD Classification task), we compare our introduced models with BERT base [16], since this research work proposes a multi-task learning model and tests its proposed approach on the ADReSS Challenge test set.

Experimental Setup: Firstly, we divide the train set provided by the Challenge into a train and a validation set (65%-35%). Next, we train the proposed architectures five times and test them using the test set provided by the Challenge. We use the Adam optimizer with a learning rate of $1e-6$. We apply *EarlyStopping* and stop training, if the validation loss has stopped decreasing for 8 consecutive epochs. Regarding MTL-BERT-DE (Multiclass), we freeze the weights of the shared BERT model. Moreover, because of the class imbalance of the MMSE categories, we apply balanced class weights to the loss function (L_{MMSE}). We set α of (3) equal to 0.1.²

Evaluation Metrics: For the primary task (AD Classification task), we evaluate our results using Accuracy, Precision, Recall, F1-score, and Specificity. All these metrics have been calculated using the dementia class as the positive one.

For the auxiliary task (MMSE Classification task), we evaluate our results using the average weighted Precision, average weighted Recall, and average weighted F1-score.

VII. RESULTS

A. Single-Task Learning Experiments

The results of the proposed models mentioned in Section V-A are reported in Table II. Also, Table II provides a comparison of our introduced models with existing research initiatives.

Regarding our proposed transformer-based models, one can easily observe that BERT obtains the highest Recall, F1-score, and Accuracy accounting for 81.66%, 86.73%, and 87.50% respectively. Specifically, BERT outperforms the other introduced transformer-based models in Recall by 1.67-13.33%, in F1-score by 2.01-10.98%, and in Accuracy by 1.25-9.17%. BioClinicalBERT achieves the second highest Accuracy and F1-score accounting for 86.25% and 84.72% respectively. Also, BioClinicalBERT obtains the highest Precision score equal to 95.03% surpassing the other transformer-based models by 4.79-15.88%. RoBERTa achieves comparable results to BERT and BioClinicalBERT yielding an Accuracy and F1-score of 84.16% and 82.81% respectively. In addition, BioBERT and ConvBERT demonstrate slight differences in Accuracy and F1-score, with

¹For BioClinicalBERT we have used the model in: https://huggingface.co/emilyalsentzer/Bio_ClinicalBERT

²We used also the experimental setup of Section VI-A. However, lower evaluation results were achieved.

TABLE II

PERFORMANCE COMPARISON AMONG PROPOSED STL MODELS AND STATE-OF-THE-ART APPROACHES ON THE ADDRESS CHALLENGE TEST SET

Architecture	Evaluation metrics				
	Prec.	Rec.	F1-score	Acc.	Spec.
Comparison with state-of-the-art approaches					
[15]	-	87.50	-	89.58	91.67
[28]	81.82	75.00	78.26	79.17	83.33
[30]	-	-	85.40	85.20	-
[20]	-	-	-	85.00	-
[27]	86.00	79.00	83.00	83.33	88.00
[21]	-	-	-	85.42	-
[22]	94.12	66.67	78.05	81.25	95.83
Proposed Transformer-based models					
<i>BERT</i>	87.19	81.66	86.73	87.50	93.33
	± 3.25	± 5.00	± 4.53	± 4.37	± 5.65
<i>BioBERT</i>	86.87	78.33	82.11	82.92	87.50
	± 6.09	± 4.86	± 2.83	± 3.06	± 6.97
<i>BioClinicalBERT</i>	95.03	76.66	84.72	86.25	95.83
	± 3.03	± 4.99	± 2.74	± 2.12	± 2.64
<i>ConvBERT</i>	83.51	79.99	81.65	82.08	84.16
	± 1.23	± 4.08	± 2.06	± 1.66	± 1.66
<i>RoBERTa</i>	90.24	76.66	82.81	84.16	91.66
	± 2.81	± 4.99	± 3.52	± 2.83	± 2.64
<i>ALBERT</i>	79.15	78.33	78.45	78.33	78.33
	± 7.89	± 3.11	± 3.12	± 3.86	± 8.89
<i>XLNet</i>	85.58	68.33	75.75	78.33	88.33
	± 2.77	± 6.77	± 4.05	± 2.82	± 3.12
Proposed Transformer-based models with co-attention mechanism					
<i>BERT</i>	83.67	84.16	83.85	83.75	83.33
<i>Co-Attention</i>	± 3.36	± 1.66	± 1.09	± 1.56	± 4.56
<i>BioBERT</i>	85.41	76.66	80.72	81.66	86.66
<i>Co-Attention</i>	± 4.91	± 3.33	± 3.16	± 3.06	± 4.86
<i>BioClinicalBERT</i>	82.60	81.66	81.99	82.08	82.50
<i>Co-Attention</i>	± 3.60	± 4.25	± 2.11	± 2.12	± 4.86
<i>ConvBERT</i>	83.78	81.66	82.43	82.50	83.33
<i>Co-Attention</i>	± 6.13	± 4.24	± 2.37	± 3.12	± 8.74
<i>RoBERTa</i>	79.39	79.16	79.06	79.16	79.16
<i>Co-Attention</i>	± 2.26	± 6.45	± 2.15	± 1.32	± 4.56
<i>ALBERT</i>	77.94	84.16	80.77	80.00	75.83
<i>Co-Attention</i>	± 3.20	± 4.86	± 1.68	± 1.66	± 5.53
<i>XLNet</i>	85.63	69.16	76.42	78.75	88.33
<i>Co-Attention</i>	± 3.45	± 5.00	± 3.75	± 3.06	± 3.12

Reported values are mean \pm standard deviation. Results are averaged across five runs.

BioBERT surpassing ConvBERT in both metrics. Specifically, BioBERT surpasses ConvBERT in F1-score by 0.46% and in Accuracy by 0.84%. Moreover, we observe that ALBERT and XLNet achieve Accuracy scores equal to 78.33%, with ALBERT surpassing XLNet in F1-score by 2.70%.

Regarding our proposed transformer-based models with a co-attention mechanism, they achieve lower performance than the proposed transformer-based models except for ConvBERT+Co-Attention, ALBERT+Co-Attention, and XLNet+Co-Attention. More specifically, ConvBERT+Co-Attention presents a slight surge of 0.42% in Accuracy in comparison with ConvBERT, ALBERT+Co-Attention presents an increase in Accuracy by 1.67% in comparison with ALBERT, and XLNet+Co-Attention demonstrates a slight increase of 0.42% in Accuracy in comparison with XLNet. BERT+Co-Attention attains the highest F1-score and Accuracy accounting for 83.85% and 83.75% respectively. BERT+Co-Attention outperforms the other models in terms of F1-score by 1.42-7.43%, and in terms of Accuracy by 1.25-5.00%. ConvBERT+Co-Attention

TABLE III

PERFORMANCE COMPARISON AMONG PROPOSED MTL MODELS AND STATE-OF-THE-ART APPROACHES ON THE ADDRESS CHALLENGE TEST SET FOR THE PRIMARY TASK (AD CLASSIFICATION TASK)

Architecture	Evaluation metrics				
	Prec.	Rec.	F1-score	Acc.	Spec.
Comparison with state-of-the-art approaches					
[16]	-	-	-	80.83	-
	-	-	-	± 1.56	-
Proposed Multi-task learning models					
<i>MTL-BERT</i>	88.59	83.33	85.84	86.25	89.16
<i>(Multiclass)</i>	± 3.05	± 2.64	± 2.12	± 2.13	± 3.33
<i>MTL-BERT-DE</i>	85.19	85.00	84.96	85.00	85.00
<i>(Multiclass)</i>	± 3.46	± 5.00	± 2.60	± 2.43	± 4.25

Reported values are mean \pm standard deviation. Results are averaged across five runs.

and BioClinicalBERT+Co-Attention demonstrate slight differences in F1-score and Accuracy, with ConvBERT+Co-Attention surpassing BioClinicalBERT+Co-Attention in F1-score by 0.44% and in Accuracy by 0.42%. BioBERT+Co-Attention and ALBERT+Co-Attention achieve almost equal F1-score results, with BioBERT+Co-Attention attaining a higher Accuracy score than ALBERT+Co-Attention by 1.66%. RoBERTa+Co-Attention and XLNet+Co-Attention demonstrate low performances attaining an Accuracy of 79.16% and 78.75% respectively.

Overall, BERT constitutes our best performing model, since it outperforms all the other introduced models in F1-score and Accuracy. Although there are models surpassing BERT in Precision and Recall, BERT outperforms all of them in F1-score, which constitutes the weighted average of Precision and Recall. In addition, there are models that outperform BERT in Specificity. However, high specificity and low recall means that the model cannot diagnose the AD patients pretty well and consequently AD patients are misdiagnosed as non-AD ones.

In comparison with the state-of-the-art approaches, one can observe that our proposed models achieve comparable performance to or outperform previous studies. More specifically, BERT outperforms all the research works, except [15], in terms of Accuracy by 2.08-8.33%, in F1-score by 1.33-8.68%, and in Recall by 2.66-14.99%. Moreover, BERT+Co-Attention surpasses [22], [27], [28] in Accuracy by 2.50%, 0.42%, and 4.58% respectively. Also, it surpasses [22], [27], [28] in Recall by 17.49%, 5.16%, and 9.16% respectively. BERT+Co-Attention outperforms [22], [27], [28] in F1-score by 5.80%, 0.85%, and 5.59% respectively.

B. Multi-Task Learning Experiments

1) *Primary Task*: The results of the introduced models described in Section V-B are reported in Table III. Also, Table III provides a comparison of our introduced approaches with state-of-the-art approaches.

With regards to our introduced models, one can easily observe that MTL-BERT (Multiclass) outperforms MTL-BERT-DE (Multiclass) in terms of all the evaluation metrics except Recall. Specifically, MTL-BERT (Multiclass) surpasses MTL-BERT-DE (Multiclass) in Precision by 3.40%, in F1-score by

TABLE IV

RESULTS OF THE PROPOSED MTL MODELS ON THE ADReSS CHALLENGE TEST SET FOR THE AUXILIARY TASK (MMSE CLASSIFICATION TASK)

Architecture	Evaluation metrics		
	Avg. W. Prec.	Avg. W. Rec.	Avg. W. F1-score
Proposed Multi-task learning models			
<i>MTL-BERT</i> (<i>Multiclass</i>)	73.62 ±2.95	69.16 ±4.04	64.75 ±3.50
<i>MTL-BERT-DE</i> (<i>Multiclass</i>)	70.50 ±5.59	70.42 ±3.06	68.57 ±2.04

Reported values are mean ± standard deviation. Results are averaged across five runs.

0.88%, in Accuracy by 1.25%, and in Specificity by 4.16%. Although MTL-BERT-DE (Multiclass) surpasses MTL-BERT (Multiclass) in Recall by 1.67%, MTL-BERT (Multiclass) obtains a higher F1-score, which constitutes the weighted average of Precision and Recall. Therefore, MTL-BERT (Multiclass) constitutes our best performing model in the MTL framework.

In comparison to the research work [16], as one can easily observe, both our introduced models attain a higher Accuracy score. To be more precise, MTL-BERT (Multiclass) outperforms BERT base [16] in Accuracy by 5.42%. In addition, MTL-BERT-DE (Multiclass) surpasses the research work [16] in Accuracy by 4.17%. These differences in performance are attributable to the fact that we adopt a different training procedure than the one adopted by [16], we consider the MMSE task as a multiclass classification task instead of a regression task, as well as to the different architectures proposed.

2) *Auxiliary Task*: The results of the introduced models mentioned in Section V-B for the auxiliary task (MMSE Classification task) are reported in Table IV.

As one can easily observe, MTL-BERT (Multiclass) obtains an average weighted Precision of 73.62% surpassing MTL-BERT-DE (Multiclass) by 3.12%. However, MTL-BERT-DE (Multiclass) outperforms MTL-BERT (Multiclass) in average weighted Recall and average weighted F1-score by 1.26% and 3.82% respectively.

VIII. ANALYSIS OF THE LANGUAGE USED IN CONTROL AND DEMENTIA GROUPS

We finally perform an extensive analysis to uncover some unique characteristics, which discriminate the AD patients from the non-AD ones, and understand the predictions made by our best performing model as well as its limits.

A. Text Statistics

We first extract some statistics, namely the syllable count, the lexicon count, the difficult words, and the sentence count, using the TEXTSTAT library in Python, in order to understand better the differences in language used between control and dementia groups. More specifically, the syllable count refers to the number of syllables, the lexicon count to the number of words, and the sentence count to the number of sentences present in the given text. With regards to the difficult words, they refer to the number of polysyllabic words with a Syllable Count > 2 that are not included in the list of words of common usage in English [48].

TABLE V

MEAN ± STANDARD DEVIATION METRICS PER TRANSCRIPT

Metric	Transcript	
	Control	Dementia
Syllable Count†	151.63 ± 79.98	119.95 ± 71.18
Lexicon Count†	107.49 ± 62.02	86.08 ± 54.10
Difficult Words†	10.58 ± 3.64	6.38 ± 3.53
Sentence Count	1.67 ± 1.03	1.92 ± 1.62

† indicates statistical significance between transcripts of control and dementia groups. All differences are significant at $p < 0.05$ after benjamini-hochberg correction.

TABLE VI

JACCARD'S INDEX BETWEEN TRANSCRIPTS OF CONTROL AND DEMENTIA GROUP

Jaccard's Index between transcripts	Result
$J(\mathcal{P}=\text{control}, \mathcal{C}=\text{dementia})$	0.4049

After extracting these statistics per transcript, we calculate the mean and standard deviation for both control and dementia groups. We test for statistical significance using an independent t-test for each metric between control and dementia groups and adjust the p-values using Benjamini-Hochberg correction [49]. As one can easily observe in Table V, the control group presents a significantly higher number of syllables, lexicon, and difficult words than the dementia group.

B. Vocabulary Uniqueness

In order to understand the vocabulary similarities and differences between control and dementia groups, we adopt the methodology proposed by [50]. Formally, let \mathcal{P} and \mathcal{C} be the sets of unique words included in the control group and dementia group respectively. Next, we calculate the Jaccard's index given by (4), in order to measure the similarity between finite sample sets. More specifically, the Jaccard's index is a number between 0 and 1, where 1 indicates that the two sets, namely \mathcal{P} and \mathcal{C} , have the same elements, while 0 indicates that the two sets are completely different.

$$J(\mathcal{P}, \mathcal{C}) = |\mathcal{P} \cap \mathcal{C}| / |\mathcal{P} \cup \mathcal{C}| \quad (4)$$

As observed in Table VI, the Jaccard's index between the control and dementia groups is equal to 0.4049, which indicates that people with dementia tend to use a different vocabulary than those in the control group.

C. Word Usage

Apart from finding the vocabulary similarities and differences, it is imperative that patterns of word usage be investigated. Thus, following the methodology introduced in [50], the main objective of this section is to explore the differences between the two classes (control and dementia) with regard to the probability of using specific words more than others. Formally, let D_1 and D_2 be two documents, where D_1 includes all the transcriptions of the control group, whereas D_2 consists of transcriptions of the dementia group. Moreover, we define S as the entire corpus

TABLE VII
KULLBACK-LEIBLER DIVERGENCE

KL divergence	Result
KL(Control Dementia)	0.2047
KL(Dementia Control)	0.2161

consisting of D_1 and D_2 . Now we can define the probability of a word w_i in the document D_1 in a collection of documents S given by (5):

$$P(w_i|D_1, S) = (1 - \alpha_D)P(w_i|D_1) + \alpha_D P(w_i|S) \quad (5)$$

Similarly, we can define the probability of a word w_i in the document D_2 in a collection of documents S given by (6):

$$P(w_i|D_2, S) = (1 - \alpha_D)P(w_i|D_2) + \alpha_D P(w_i|S) \quad (6)$$

We employ the Jelinek-Mercer smoothing method and consider that $\alpha_D \in [0, 1]$. More specifically, α_D is a parameter that controls the probability of words included only in one document (D_1 or D_2). In our experiments, we set α_D equal to 0.2.

Moreover, we define $P(w_i|S) = \frac{s_{w_i}}{|S|}$, where s_{w_i} denotes the number of times a word w_i is included in the collection, whereas $|S|$ is the total number of words occurrences in the collection. Similarly, $P(w_i|D_1) = \frac{d_{w_i}}{|D_1|}$, where d_{w_i} denotes the number of times a word w_i is presented in the document D_1 , whereas $|D_1|$ is the total number of words occurrences in the document D_1 . The same methodology has been adopted for calculating the $P(w_i|D_2)$.

After having calculated the two distributions, i.e., $P(w_i|D_1, S)$ and $P(w_i|D_2, S)$, we exploit the Kullback-Leibler (KL) divergence, in order to measure the difference of these two distributions. KL-divergence is always greater than zero and is given by (7). The larger it gets, the more different the two distributions are.

$$KL(P||C) = \sum_x P(x) \log \frac{P(x)}{C(x)} \quad (7)$$

As one can easily observe in **Table VII**, the KL divergence between control and dementia groups is high indicating that these two groups present differences regarding the probability of using some words more than others. Our findings agree with the ones in [50], where the authors state that there are clear differences in terms of language use between positive (depression and self-harm) and control group, where the values of KL-divergence range from 0.18 to 0.21.

D. Linguistic Feature Analysis

Following the method introduced by [51], the main objective of this section is to shed light on which unigrams and pos-tags are mostly correlated with each class separately. To facilitate this, we compute the point-biserial correlation between each feature (unigram and pos-tag) across all the transcriptions and a binary label (0 for the control and 1 for the dementia group). Before computing the correlation, we normalize features so that they sum up to 1 across each transcription. We use the point-biserial correlation, since it is a correlation used between continuous and binary variables. It returns a value between -1 and 1. Since we are

TABLE VIII
FEATURES ASSOCIATED WITH CONTROL AND DEMENTIA SUBJECTS,
SORTED BY POINT-BISERIAL CORRELATION

Control		Dementia	
Unigrams	corr.	Unigrams	corr.
is	0.364	here	0.310
curtains	0.361	-	-
window	0.301	-	-
are	0.300	-	-

POS	corr.	POS	corr.
VBG	0.285	RB	0.388
DT	0.216	PRP	0.354
NN	0.210	VBD	0.275
-	-	UH	0.242

All correlations are significant at $P < 0.05$ after benjamini-hochberg correction.

only interested in the strength of the correlation, we compute the absolute value, where negative correlations refer to the control group (label 0) and positive correlations refer to the dementia one (label 1). We report our findings in **Table VIII**, where all correlations are significant at $p < 0.05$, with Benjamini-Hochberg correction [49] for multiple comparisons.

As one can easily observe, the pos-tags associated with the dementia group are the following: RB (adverbs), PRP (personal pronoun), VBD (verb in past tense), and UH (interjection). On the other hand, people in the control group tend to use VBG (verb, gerund, or present participle), DT (determiner), and NN (noun). These findings can be justified in **Table IX**, where we present three examples of transcripts belonging to the control group and three examples of transcripts belonging to the dementia one. More specifically, we have assigned colours to different pos-tags, so as to render the differences in the language patterns used by each group easily understandable to the reader. To be more precise, red colour indicates the VBG pos-tag, yellow refers to the DT pos-tag, fuchsia to the RB pos-tag, apricot to the PRP pos-tag, navy blue to the VBD pos-tag, and the pine green to the UH pos-tag.

We observe that people in the dementia group tend to use personal pronouns (he, she, I, them etc.) very often, since they are unable to remember the specific terms (mom, boy, etc.). This finding agrees with the research conducted by [52], where the authors state that personal pronouns present a high frequency in the speech of AD patients, since these people cannot find the target word. To be more precise, in a conversation people have to remember what they have said during the entire conversation. However, this is not possible in AD patients, who present working memory impairment and thus tend to produce empty conversational speech (use of personal pronouns). On the other hand, people in the control group tend to use more nouns instead of personal pronouns, since they are able to maintain various kinds of information.

Moreover, AD patients tend to use verbs in the past tense (were, forgot, did, started) in contrast to people who are not suffering from dementia and use verbs in the present participle. One typical example that can illustrate this difference can be seen in the fifth transcription in **Table IX**, i.e., "*oh have you heard of that new game that they started to play after christmas? did you*". The AD patient perhaps remembers a personal story from the

TABLE IX
EXAMPLES OF TRANSCRIPTS ALONG WITH THEIR LABELS

Transcript	Label
" well the girl is watching the boy go into the cookie jar . he has a cookie in his hand . he's on the stool . the stool is falling . the mother is drying dishes . has a plate in her hand . sink is overflowing . there's water on the floor . she's stepping in the water . something that's going on you said ? the little girl looks like she's motioning to the boy to be quiet . and I don't know what else . the woman's looking out the window . the window's open . "	Control
" action . alright . a lady's drying dishes . the boy was standing on a stool but the action is that the stool has slipped and he is falling . and the girl has her hand raised reaching for a cookie . and there's a lot of action in the sink here . the water is flowing out . she is apparently so daydreaming that she doesn't realize that the sink is overflowing . any more action ? or is that enough action ? "	Control
" touching lip . raising arm . is that what you mean ? reaching for cookie . handing cookie down . slipping from stool . stool falling over . wiping dishes . water running . water overflowing . breeze . I don't know if that's action . stepping out from water . I guess that's it . "	Control
" alright . I see the little boy stealing cookies from the cookie jar . and he gave some to the little girl and she's eating some of the cookies . and I guess this is mama and she's washing the dishes . and she dropped a dish . no she didn't drop a dish . the water that she's washing the dishes with she let run . and it's overflowed . that doesn't sound right . did it ? we forgot to turn off the spigot . and so the water is running off onto the floor here . and mom apparently is washing the dishes . and here's this little boy stealing the cookies . he's gonna fall because the four legged stool is gonna fall over with him and the cookie jar . and mama's drying the dishes as usual for mamas if they don't have a husband that dries them or washes them or whatever . let's see now . I guess there's more things I'm sposta see . let's see here now . oh and the water is flowing out of the sink they forgot to turn off whoever's doing the dishwashing . mom apparently here , she forgot to turn off the water and the water is spilling out onto the kitchen floor . and the little girl has pushed over the stool with the boy that was reaching up to get the cookies . either she pushed it over or he fell over with it . you know it excuse me but you know I was ... "	Dementia
" mhm . oh I see a part of the whole kitchen . is that all the kitchen or isn't it ? oh I can't read ... a lady a mother were in her kitchen . in her kitchen doing some work I suppose . and there's another woman there sharing their pleasures or whatever . oh have you heard of that new game that they started to play after christmas ? did you ? is a . well it looks like ... I'd say this is ... well let's see . it looks like ... oh ... my wife will beat me by a couple rows of this . that's like the washing machine ? or let me see . I can't ... oh that's the son come from school maybe or something . that's a youngster there . well that's just as though they getting ready to go to school or they're just coming out from school . and right there he's same as back there except for down there in the bottom I think it's ... that's a little . "	Dementia
" yes . the water ? well let's see . there's something hasta be where the water goes down over . there's probably something that's ... or they don't have it open or something might have . I don't know . what ..? when the water goes down what do you call that ? this here . right here . this . what do you call that ? what is that ? what is that ? I don't know ! that's what I'm saying . I don't know what that is . the what ? a pipe . oh water pipe ! oh yeah . okay . well then maybe the water pipe is not broke but there must be things in there . that the water will not go down . I don't know . huh ? what's happening to the water ? well the water is going down in the ... I don't know . what would you call this ? floor ! yeah okay . yeah . well down on this side of the picture . well this thing here is turning over . yeah . no , uhuh . I don't know what's going on . well he's probably getting ... what's this here ? cocoa jar ? what's this cocoa ? c o o k i e . I don't know . I don't know what ..? huh ? cookie , oh a cookie . oh ! oh okay . mhm . well he's getting it out . and he's gonna give it to the girl / down here . mhm . going on in the picture ? well the boy is giving her the girl the cookie . this probably is broke . so the water will not go down in and it's coming up and going in here huh . well it looks like she was gonna wash . what they eat with , all that . what do you call that ? what do you call this ? a plate ? oh yeah . what you eat on . is that what you call them a plate ? oh this is a cup ? oh maybe , I don't know . mhm . okay . "	Dementia

Red colour indicates the VBG pos-tag, yellow refers to the DT pos-tag, fuchsia to the RB pos-tag, apricot to the PRP pos-tag, navy blue to the VBD pos-tag, and the pine green to the UH pos-tag.

past that wants to narrate, instead of the task he has been assigned to conduct. Therefore, the patient is not able to stay focused on describing the picture. This finding is consistent with [53], [54], where the authors state that AD patients present difficulty in maintaining and continuing the development of a topic and thus demonstrate unexpected topic shifts. Also, this finding reveals a difference in language used by the AD patients and the agrammatic aphasics. Specifically, patients with agrammatic aphasia typically have problems using past tense inflection and instead rely on infinitive or present tense verb forms [55].

In addition, AD patients tend to use the UH (oh, yeah, well) and the RB (maybe, probably) pos-tags, since they are not certain of what they are describing due to the cognitive impairment. Concurrently, the UH pos-tag constitutes an example of empty speech. More specifically, this pos-tag is used as filler at the beginning of each utterance, since AD patients are thinking of what to say.

E. Explainability - Error Analysis

In this section, we employ LIME [18] (using 5000 samples) to explain the predictions made by our best performing model,

namely BERT, and shed more light regarding the differences in language between AD and non-AD patients. More specifically, LIME generates local explanations for any machine learning classifier by introducing an interpretable model, which is trained on data generated through observing differences in the classification performance when removing tokens from the input string.

Examples of explanations generated by LIME are illustrated in Figs. 1-4. More specifically, Fig. 1 illustrates two transcripts, whose ground-truth label is dementia, while our model predicts them as belonging to non-AD patients. Fig. 2 refers to transcripts with both ground-truth label and prediction corresponding to dementia. In Fig. 3, two transcripts are presented, whose prediction is control and true label is control too. Finally, Fig. 4 illustrates transcripts, which are misclassified. The ground-truth is control, whereas the prediction is dementia. Moreover, as one can observe, each token has been assigned a colour, either blue or orange. To be more precise, the blue colour indicates which tokens are indicative of the control group, whilst the orange colour indicates tokens, which are used mainly by AD patients. The more intense the colours are, the more important these tokens are towards the final classification of the transcript.

I see two kids up at the cookie jar , one on a stool the other standing on the floor . cupboard door is opened . mother washing the dishes . the water is overflowing the sink . and there two cups and a plate on the counter . and she holding a plate in her hand . curtains at the windows . the cookie jar has the lid off . that about it . cupboards underneath the sink . cupboards underneath the other cupboards . kid falling off the stool . the girl laughing at him . cookies in the cookie jar with the lid off . he has a cookie in his hand . and that it .

(a)

a little boy is getting himself some cookies out of the jar . and the stool is turning over on him that he standing on . he handing the little girl some cookies out of the jar . is that it ? well that in the first one . he handing the little girl cookies out of the cookie jar . and he on a stool and it turning over . and he handing her some cookies . and in the second picture the little girl looks as though she drying dishes . and the sink is running over on the floor . she has the spigots on . a nice view from the window . is that it ? that it .

(b)

Fig. 1. Label: Dementia, Prediction: Control.

yeah I see the woman in a kitchen . and / now it looks like she ... I can't really pick it out but ... oh and there a little girl here talking and a little boy I assume on this side here . and this is a stool here or some kind of a chair . and I don't know what this is here . I can't see what that is . oh there another . did I talk about this girl up here ? she ... I can't see too plain what she doing . oh yes I think so . where was she ? this girl ? I really can't see what she doing . no I don't . yeah , that awfully hard for me to distinguish .

(a)

hm ... it's a little boy climbing up getting some cookies out of the cookie jar . and his little sister reaching for some . and the little boy is standing on a stool . and his big sister washing the dishes at the sink . big sister washing the dishes and then she got dishes sitting on the sink . and I think she's running water . and I said Johnny he is up on the ladder getting some cookies and the little sister reaching up after some . he's passing it down to her . and the stool about to turn over . the cups maybe she going to wash them and she got them sitting on the sink . and maybe running water on the sink and if she got a curtain to pull that she might get some light in there . since the dishes stacked up . they might be on the sink . no that be about all .

(b)

all the action ? okay it's a boy and a girl and their mom . and well they're falling down in through here . and then this here when the water it should be going down in there but it's going down on the side here . it's going all the way down in there . they're getting something to eat here . cookiejar . and they're getting something to eat here . and this is a nice place what they have . but they put that stuff around in there . it looks nice . and then here when they had some stuff in through here . and ... I like these things in through here too . yeah .

(c)

Fig. 2. Label: Dementia, Prediction: Dementia.

I see a little boy on a stool almost falling over , taking cookies out of the cookie jar . and the little girl is putting her finger to her mouth to keep it quiet . the mother is washing dishes . she drying the dishes and letting the water keep on running in the sink . and then water is running over and she is standing in the water that running over . there a window there she looking at , at the grass and the flowers . and the curtains seem to be shaking from the wind and the air that blowing in . the dishes that she through drying are sitting on the sink top . and the little girl raising her hands for the little boy to hand her a cookie . and he has one cookie in his hand and he going after another one . he ready to hand her a cookie . mother is holding a dish cloth that she drying the dishes with . she has a platter that she drying . I don't see any other action .

(a)

well let's see . the girl is whispering to be quiet because mother might find out that the he's is standing on a stool which is bending over . and he's reaching in a cookie jar and he has a cookie . and she's grabbing for the one that he has in his left hand . and the sink is running over with water for some reason or other while she's drying a dish and looking out the window and stepping in a puddle of water . and the race horse is jumping through the window . no .

(b)

Fig. 3. Label: Control, Prediction: Control.

the sink's running over . the water's going all over the floor . here the stepstool is turning under his legs and he's stealing cookies out of the cookie jar . and she's begging for cookies the girl is coming back to the sink let's see here . mama's stepping in the water . and I said the sink was running over . she's drying dishes . wait a minute . what the devil is ? there is something but I don't know what it is written on the grass it seems . what is that ? and the curtains . that's a p p something there . hm let's see . I don't see anything else there . she's stepping in the water . the sink's running over . that spells something down there but I can't see it . so far . and he's on a stool that's gonna fall over while they're stealing cookies . and there's a plate and two cups on the sink and she's got a plate in her hand . I don't see anything else .

(a)

okay the kid on the bench who's got his hand in the cookie jar and he's falling off and his sister wants one . his mother is standing in a puddle of water because she didn't turn off the faucet and she's dry a dish . she oughta dry her feet instead . the window is open . well the sink is overflowing . it's obviously summer because the window as I said was open . there's supposedly leaves on the trees . anything else that I'm sposta pick up ? well the kid's gonna fall off . and it . the lid is off the cookie jar . and he's got one in his hand and handing it to his sister . and one and he's sneaking another one . not sneaking . the water is still running in the sink . and splashing on the floor .

(b)

okay . well in the first place the mother forgot to turn off the water and the water's running out the sink . and she's standing there . it's falling on the floor . the child is got a stool and reaching up into the cookie jar . and the stool is tipping over . and he's sort of put down the plates . and she's reaching up to get it but I don't see anything wrong with her though . yeah that's it . I can't see anything .

(c)

Fig. 4. Label: Control, Prediction: Dementia.

As one can easily observe in Fig. 2, tokens belonging to the UH pos-tag, such as yeah and oh, are identified as important for the dementia class by our best performing model. Moreover, personal pronouns (she, they) and verbs in the past tense (got, had) are also indicative of dementia. Also, our model considers the token “here,” which corresponds to the RB pos-tag, indicative of the dementia class. These findings are consistent with the ones in Section VIII-D, where we have found that PRP, VBD, UH pos-tags as well as the unigram “here” are significantly correlated with the dementia class. In addition, our model identifies the repetition of token “and” as important for the dementia class. This finding agrees with previous research works [17], where the word “and” indicates a short answer and burst of speech.

Regarding Fig. 3, one can easily observe that our model identifies tokens belonging to the VBG (putting, drying, blowing, standing, etc.), DT (the, a), and NN (cookie, action, stool, etc.) pos-tags as significant for the control class. Concurrently, in consistency with the findings in Section VIII-D, the unigrams “curtain” and “window” are used mainly by non-AD patients.

With regards to Figs. 1 and 4, our model is not able to classify these transcripts correctly. One possible reason for such misclassifications has to do with the fact that these transcripts include pos-tags which are indicative of both the control and the dementia class. To be more precise, in Fig. 1, the majority of tokens in both transcripts belong to the VBG, NN, and DT pos-tags, which are correctly identified by our model as significant for the control group. Words, like “and,” “him,” and “well” are used in a low frequency. Similarly to Fig. 1, in Fig. 4, the majority of tokens in each transcript belong to the pos-tags which are significantly correlated with the dementia class. This can be illustrated in Fig. 4(c), where we observe the usage of words, like “and,” “yeah,” “well” & “got”.

IX. CONCLUSION AND FUTURE WORK

We introduced both single-task and multi-task learning models. Regarding single-task learning models, we employed several transformer-based networks and compared their performances. Results showed that BERT achieved the highest classification performance with accuracy accounting for 87.50%. Concurrently, we introduced siamese networks coupled with a co-attention mechanism which can detect AD patients with an accuracy up to 83.75%. In terms of the multi-task learning setting, it consisted of two tasks, the primary and the auxiliary one. The primary task was the identification of dementia (binary classification), whereas the auxiliary task was the categorization of the severity of dementia into one of the four categories -healthy, mild/moderate/severe dementia- (multiclass classification). Specifically, we proposed two multi-task learning models. Results showed that our model achieves competitive results in the MTL framework reaching accuracy up to 86.25% on the detection of AD patients. Next, we performed an in-depth linguistic analysis, in order to understand better the differences in language between AD and non-AD patients. Finally, we employed LIME, in order to shed light on how our best performing model works. Findings suggest that AD patients tend to use personal pronouns, interjection, adverbs, verbs in the past tense, and the token

“and” at the beginning of utterances in a high frequency. On the contrary, healthy people use verbs in present participle or gerund, nouns as well as determiners.

One limitation of the current research work is pertinent to the small dataset used for conducting our experiments. However, we opted for this dataset, in order to mitigate different kinds of biases that could otherwise influence the validity of the proposed approaches.

We conducted our experiments on the ADReSS Challenge dataset, which is matched for gender and age and consists of a statistically balanced, acoustically enhanced set of recordings of spontaneous speech. Therefore, the results of this study could be integrated into an application, which will predict whether a person is an AD patient and will provide at the same time the reasons for this prediction via the explainability method.

In the future, we plan to investigate multimodal deep learning models incorporating both text and audio. Specifically, we plan to propose end-to-end trainable deep neural networks in contrast to existing research initiatives, which train multiple models separately and then use majority-voting approaches. In addition, our aim is to investigate fusion methods, in order to assign more importance to the most relevant modality and suppress the irrelevant information. Another future plan is to exploit further explainability techniques, such as anchor explanations [56].

REFERENCES

- [1] Alzheimer’s Association (2021), “What is dementia? Alzheimer’s disease and dementia.” Accessed: Jul. 30, 2021. [Online]. Available: <https://www.alz.org/alzheimers-dementia/what-is-dementia>
- [2] F. Zhang, Z. Li, B. Zhang, H. Du, B. Wang, and X. Zhang, “Multi-modal deep learning model for auxiliary diagnosis of Alzheimer’s disease,” *Neurocomputing*, vol. 361, pp. 185–195, 2019.
- [3] A. Syed Hassan and T. Khan, “A machine learning model to predict the onset of Alzheimer disease using potential cerebrospinal fluid (CSF) biomarkers,” *Int. J. Adv. Comput. Sci. Appl.*, vol. 8, no. 12, pp. 124–131, 2017.
- [4] C. Davatzikos, P. Bhatt, L. M. Shaw, K. N. Batmanghelich, and J. Q. Trojanowski, “Prediction of MCI to AD conversion, via MRI, CSF biomarkers, and pattern classification,” *Neurobiol. Aging*, vol. 32, no. 12, pp. 2322.e19–2322.e27, 2011.
- [5] C. Ieracitano, N. Mammone, A. Hussain, and F. C. Morabito, “A novel multi-modal machine learning based approach for automatic classification of EEG recordings in dementia,” *Neural Netw.*, vol. 123, pp. 176–190, 2020.
- [6] S. Luz, F. Haider, Sofia de la Fuente, D. Fromm, and B. MacWhinney, “Alzheimer’s dementia recognition through spontaneous speech: The ADReSS challenge,” in *Proc. InterSpeech*, Shanghai, China, 2020, pp. 2172–2176.
- [7] S. Luz, F. Haider, S. de la Fuente, D. Fromm, and B. MacWhinney, “Detecting cognitive decline using speech only: The ADReSSo challenge,” in *Proc. InterSpeech*, 2021, pp. 3780–3784.
- [8] J. Weiner and T. Schultz, “Selecting features for automatic screening for dementia based on speech,” in *Proc. Int. Conf. Speech Comput.*, 2018, pp. 747–756.
- [9] L. Calzá, G. Gagliardi, R. R. Favretti, and F. Tamburini, “Linguistic features and automatic classifiers for identifying mild cognitive impairment and dementia,” *Comput. Speech Lang.*, vol. 65, 2021, Art. no. 101113.
- [10] K. C. Fraser, K. L. Fors, M. Eckerström, F. Öhman, and D. Kokkinakis, “Predicting MCI status from multimodal language data using cascaded classifiers,” *Front. Aging Neurosci.*, vol. 11, 2019, Art. no. 205.
- [11] S. Nasreen, M. J. R. Hough, and M. Purver, “Alzheimer’s dementia recognition from spontaneous speech using disfluency and interactional features,” *Front. Comput. Sci.*, vol. 3, 2021. [Online]. Available: <https://www.frontiersin.org/article/10.3389/fcomp.2021.640669>

- [12] A. Khodabakhsh, S. Kuşuoğlu, and C. Demiroğlu, "Natural language features for detection of Alzheimer's disease in conversational speech," in *Proc. IEEE-EMBS Int. Conf. Biomed. Health Informat.*, 2014, pp. 581–584.
- [13] J. Chen, J. Zhu, and J. Ye, "An attention-based hybrid network for automatic detection of Alzheimer's disease from narrative speech," in *Proc. InterSpeech*, 2019, pp. 4085–4089.
- [14] D. F. Palo and N. Parde, "Enriching neural models with targeted features for dementia detection," in *Proc. 57th Annu. Meeting Assoc. Comput. Linguistics, Student Res. Workshop*, Florence, Italy, 2019, pp. 302–308.
- [15] Z. S. Syed, M. S. S. Syed, M. Lech, and E. Pirogova, "Automated recognition of Alzheimer's dementia using bag-of-deep-features and model ensembling," *IEEE Access*, vol. 9, pp. 88377–88390, 2021.
- [16] Y. Zhu, X. Liang, J. A. Batsis, and R. M. Roth, "Exploring deep transfer learning techniques for Alzheimer's dementia detection," *Front. Comput. Sci.*, vol. 3, 2021. [Online]. Available: <https://www.frontiersin.org/article/10.3389/fcomp.2021.624683>
- [17] S. Karlekar, T. Niu, and M. Bansal, "Detecting linguistic characteristics of Alzheimer's dementia by interpreting neural models," in *Proc. Conf. North Amer. Chapter Assoc. Comput. Linguistics, Hum. Lang. Technol.*, New Orleans, Louisiana, 2018, vol. 2, pp. 701–707.
- [18] M. T. Ribeiro, S. Singh, and C. Guestrin, "Why should I trust you?" explaining the predictions of any classifier," in *Proc. 22nd ACM SIGKDD Int. Conf. Knowl. Discov. Data Mining*, 2016, pp. 1135–1144.
- [19] R' mani Haulcy and J. Glass, "Classifying Alzheimer's disease using audio and text-based representations of speech," *Front. Psychol.*, vol. 11, 2021, Art. no. 3833.
- [20] Z. Shah, J. Sawalha, M. Tasnim, S. Qi, E. Stroulia, and R. Greiner, "Learning language and acoustic models for identifying Alzheimer's dementia from speech," *Front. Comput. Sci.*, vol. 3, no. 4, 2021, Art. no. 624659.
- [21] M. S. S. Syed, Z. S. Syed, M. Lech, and E. Pirogova, "Automated screening for Alzheimer's dementia through spontaneous speech," in *Proc. InterSpeech*, 2020, pp. 2222–2226.
- [22] A. Pompili, T. Rolland, and A. Abad, "The INESC-ID multi-modal system for the ADReSS 2020 challenge," in *Proc. InterSpeech*, 2020, pp. 2202–2206.
- [23] Y. Pan, B. Mirheidari, M. Reuber, A. Venneri, D. Blackburn, and H. Christensen, "Automatic hierarchical attention neural network for detecting AD," in *Proc. InterSpeech*, 2019, pp. 4105–4109.
- [24] W. Kong, H. Jang, G. Carenini, and T. Field, "A neural model for predicting dementia from language," in *Proc. Mach. Learn. HealthCare Conf.*, PMLR, 2019, pp. 270–286.
- [25] Y. Pan, V. S. Nallanthighal, D. Blackburn, H. Christensen, and A. Härmä, "Multi-task estimation of age and cognitive decline from speech," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process.*, 2021, pp. 7258–7262.
- [26] A. Balagopalan, B. Eyre, F. Rudzicz, and J. Novikova, "To BERT or not to BERT: Comparing speech and language-based approaches for Alzheimer's disease detection," in *Proc. InterSpeech*, 2020, pp. 2167–2171.
- [27] A. Meghanani, C. S. Anoop, and A. G. Ramakrishnan, "Recognition of Alzheimer's dementia from the transcriptions of spontaneous speech using fasttext and CNN models," *Front. Comput. Sci.*, vol. 3, no. 7, 2021, Art. no. 624558.
- [28] M. Rohanian, J. Hough, and M. Purver, "Multi-modal fusion with gating using audio, lexical and disfluency features for Alzheimer's dementia recognition from spontaneous speech," in *Proc. InterSpeech*, 2020, pp. 2187–2191.
- [29] M. Rohanian, J. Hough, and M. Purver, "Alzheimer's dementia recognition using acoustic, lexical, disfluency and speech pause features robust to noisy inputs," in *Proc. InterSpeech*, 2021, pp. 3820–3824.
- [30] N. Cummins *et al.*, "A comparison of acoustic and linguistics methodologies for Alzheimer's dementia recognition," in *Proc. InterSpeech*, 2020, pp. 2182–2186.
- [31] J. T. Becker, F. Boiler, O. L. Lopez, J. Saxton, and K. L. McGonigle, "The natural history of Alzheimer's disease: Description of study cohort and accuracy of diagnosis," *Arch. Neurol.*, vol. 51, no. 6, pp. 585–594, Jun. 1994.
- [32] B. MacWhinney, *The CHILDES Project: Tools for Analyzing Talk, Volume II: The Database*. New York, NY, USA: Psychology Press, 2014.
- [33] J. L. Lee, R. Burkholder, G. B. Flinn, and E. R. Coppess, "Working with chat transcripts in Python," Dept. Comput. Sci., Univ. Chicago, Chicago, IL, USA, Technical Rep. TR-2016-02, 2016.
- [34] J. Devlin, M. Chang, K. Lee, and K. Toutanova, "BERT: Pre-training of deep bidirectional transformers for language understanding," in *Proc. Conf. North Amer. Chapter Assoc. Comput. Linguistics: Hum. Lang. Technol.*, Minneapolis, MI, 2019, vol. 1, pp. 4171–4186.
- [35] J. Lee *et al.*, "BioBERT: A pre-trained biomedical language representation model for biomedical text mining," *Bioinformatics*, vol. 36, no. 4, pp. 1234–1240, 2020.
- [36] E. Alsentzer *et al.*, "Publicly available clinical BERT embeddings," in *Proc. 2nd Clin. Natural Lang. Process. Workshop*, Minneapolis, Minnesota, USA, 2019, pp. 72–78.
- [37] Z. Jiang, W. Yu, D. Zhou, Y. Chen, J. Feng, and S. Yan, "ConvBERT: Improving bert with span-based dynamic convolution," in *Proc. Adv. Neural Inf. Process. Syst.*, 2020, vol. 33, pp. 12837–12848.
- [38] Y. Liu *et al.*, "Roberta: A robustly optimized BERT pretraining approach," 2019. [Online]. Available: <https://arxiv.org/abs/1907.11692>
- [39] Z. Lan, M. Chen, S. Goodman, K. Gimpel, P. Sharma, and R. Soricut, "ALBERT: A lite bert for self-supervised learning of language representations," in *Proc. Int. Conf. Learn. Representations*, 2020. [Online]. Available: <https://openreview.net/forum?id=H1eA7AEtVS>
- [40] Z. Yang, Z. Dai, Y. Yang, J. Carbonell, R. R. Salakhutdinov, and Q. V. Le, "XLNet: Generalized autoregressive pretraining for language understanding," in *Proc. Adv. Neural Inf. Process. Syst.*, 2019, vol. 32, Art. no. 517.
- [41] J. Lu, J. Yang, D. Batra, and D. Parikh, "Hierarchical question-image co-attention for visual question answering," in *Proc. Adv. Neural Inf. Process. Syst.*, 2016, vol. 29, pp. 289–297.
- [42] K. Shu, L. Cui, S. Wang, D. Lee, and H. Liu, "DEFEND: Explainable fake news detection," in *Proc. 25th ACM SIGKDD Int. Conf. Knowl. Discov. Data Mining*, 2019, pp. 395–405.
- [43] Y. Lu and C. Li, "GCAN: Graph-aware co-attention networks for explainable fake news detection on social media," in *Proc. 58th Annu. Meeting Assoc. Comput. Linguistics*, 2020, pp. 505–514.
- [44] R. Caruana, "Multitask learning," *Mach. Learn.*, vol. 28, no. 1, pp. 41–75, 1997.
- [45] S. Rajamanickam, P. Mishra, H. Yannakoudakis, and E. Shutova, "Joint modelling of emotion and abusive language detection," in *Proc. 58th Annu. Meeting Assoc. Comput. Linguistics*, 2020, pp. 4270–4279.
- [46] M. Jin and N. Aletras, "Modeling the severity of complaints in social media," in *Proc. Conf. North Amer. Chapter Assoc. Comput. Linguistics: Hum. Lang. Technol.*, 2021, pp. 2264–2274.
- [47] T. Wolf *et al.*, "Transformers: State-of-the-art natural language processing," in *Proc. Conf. Empirical Methods Natural Lang. Process., Syst. Demonstrations*, 2020, pp. 38–45.
- [48] B. Portelli, E. Lenzi, E. Chersoni, G. Serra, and E. Santus, "BERT prescriptions to avoid unwanted headaches: A comparison of transformer architectures for adverse drug event detection," in *Proc. 16th Conf. Eur. Chapter Assoc. Comput. Linguistics: Main Volume*, 2021, pp. 1740–1747.
- [49] Y. Benjamini and Y. Hochberg, "Controlling the false discovery rate: A practical and powerful approach to multiple testing," *J. Roy. Stat. Soc., Ser. B (Methodological)*, vol. 57, no. 1, pp. 289–300, 1995.
- [50] E. Rissola, M. Aliannejadi, and F. Crestani, "Beyond modelling: Understanding mental disorders in online social media," in *Proc. Eur. Conf. Inf. Retrieval*, Springer, 2020, pp. 296–310.
- [51] H. A. Schwartz *et al.*, "Personality, gender, and age in the language of social media: The open-vocabulary approach," *PLoS One*, vol. 8, no. 9, 2013, Art. no. e73791.
- [52] A. Almor, D. Kempler, M. C. MacDonald, E. S. Andersen, and L. K. Tyler, "Why do alzheimer patients have difficulty with pronouns? working memory, semantics, and reference in comprehension and production in Alzheimer's disease," *Brain Lang.*, vol. 67, no. 3, pp. 202–227, 1999.
- [53] C. M. Watson, "An analysis of trouble and repair in the natural conversations of people with dementia of the Alzheimer's type," *Aphasiology*, vol. 13, no. 3, pp. 195–218, 1999.
- [54] L. J. Garcia and Y. Joannette, "Analysis of conversational topic shifts: A multiple case study," *Brain Lang.*, vol. 58, no. 1, pp. 92–114, 1997.
- [55] R. Bastiaanse, "Why reference to the past is difficult for agrammatic speakers," *Clin. Linguistics Phonetics*, vol. 27, no. 4, pp. 244–263, 2013.
- [56] M. T. Ribeiro, S. Singh, and C. Guestrin, "Anchors: High-precision model-agnostic explanations," in *Proc. 32nd AAAI Conf. Artif. Intell.*, 2018.