

# Context-Aware Attention Layers coupled with Optimal Transport Domain Adaptation and Multimodal Fusion methods for recognizing dementia from spontaneous speech

Loukas Ilias<sup>a,\*</sup>, Dimitris Askounis<sup>a</sup>

<sup>a</sup>Decision Support Systems Laboratory, School of Electrical and Computer Engineering, National Technical University of Athens, Zografou, Athens, 15780, Greece

## ARTICLE INFO

### Keywords:

Dementia  
log-Mel spectrogram  
BERT  
DeiT  
Optimal Transport  
Context-based Self-Attention  
Model Calibration  
label smoothing

## ABSTRACT

Alzheimer's disease (AD) constitutes a complex neurocognitive disease and is the main cause of dementia. Although many studies have been proposed targeting at diagnosing dementia through spontaneous speech, there are still limitations. Existing state-of-the-art approaches, which propose multimodal methods, train separately language and acoustic models, employ majority-vote approaches, and concatenate the representations of the different modalities either at the input level, i.e., early fusion, or during training. Also, some of them employ self-attention layers, which calculate the dependencies between representations without considering the contextual information. In addition, no prior work has taken into consideration the model calibration. To address these limitations, we propose some new methods for detecting AD patients, which capture the intra- and cross-modal interactions. First, we convert the audio files into log-Mel spectrograms, their delta, and delta-delta and create in this way an image per audio file consisting of three channels. Next, we pass each transcript and image through BERT and DeiT models respectively. After that, context-based self-attention layers, self-attention layers with a gate model, and optimal transport domain adaptation methods are employed for capturing the intra- and inter-modal interactions. Finally, we exploit two methods for fusing the self and cross-attention features. For taking into account the model calibration, we apply label smoothing. We use both performance and calibration metrics. Experiments conducted on the ADReSS and ADReSSo Challenge datasets indicate the efficacy of our introduced approaches over existing research initiatives with our best performing model reaching Accuracy and F1-score up to 91.25% and 91.06% respectively.

## 1. Introduction

Alzheimer's disease (AD) is a progressive neurologic disorder and constitutes the most common cause of dementia. According to the World Health Organization, around 55 million people have dementia worldwide with over 60% living in low- and middle-income countries [84]. In addition, dementia affects the ability of a person to communicate. More specifically, people with dementia may not be capable of finding the right words or may not be able to find any word at all. Concurrently, they are not able to stay focused on a discussion and tend to use words without meaning, thus being unable to communicate with other people [2]. This fact entails physical, psychological, social, and economic impacts not only for people living with dementia, but also for their carers, families, and society at large. Due to the fact that dementia becomes worse over time, it is important to be diagnosed early. For this reason, there have been several studies proposed, which distinguish AD patients from non-AD ones using speech and transcripts.

Although many methods have been proposed for fusing representation vectors from different modalities in many tasks, the task of multimodal dementia detection using speech and transcripts has still substantial limitations. More

specifically, most previous research works exploit early fusion, late fusion approaches, or add/concatenate the representations obtained by the different modalities. For the early fusion approaches, the authors simply concatenate features from the different modalities at the input level [58, 43]. For the late fusion approaches, the authors train textual and acoustic models separately and perform final decision voting in a weighted manner [47]. Other approaches train textual and acoustic models separately and apply a majority-vote approach for the final prediction [13]. It is obvious that these approaches do not account for both intra- and inter-modality interactions. Concurrently, majority-vote approaches increase the training time, since multiple models must be trained and tested separately. In addition, early fusion strategies or the add/concatenation operation give equal importance to the different modalities and do not capture the inherent correlations between the two modalities.

Recently, some studies [26, 27] address the limitations of fusing the different modalities. However, some limitations still exist. Specifically, the work in [26] concatenates the representation vectors of the two modalities and exploits a self-attention layer incorporating a gated model. However, in terms of the textual modality recent studies have shown that Self-Attention layers treat the input sequence as a bag-of-word tokens and each token individually performs attention over the bag-of-word tokens. Consequently, the contextual information is not taken into account in the calculation of dependencies between elements. There have been proposed

\*Corresponding author

✉ [lilias@epu.ntua.gr](mailto:lilias@epu.ntua.gr) (L. Ilias); [askous@epu.ntua.gr](mailto:askous@epu.ntua.gr) (D. Askounis)

ORCID(s): 0000-0002-4483-4264 (L. Ilias); 0000-0002-2618-5715 (D. Askounis)

a number of studies enhancing the self-attention layers with contextual information [76, 90, 81, 80]. In addition, the study in [27] employs a tensor fusion layer. However, the output tensor is high dimensional.

In addition, the reliability of a machine learning model's confidence in its predictions, denoted as calibration [14, 49], is critical for high risk applications, such as deciding whether to trust a medical diagnosis prediction [12, 29, 62]. However, no prior work has taken into account the calibration of the models, creating in this way overconfident models. According to [23], modern neural networks are not well-calibrated, while they are overconfident at the same time.

In order to tackle the aforementioned limitations, we introduce deep neural networks, which are trained in an end-to-end trainable manner and capture both the inter- and intra-modal interactions. First, we convert the audio files into images consisting of three channels, namely log-Mel spectrograms, their delta, and delta-delta. Next, each transcript and image are passed through BERT [16] and DeiT [74] models respectively. In order to ensure that the sequence length of the vectors obtained by BERT and DeiT is the same, we exploit an Optimal Transport Kernel (OTK) Embedding. We pass the textual representation through an enhanced self-attention layer with contextual information. We exploit three main methods for the contextualization, including the global context, deep context, and deep-global context [8, 85]. Next, we pass the image representation through a self-attention mechanism with a novel gating model proposed by [87] to model the intra-modal interactions. Motivated by the study of [59], we use optimal transport based domain adaptation [79] methods for capturing the inter-modal interactions. Then, we propose two attention-based methods for fusing the self and cross-attention features. Finally, for preventing models becoming too overconfident, we use label smoothing. We use metrics for assessing both the performance and the calibration of our model. We verify the effectiveness and robustness of our approaches by conducting experiments on two publicly available datasets, namely ADReSS and ADReSSo Challenge datasets, and using both manual and automatically generated transcripts. We show that our introduced approaches obtain multiple advantages over the state-of-the-art approaches.

Our main contributions can be summarized as follows:

- To the best of our knowledge, this is the first study utilizing DeiT, optimal transport kernel, and optimal transport domain adaptation methods in the task of dementia detection from spontaneous speech.
- This is the first study in the task of dementia detection from spontaneous speech exploiting label smoothing for preventing the models become too overconfident. We also evaluate our proposed models in terms of both the performance and the calibration.
- This is the first study in the task of dementia detection from speech data exploiting context-aware self-attention mechanisms and comparing two different

approaches for fusing the self- and cross-attention features.

- We conduct a series of ablation experiments to demonstrate the effectiveness of the introduced approach. We evaluate our approaches on the ADReSS and ADReSSo Challenge datasets and show that they achieve competitive results to the existing research initiatives.

## 2. Related Work

### 2.1. Dementia Detection

#### 2.1.1. Unimodal Approaches

Bertini et al. [6] introduced a unimodal approach using only speech data to detect AD patients. First, the authors applied a data augmentation technique, namely *SpecAugment* [56], for increasing the size of the data. Next, the authors used an autoencoder, called *auDeep* [21], and passed the latent vector representation through a multilayer perceptron for the final prediction. The main limitation of this approach is pertinent to the representation of the speech signal as a log-Mel spectrogram. However, appending delta and delta-delta features to log-Mel spectrogram is proven to be more beneficial, since delta features add dynamic information [33].

The research work proposed by [24] employed unimodal approaches by using only either speech or text to classify subjects into AD patients or non-AD ones. For the text modality, the authors extracted embeddings by using fast-Text, BERT, LIWC, and CLAN. For the acoustic modality, the authors extracted i-vectors and x-vectors. For both modalities, they employed dimensionality reduction techniques and trained shallow machine learning classifiers and neural networks (CNNs and LSTMs). The authors claimed that the Support Vector Machine and the Random Forest Classifiers trained on BERT embeddings achieved the highest accuracy. One limitation of this study is the fact that the authors used BERT embeddings as features for training additional algorithms. They did not experiment with extracting the [CLS] token and passing it to a dense layer for performing the classification.

Karlekar et al. [30] applied three deep neural networks based on CNNs, LSTM-RNNs, and their conjunction to distinguish AD patients from non-AD ones utilizing only transcripts. Next, they proposed explainability techniques by applying automatic cluster pattern analysis and first derivative saliency heat maps, in order to uncover differences in language between AD patients and healthy control groups. The main limitation of this paper is the fact that the authors did not experiment with language models based on transformers, i.e., BERT, RoBERTa, and so on.

The authors in [1] extracted a large number of acoustic features for detecting people in healthy control, mild cognitive impairment, and AD groups and used a longitudinal dataset for conducting their experiments. After extracting the acoustic features, they applied feature selection techniques for finding the optimal set of features. Also, they

applied sampling techniques, namely the Synthetic Minority Over-sampling Technique, for dealing with the imbalanced dataset. The Support Vector Machine and the Stochastic Gradient Descent were used for the classification purposes. The limitation of this paper lies on the feature extraction process, which is a time-consuming and tedious procedure. Additionally, the optimal feature set may not be found, since some level of domain expertise is required.

Similarly, the work proposed by [31] extracted seventeen features from transcripts for detecting AD patients. Specifically, the authors extracted the rate of pauses in utterances, filler sounds, number of no answers, part-of-speech tags, intelligibility of speech, diversity and complexity of the words, and many more. Next, they trained Support Vector Machines, Linear Discriminant Analysis, and Decision Trees. Results indicated that 90% prediction accuracy can be obtained using only phone entropy, silence rate per utterance, and word entropy with a Decision Tree classifier. The limitation of this paper lies on the feature extraction process, which is a time-consuming and tedious procedure. Additionally, the optimal feature set may not be found, since some level of domain expertise is required.

An augmented adversarial self - supervised learning method was proposed by [86]. Specifically, the introduced approach was based on contrastive predictive encoding. For dealing with the imbalanced dataset, i.e., limited number of speech samples corresponding to AD patients, the authors applied three augmentation schemes, including speed based augmentation, tempo based augmentation, and tremolo based augmentation. Findings indicated that the proposed methods improved the performance for AD detection to a large margin compared to other models.

### 2.1.2. Multimodal Approaches

Several approaches have been introduced which fuse the representation vectors or features of the different modalities at the input level. This strategy is known as an early fusion approach and does not capture effectively the inter-modal interactions. Edwards et al. [18] proposed a multimodal (audio and text) and multiscale (word and phoneme levels) approach. For the acoustic modality, the authors extracted features using the OpenSMILE toolkit, applied feature selection techniques, and trained shallow machine learning classifiers, including SVM, latent discriminant analysis (LDA), and LR. In terms of the language models, the authors trained a Random Forest Classifier on Word2Vec and GloVe embeddings. Also, they trained a FastText classifier from scratch. In addition, pretrained embeddings obtained by Sent2Vec, RoBERTa, ELECTRA, and so on were fine-tuned with the FastText classifier. The authors transcribed the segmented text into phoneme written pronunciation using CMUDict and stated that the FastText classifier was the best performing model trained on the phoneme representation. Results also showed that the combination of phonemes and audio yielded to the highest accuracy accounting for 79.17%. Martin and Pollak [43] proposed also an early fusion approach.

The authors extracted a large number of features corresponding to the textual and acoustic modality. They fused the feature sets via an early fusion method. Finally, they trained four machine learning classifiers, namely XGBoost, Random Forest, SVM, and Logistic Regression. Findings showed that the logistic regression and SVMs were proved to be better than XGBoost and Random Forest. Also, the authors stated that the readability features led to a surge in the classification performance. In terms of the audio features, the duration was the best performing one. Pompili et al. [58] proposed an early fusion approach for fusing the modalities of speech and transcript. Specifically, for the text modality, the authors employed the BERT model first and then trained three deep neural models on top of the BERT embeddings, namely (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. For the audio modality, the authors extracted the x-vectors. Finally, the authors merged the feature sets corresponding to the two different modalities and trained a Support Vector Machine classifier. Results showed that the fusion of the two modalities increased the performance obtained by unimodal approaches exploiting only speech or text.

Other approaches employ late-fusion strategies. This means that multiple models, i.e., acoustic and language, are trained separately and the final result/prediction is often taken after a majority vote approach. In this way, the inter-modal interactions are not captured. The authors in [70] proposed a majority-level approach for classifying AD patients using the audio and textual modalities. In terms of the textual modality, the authors extracted handcrafted textual features and deep textual embeddings of transcripts. For the extraction of deep textual embeddings, they used BERT, RoBERTa, and distilled versions of BERT and RoBERTa. Next, they exploited feature aggregation techniques and classified the subject as AD or non-AD patient by training either a Logistic Regression (LR) or a Support Vector Machine (SVM) classifier. In terms of the audio modality, the authors extracted handcrafted acoustic features, i.e., ComParE, COVAREP, etc. and deep acoustic embeddings, i.e., YAMNet, VGGish, etc. Similarly to the textual modality, they used feature aggregation techniques and trained a LR and SVM classifier. Results indicated that the majority-level approach of text models yielded the highest evaluation results, while the fusion of textual and acoustic modalities led to a degradation in performance. Shah et al. [69] introduced a weighted majority-vote ensemble meta-algorithm for classification utilizing the modalities of speech and transcripts. For the textual modality, the authors extracted language and fluency features, including the type-token ratio, the number of verbs per utterance, etc. and n-gram features. For the acoustic modality, the authors extracted four feature sets using the OpenSMILE v2.1 toolkit. After that, the authors applied dimensionality reduction techniques, i.e., Principal Component Analysis, and feature selection techniques, i.e., ANOVA F-values. Finally, shallow machine

learning classifiers were trained. Best results were obtained by using only the textual modality, while the majority vote approach by combining textual and acoustic modalities led to a decrease in the classification performance. Cummins et al. [13] exploited also a majority level approach for detecting AD patients. For the acoustic modality, the authors exploited bag-of-audio-words, siamese networks, and end-to-end convolutional neural networks, while for the textual modality, a bidirectional Hierarchical Attention Network and a BiLSTM with an attention mechanism were used. Findings indicated that the majority vote approach achieved the highest accuracy accounting for 85.20%. Sarawgi et al. [68] trained acoustic and language models separately and proposed three kinds of ensemble modules for classification. Specifically, the authors experimented with hard ensemble, meaning that a majority vote was taken between the predictions of the three individual models. A soft ensemble was also proposed, where a weighted sum of the class probabilities was computed for final decision, in order to leverage the confidence of the predictions. Also, a learnt ensemble was exploited, where a logistic regression classifier was trained using class probabilities as inputs. Results showed that the hard ensemble approach yielded the best results. Mittal et al. [47] proposed a late fusion strategy using the modalities of speech and transcripts. Firstly, they trained separately acoustic and language models. For the acoustic modality, the authors trained a VGGish model with log-mel spectrograms. For the textual modality, the authors concatenated the representation obtained by BERT, Sentence-BERT, and fastText-CNN. Finally, the probabilities calculated by the audio and text-based model were combined in a weighted manner, and a threshold was fixed for classifying the persons into AD and healthy control.

There are also approaches, which add or concatenate the representation vectors of different modalities during training. However, in this way, the inherent correlations between the different modalities are not captured. On the contrary, equal importance is assigned to the different modalities. Mahajan and Baths [42] introduced both unimodal and multimodal approaches for detecting AD patients. In terms of the multimodal approach, the introduced deep learning architecture consisted of GRU, Dense, CNN, BiLSTM, and Attention layers. For fusing the two modalities, the authors exploited the concatenation operation. Results indicated that the multimodal architecture achieved equal accuracy with the unimodal approach utilizing only text. Also, none of the proposed architectures surpassed the results of the baseline study [39]. Also, the authors in [91] proposed both unimodal and multimodal approaches. With regards to the multimodal approaches, the authors used the add and concatenation operation for fusing the modalities of speech and transcripts. Best results were obtained by concatenating the representations obtained by BERT and Speech BERT. The authors in [53] proposed both unimodal and multimodal approaches. Regarding unimodal approaches using speech data, the authors extracted acoustic features and trained four shallow machine learning classifiers. For the language

modality, the authors trained a BERT model. In terms of the multimodal approach, the authors simply concatenated the representations obtained by BERT and acoustic modality. Results on the test set indicated that the fusion approach achieved lower performance than the unimodal one using the textual modality.

Pappagari et al. [54] trained acoustic and language models separately and used the output scores as inputs to a Logistic Regression classifier for obtaining the final prediction. For the language models, the authors used automatic speech recognition models for transcribing the recordings and employed a BERT model. For the acoustic modality, the authors used x-vectors for classifying subjects into AD patients and non-AD ones. Also, they extracted eGeMAPS, VGGish, prosody features, etc. and trained Logistic Regression and XGBoost classifiers. The authors stated that the combination of the different models and the BERT model trained on automatic transcripts achieved equal accuracy on the test set. Similarly, the authors in [55] trained also acoustic and language models separately. In terms of the acoustic models, the authors extracted the x-vectors and trained a Probabilistic Linear Discriminant Analysis classifier. For the textual modality, the authors employed a BERT model. For fusing the two modalities, the authors employed the scores from the whole training subset to train a final fusion GBR model that was used to perform the fusion of scores coming from the acoustic and transcript-based models for the challenge evaluation. Results showed that the proposed approach was the best performing one.

A different approach was proposed by [64]. More specifically, the authors extracted textual and acoustic features and passed them through two different branches of BiLSTM layers. A gating mechanism consisting of highway networks was proposed for fusing the two modalities. However, the authors did not experiment with replacing the proposed fusion method with a concatenation operation via an ablation study. Thus, this fusion method cannot guarantee performance improvement. Similarly, [65] used BERT instead of BiLSTM for extracting the text representation and stated that the BiLSTM performed better than BERT due to the fewer parameters used.

In [26], the authors introduced three approaches for fusing the textual and visual modalities. Specifically, the authors passed the transcripts through BERT model and the images (log-Mel spectrogram, delta, delta-delta of the audio files) into a Vision Transformer. In terms of the first method for fusing the two modalities, the authors used a co-attention mechanism. Regarding the second method, the authors concatenated the representation matrices of the two modalities and passed the concatenated matrix through a gated-self attention layer. Finally, the authors used a Multimodal Shifting Gate, where they injected extra information to the BERT model, instead of capturing cross-modal interactions. Our work is different from [26], since we exploit context-based self-attention, optimal transport domain adaptation methods, Optimal Transport Kernel, DeiT, label smoothing, and one

different method for fusing the representation matrices of self and cross-attention features.

The authors in [27] exploited a Tensor Fusion Layer for explicitly aggregating unimodal, bimodal, and trimodal interactions. However, a high-dimensional output vector is generated through this method leading to the problem of overfitting and utilization of many resources.

Ilias et al. [28] introduced both unimodal and multimodal approaches. In terms of the unimodal approaches, the authors converted the audio files into log-Mel spectrograms, delta, and delta-delta and exploited several pretrained models, including Vision Transformer, ResNet50, VGG16, EfficientNet-B2, etc. Regarding the multimodal approach, the authors experimented with three methods for fusing the two modalities. Firstly, they used the concatenation operation. Secondly, they employed a Gated Multimodal Unit for controlling the influence of each modality. Finally, they exploited the crossmodal attention and stated that crossmodal models outperform the competitive multimodal ones.

## 2.2. Other Multimodal Tasks

Villegas et al. [67] introduced multimodal approaches for inferring the political ideology of an ad sponsor and identifying whether the sponsor is an official political party of a third-party organization. The authors employed BERT and EfficientNet [73] for extracting textual and visual representations respectively. They concatenated these two representations and passed the resulting vector to an output layer for binary classification. Results suggested that the combination of both modalities led to a surge in the classification performance.

Villegas and Aletras [66] proposed multimodal approaches for the task of point-of-interest type prediction. Specifically, the authors exploited BERT and Xception [10] for extracting text and visual representations respectively. Next, they introduced three different architectures for fusing the two modalities. First, they exploited the Gated Multimodal Unit introduced by [3]. Secondly, inspired by [75], they proposed a model for modeling the cross-modal interactions. Finally, the authors introduced an architecture, which includes the gated multimodal mechanism and the cross-attention layers on the top of the gated multimodal mechanism. Findings suggested that the proposed architecture yielded new state-of-the-art results outperforming significantly the previous text-only models.

Gu et al. [22] presented a deep multimodal network with both feature attention and modality attention to classify utterance-level speech data. The authors used the modalities of audio signal and text data as input to the deep neural network. In terms of the modality fusion approach proposed, it consisted of three main parts, namely the modality attention module, the weighted operation, and the decision making module. Findings showed that the multimodal system achieved state-of-the-art performance and was tolerant to noisy data indicating in this way its generalizability.

Pan et al. [52] proposed a multimodal architecture for detecting sarcasm in Twitter. More specifically, the authors

exploited the ResNet-152 model and obtained a visual representation. Regarding the textual modality, they used a pretrained BERT model. After obtaining embeddings for the input sequence and the hashtags included in the sequence, the authors passed the corresponding embeddings through encoders of the transformer. For modeling the cross-modal interactions, an additional encoder was used, where the visual representation corresponded to the key and value, while the sequence representation corresponded to the query. In addition, an intra-modality attention approach was used, which gets as input the sequence and the hashtag representations. The outputs obtained were concatenated and passed to an output layer for the final prediction. Findings stated that the proposed architecture achieved state-of-the-art results.

Inspired by the transformer model in machine translation [78], the authors in [88] presented some multimodal approaches for the task of visual question answering. More specifically, the authors employed a self-attention and a guided-attention unit for capturing the intra- and inter-modal interactions respectively. Next, they obtained a Modular Co-Attention layer, which constitutes the modular composition of the self-attention and guided-attention units. Finally, the authors proposed a deep Modular Co-Attention Network consisting of cascaded Modular Co-Attention layers. Results indicated that the introduced approach surpassed the existing co-attention models.

Zadeh et al. [89] introduced a novel model, termed Tensor Fusion Network, for the task of multimodal sentiment analysis. The authors used visual, language, and acoustic modalities. For capturing the intra-modal interactions, the authors proposed three Modality Embedding Subnetworks. For capturing the inter-modal interactions, the Tensor Fusion layer has been used. Finally, the authors employed the Sentiment Inference Subnetwork, which is conditioned on the output of the Tensor Fusion layer and performs sentiment inference. Results indicated a surge in performance in comparison with existing research initiatives.

Cai et al. [7] presented a multimodal approach for sarcasm detection in Twitter. The authors used the modalities of text features, image features, and image attributes. After extracting image features and attributes, the authors leveraged attribute features and BiLSTM layers for extracting the text features. Next, the authors employed a representation fusion approach for reconstructing the features of the three modalities. Finally, they proposed a modality fusion approach motivated by [22]. Results showed the effectiveness of the proposed architecture and the usefulness of the three modalities.

A different approach was proposed by [59], where the authors utilized optimal transport for capturing the cross-modal interactions and self attention mechanisms for capturing the intra-modal correspondence. Specifically, they exploited three different modalities, namely visual, language, and acoustic modalities. After utilizing self-attention and optimal transport methods, they used the multimodal attention fusion method introduced by [22]. Experiments conducted

towards the sarcasm and humor detection tasks demonstrated valuable advantages over existing research initiatives.

Yu et al. [87] introduced an approach for capturing both the inter- and intra-modal interactions for the visual question answering and the visual grounding tasks using the modalities of text and image. Specifically, after obtaining text and visual representations, they passed these two representations through a unified attention block. The authors proposed also a variation of the self-attention mechanism by introducing a novel gating model. Findings showed the effectiveness of the proposed approach on five datasets.

### 2.3. Related Work Review Findings

From the research works mentioned above, it is obvious that despite the fusion methods, which have been proposed for many tasks, little work has been done in terms of the multimodal approaches for the dementia detection task. More specifically, for the task of dementia detection, existing research initiatives employ majority-vote approaches or propose early and late fusion strategies. Applying majority-vote approaches and late fusion strategies entails significant increase in the computational time, since multiple models must be trained separately. For early fusion approaches, features obtained by different modalities are concatenated at the input level. All these approaches do not capture either the inter- or the intra-modal interactions effectively. Recently, there have been proposed some studies trying to address the limitations of fusing the different modalities. However, limitations still exist. Finally, no study has addressed the problem of creating overconfident models, which is crucial in health-related tasks.

Therefore, the present study differs from the existing state-of-the-art approaches, since we (i) exploit DeiT for extracting a visual representation, (ii) use optimal transport domain adaptation methods for capturing the inter-modal interactions and a self-attention mechanism with a gating model for capturing the intra-modal interactions, (iii) exploit a context-based self-attention mechanism, where we enhance the self-attention layer with contextual information by using three approaches, (iv) employ label smoothing for calibrating our proposed approaches, and (v) introduce a new method in the task of dementia recognition from spontaneous speech for fusing the self and cross-attention features.

## 3. Data & Task

### 3.1. ADReSS Challenge Dataset

We use the ADReSS Challenge Dataset [39] for conducting our experiments. The data corresponds to spoken picture descriptions elicited from participants through the Cookie Theft picture from the Boston Diagnostic Aphasia Exam [5]. We choose the specific dataset, since it minimizes several kinds of biases, which could influence the validity of the proposed approaches during the training and evaluation procedure. Specifically, in contrast to other datasets, the ADReSS Challenge dataset is matched for gender and age. In addition, it is balanced, since it includes 78 AD and 78-non

AD patients. What is also worth noting is the fact that the ADReSS Challenge dataset has been carefully selected so as to mitigate common 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. To be more precise, recordings have been acoustically enhanced with stationary noise removal and audio volume normalization has been applied across all speech segments to control for variation caused by recording conditions, such as microphone placement. The ADReSS Challenge dataset has been divided by the organizers into a train and a test set. The train set consists of 54 AD and 54 non-AD patients, while the test set comprises 24 AD patients and 24 non-AD ones.

### 3.2. ADReSSo Challenge Dataset

To further verify the effectiveness of our proposed approaches, we use the ADReSSo Challenge Dataset [40]. Similar to the ADReSS Challenge Dataset, this dataset has been created in a way for eliminating various kinds of biases. Additionally, each participant is asked to describe the Cookie Theft picture from the Boston Diagnostic Aphasia Examination. It is divided into a train and a test set. The train set consists of 87 AD patients and 79 non-AD ones, while the test set includes 35 AD and 36 non-AD patients. Also, this dataset includes only audio files. No transcripts are provided. Therefore, one should convert the speech into text automatically via Automatic Speech Recognition (ASR) methods. Specifically, we use whisper<sup>1</sup> [61] and get the automatically generated transcripts per audio file.

### 3.3. Task

Let a labeled dataset consist of audio files and their corresponding transcripts. Each transcript along with its audio file belongs to an AD patient or non-AD patient. The task is to identify if a specific transcript along with its audio file corresponds to an AD patient or to a person belonging to the healthy control group (binary classification problem).

## 4. Predictive Models

### 4.1. Architecture

In this section, we describe our proposed deep learning architectures for detecting AD patients. The proposed architectures are illustrated in Fig. 4. Due to the fact that the manual transcripts have been annotated using the CHAT coding system [41], we use the PyLangAcq library [34] for having access to these transcripts. In addition, we use the Python library, called librosa [44, 45], and convert each audio file into a log-Mel spectrogram, its delta, and delta-delta. In this way, we create an image consisting of three channels. For all the experiments conducted, we use 224 Mel bands, hop length equal to 1024, and a Hanning window. Each image is resized to (224 × 224) pixels.

Firstly, we pass each transcript through a BERT [16] model and the corresponding image through a DeiT [74]

<sup>1</sup><https://github.com/openai/whisper>

model. Formally, let  $X \in \mathbb{R}^{n \times D}$  and  $Y \in \mathbb{R}^{T \times D}$  be the outputs of the BERT and DeiT pretrained models respectively. Next, we pass  $Y$  through an Optimal Transport Kernel introduced by [46], in order to ensure that the sequence length of  $Y$  is equal to the sequence length of  $X$ , i.e.,  $T = n$ . Let  $S \in \mathbb{R}^{T \times D}$ , where  $T = n$ , denote the output representation of the Optimal Transport Kernel.

**Context-Aware Self Attention for the textual modality:**

Fig. 2a illustrates the conventional self-attention mechanism, which individually calculates the attention weight of two items, i.e., "the" and "tomorrow", ignoring the contextual information. In this study, we aim to enhance the self-attention layer by adding contextual information. Therefore, we exploit the context-based self-attention layer [8], which is illustrated in Fig. 1. We observe that this layer receives as input the input sequence denoted by  $X$  and the contextual information vector denoted by  $C$ .

We transform the input sequence  $X$  into a query, key, and value matrix, as described via the Equation 1:

$$Q = XW_q, K = XW_k, V = X \quad (1)$$

, where  $W_q \in \mathbb{R}^{D \times D_q}$ ,  $W_k \in \mathbb{R}^{D \times D_k}$  are learnable weight matrices.

As described in Equations 2 and 3, the context vector  $C \in \mathbb{R}^{n \times D_c}$  is transformed to a contextual query matrix  $Q_c \in \mathbb{R}^{n \times D_q}$  and a contextual key matrix  $K_c \in \mathbb{R}^{n \times D_k}$ :

$$Q_c = CW_q^c \quad (2)$$

, where  $W_q^c \in \mathbb{R}^{D_c \times D_q}$  is a learnable weight matrix.

$$K_c = CW_k^c \quad (3)$$

, where  $W_k^c \in \mathbb{R}^{D_c \times D_k}$  is a learnable weight matrix.

Next, we exploit gated sum, as illustrated in Fig. 1b, for quantifying the contribution of the input sequence  $X$  and the contextual vector  $C$  to the attention weight prediction. Finally, we get new query and key matrices denoted by  $\bar{Q} \in \mathbb{R}^{n \times D_q}$  and  $\bar{K} \in \mathbb{R}^{n \times D_k}$  respectively. We describe the equations governing the gated sum below:

$$g_q = \sigma(QW_g^Q + Q_cW_g^{Q_c}) \quad (4)$$

, where  $W_g^Q, W_g^{Q_c} \in \mathbb{R}^{D_q \times 1}$  are learnable parameters.

$$g_k = \sigma(KW_g^K + K_cW_g^{K_c}) \quad (5)$$

where  $W_g^K, W_g^{K_c} \in \mathbb{R}^{D_k \times 1}$  are learnable parameters.

$g_q$  and  $g_k$  indicate the weight of the importance of the contextual information.

$$\bar{Q} = (1 - g_q)Q + g_qQ_c \quad (6)$$

$$\bar{K} = (1 - g_k)K + g_kK_c \quad (7)$$

Therefore, we obtain new query and key matrices. Finally, we calculate the self-attention via the equation mentioned below:

$$Attention(\bar{Q}, \bar{K}, V) = softmax\left(\frac{\bar{Q} \cdot \bar{K}^T}{\sqrt{D_k}}\right)V \quad (8)$$

Next, we describe three methods, namely Global Context, Deep Context, and Deep-Global Context, for calculating the contextual vector  $C$ . Specifically, we follow [8, 85] to represent the context vector ( $C$ ), which is composed of internal representation.

- **Global Context:** Fig. 2b illustrates the global context strategy. More specifically, the global context indicates the mean operation over the input sequence for summarizing the input representation. Let  $X = [x_1, x_2, \dots, x_n] \in \mathbb{R}^{n \times D}$ . We calculate the context representation  $C$  as defined in Eq. 9. Note that the output of Eq. 9 is a vector, i.e.,  $C \in \mathbb{R}^D$ , instead of a matrix. To facilitate subsequent calculation operations, we use Eq. 10, where we obtain the contextual matrix  $\mathbf{C} \in \mathbb{R}^{n \times D}$ .

$$C = \bar{X} = Avgpool(X) = \frac{1}{n} \sum_1^n x_i \quad (9)$$

$$\mathbf{C} = stack(C, C, \dots, C) \quad (10)$$

- **Deep Context:** By deeply stacking self-attention layers, the model captures only high-level syntactic and semantic information neglecting the lower-level information. Therefore, as shown in Fig. 2c, the deep context strategy enables the layer to fuse different types of syntactic and semantic information captured by different layers.

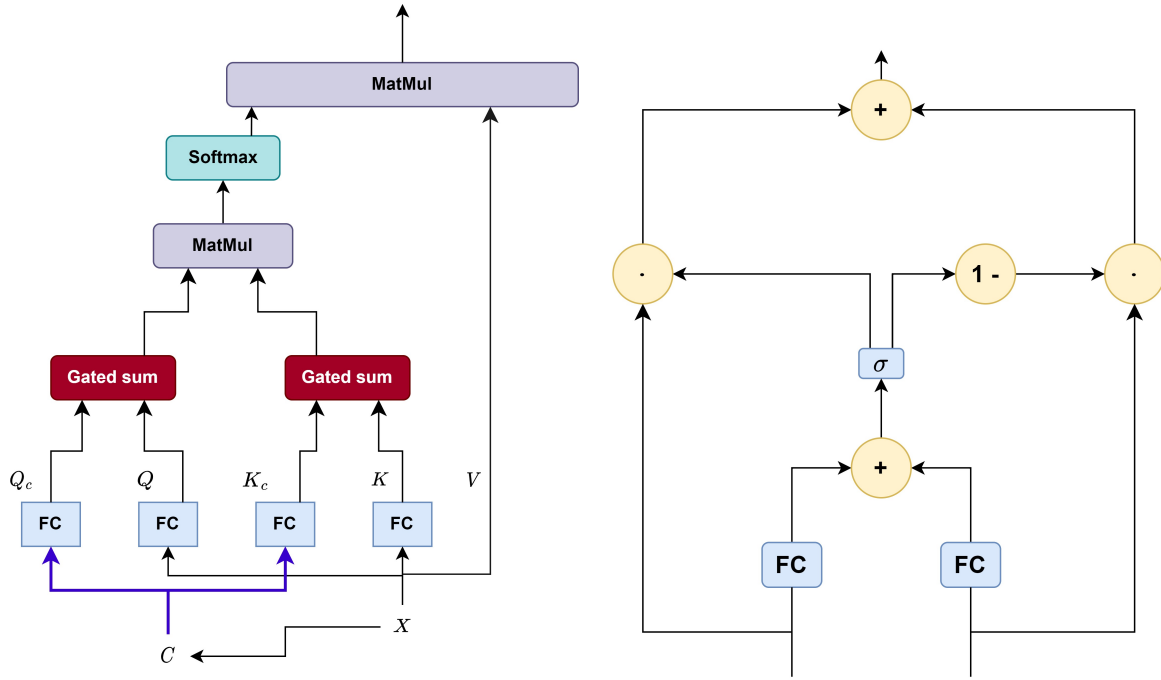
Formally, taking  $X = [x_1, x_2, \dots, x_n] \in \mathbb{R}^{n \times D}$  as the initial input sequence  $X^0$ , and the output of the  $L^{th}$  layer is  $X^L = [x_1^L, x_2^L, \dots, x_n^L] \in \mathbb{R}^{n \times D}$ , the deep context matrix  $C \in \mathbb{R}^{n \times D}$  can be represented as follows:

$$C = \hat{X}W_c^0 \quad (11)$$

$$\hat{X} = concat(X^0, X^1, \dots, X^{L-1}) \in \mathbb{R}^{n \times LD} \quad (12)$$

, where  $W_c^0 \in \mathbb{R}^{LD \times D}$  is a learnable parameter matrix.  $concat(.)$  denotes join operation.

- **Deep-Global Context:** The deep-global context strategy combines the strategies of global context and deep context as described before. The deep-global context strategy is illustrated in Fig. 2d and is described via the equations below:



(a) Context-based Self-Attention. This method is different from the conventional self-attention mechanism, since it exploits a contextual information vector  $C$ . (b) Gated sum. This unit is used for quantifying the contribution of the original representation  $X$  and the context vector  $C$  to the attention weight prediction.

Figure 1: Context-based Self-Attention

$$C = \bar{C}W_C^0 \quad (13)$$

$$\bar{C} = \text{concat}(C^0, C^1, \dots, C^{l-1}) \quad (14)$$

, where  $C^j = \text{Avgpool}(X^j)$ ,  $C^j \in \mathbb{R}^D$ . Therefore,  $\bar{C} \in \mathbb{R}^{lD}$ . In addition,  $W_C^0 \in \mathbb{R}^{lD \times D}$ . Thus, we obtain  $C$  of Eq. 13, as  $C \in \mathbb{R}^D$ .

As mentioned before, to facilitate subsequent calculation operations, matrix  $\mathbf{C} \in \mathbb{R}^{n \times D}$  is obtained through the stack operation, as follows:  $\mathbf{C} = \text{stack}(C, C, \dots, C)$ .

Let  $F$  be the output of the context-based self-attention mechanism corresponding to the textual modality denoted by  $X$ .

**Gated Self-Attention for the image modality:** Motivated by the work of [87], we pass  $S$  through a self-attention mechanism, which incorporates a novel gating model, for capturing the intra-modal interactions. This gated self-attention mechanism is illustrated in Fig. 3. The self-attention mechanism including the gating model is described via the equations below:

$$Q = S, K = S, V = S \quad (15)$$

$$M = \sigma \left( FC^g \left( FC_q^g(Q) \odot FC_k^g(K) \right) \right) \quad (16)$$

where  $FC_q^g, FC_k^g \in \mathbb{R}^{D \times d_g}$ ,  $FC^g \in \mathbb{R}^{d_g \times 2}$  are three fully-connected layers, and  $d_g$  denotes the dimensionality of the projected space and is equal to 64 units.  $\odot$  denotes the element-wise product function and  $\sigma$  the sigmoid function. In addition,  $M \in \mathbb{R}^{T \times 2}$  corresponds to the two masks  $M_q \in \mathbb{R}^T$  and  $M_k \in \mathbb{R}^T$  for the features  $Q$  and  $K$  respectively.

Next, the two masks  $M$  and  $K$  are tiled to  $\tilde{M}_q, \tilde{M}_k \in \mathbb{R}^{T \times D}$  and then used for computing the attention map as following:

$$A^g = \text{softmax} \left( \frac{(Q \odot \tilde{M}_q)(K \odot \tilde{M}_k)^T}{\sqrt{D}} \right) \quad (17)$$

$$H = A^g V \quad (18)$$

Let  $H$  be the output of the self-attention mechanism corresponding to the visual modality denoted by  $S$ .

**Optimal Transport:** Next, we use optimal transport-based domain adaptation methods [79, 11, 19], i.e., Earth Mover's Distance (EMD) Transport, for transporting between each pair of modalities, which can be interpreted as domain adaptation across two modalities. Formally:

$$X' = OT(S \rightarrow X) \quad (19)$$

$$S' = OT(X \rightarrow S) \quad (20)$$



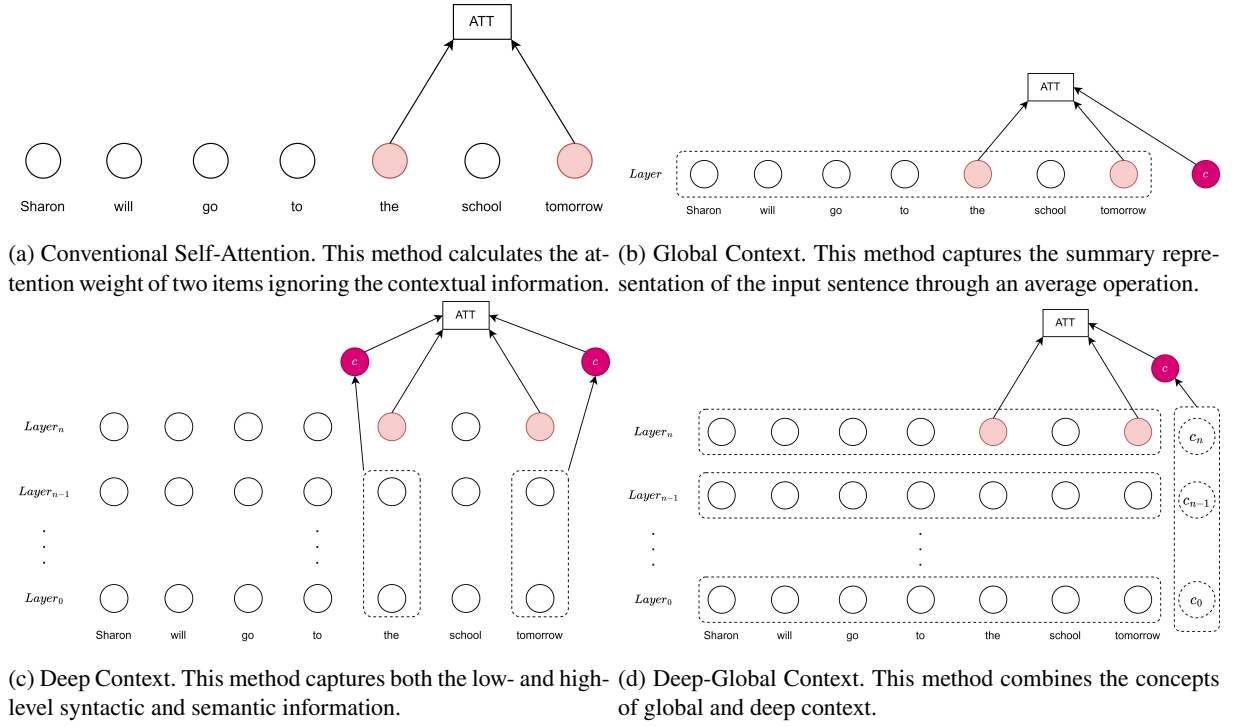


Figure 2: Self-Attention based on different context-vectors

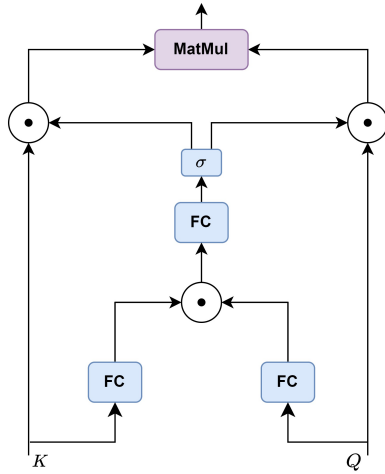


Figure 3: Gated Dot-product. This gating model is incorporated in the conventional self-attention mechanism for improving the quality of the learned attention. This method is based on low-rank bilinear pooling.

**Concatenation:** After that, we concatenate transported and self-attended features as follows:

$$C = [F, X'] \quad (21)$$

$$S = [H, S'] \quad (22)$$

**Fusion:** Next, we describe two methods for fusing  $C$  and  $S$ :

- **(i) Co-Attention Mechanism:** We exploit the fusion method proposed by [35] and implemented in prior work [26]. Specifically, given  $(C \in \mathbb{R}^{d' \times n})$  and  $(S \in \mathbb{R}^{d' \times T})$ , where  $d' = 2 \cdot D$ , the affinity matrix  $F \in \mathbb{R}^{n \times T}$  is calculated using the equation presented below:

$$F = \tanh(C^T W_l S) \quad (23)$$

where  $W_l \in \mathbb{R}^{d' \times d'}$  is a matrix of learnable parameters. By treating the affinity matrix as a feature, we learn to predict the attention maps via the following,

$$H^s = \tanh(W_s S + (W_c C) F) \quad (24)$$

$$H^c = \tanh(W_c C + (W_s S) F^T) \quad (25)$$

where  $W_s, W_c \in \mathbb{R}^{k \times d'}$  are matrices of learnable parameters. We set  $k$  equal to 40. Then, we generate the attention weights through the softmax function as follows,

$$a^s = \text{softmax}(w_{hs}^T H^s) \quad (26)$$

$$a^c = \text{softmax}(w_{hc}^T H^c) \quad (27)$$

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 are obtained via the following equations:

$$\hat{s} = \sum_{i=1}^T a_i^s s^i, \quad \hat{c} = \sum_{j=1}^n a_j^c c^j \quad (28)$$

where  $\hat{s} \in \mathbb{R}^{1 \times d'}$  and  $\hat{c} \in \mathbb{R}^{1 \times d'}$ .

Finally, these vectors are concatenated  $p = [\hat{c}, \hat{s}]$ . We apply a dropout layer with a rate of 0.5. Then, this vector is passed through a Dense Layer consisting of 128 units with a ReLU activation function. We apply also a dropout layer with a rate of 0.2. Finally, we use a dense layer consisting of two units, which gives the final output.

The proposed architecture is illustrated in Fig. 4a.

- **(ii) Attention-based fusion:** Motivated by the work of [88], we design an attentional reduction model for  $C$ , as defined in Equation 21 (or  $S$ , as defined in Equation 22), for obtaining its attended feature  $\tilde{c}$  (or  $\tilde{s}$ ). To the best of our knowledge, this is the first study utilizing this fusion method in the task of dementia detection from spontaneous speech. Taking  $C$  as an example, we describe the attention reduction model used in this study via the equations presented below:

$$\alpha^c = \text{softmax}(MLP(C)) \quad (29)$$

, where  $\alpha^c$  refers to the learned attention weights and  $MLP$  is given by the equation below:

$$MLP = FC(128) - ReLU - Dropout(0.1) - FC(1) \quad (30)$$

$$\tilde{c} = \sum_{i=1}^n \alpha_i^c c_i \quad (31)$$

, where we obtain the attended feature  $\tilde{c}$  for  $C$ .

We obtain the attended feature  $\tilde{s}$  using an independent attention reduction model in the same way. Having computed  $\tilde{c}$  and  $\tilde{s}$ , we design the linear multimodal fusion function as follows:

$$z = \text{LayerNorm}(W_c^T \tilde{c} + W_s^T \tilde{s}) \quad (32)$$

, where  $W_c, W_s \in \mathbb{R}^{d' \times d_z}$  are two linear projection matrices,  $d_z$  is the common dimensionality of the fused feature and is equal to 128, and LayerNorm [4] is used for stabilizing the training. Finally, we pass  $z$  to a dense layer consisting of two units, which gives the final prediction.

The proposed architecture is illustrated in Fig. 4b.

## 4.2. Model Calibration

To prevent the model becoming too overconfident, we use label smoothing [72, 48]. Specifically, label smoothing calibrates learned models so that the confidences of their predictions are more aligned with the accuracies of their predictions.

For a network trained with hard targets, the cross-entropy loss is minimized between the true targets  $y_k$  and the network's outputs  $p_k$ , as in  $H(y, p) = \sum_{k=1}^K -y_k \log(p_k)$ , where  $y_k$  is "1" for the correct class and "0" for the other. For a network trained with label smoothing, we modify the true targets  $y_k$  to  $y_k^{LS_u}$  as shown in Eq. 33:

$$y_k^{LS_u} = y_k \cdot (1 - \alpha) + \frac{\alpha}{K} \quad (33)$$

, where  $\alpha$  is the smoothing parameter and  $K$  is the number of classes.

Finally, we minimize the cross-entropy between the modified targets  $y_k^{LS_u}$  and the network's outputs  $p_k$ , as shown in Eq. 34:

$$H(y, p) = \sum_{k=1}^K -y_k^{LS_u} \cdot \log(p_k) \quad (34)$$

## 5. Experiments

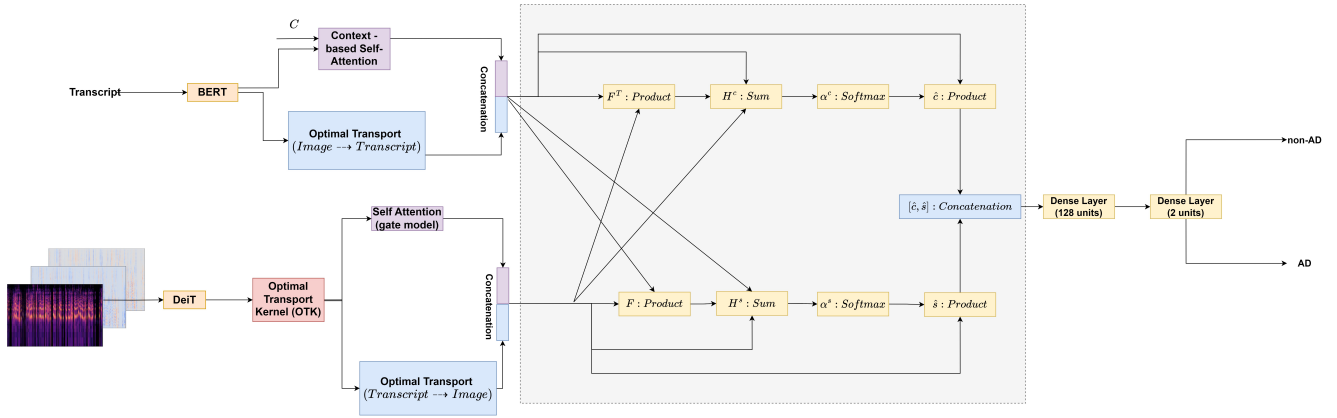
### 5.1. Baselines

We compare our introduced approaches with the following research works reported in Tables 1 and 2, since these research works have conducted their experiments on the ADReSS and ADReSSo test set. Specifically, Table 1 describes the baselines used in terms of the ADReSS Challenge dataset, while Table 2 reports the baselines used regarding the ADReSSo Challenge dataset. Regarding Table 1, we are using existing published results for all the baselines except for *Introduced Approaches without Label Smoothing*. In terms of Table 2, we are using existing published results for all the baselines except for: (i) *BERT*, and (ii) *Introduced Approaches without Label Smoothing*.

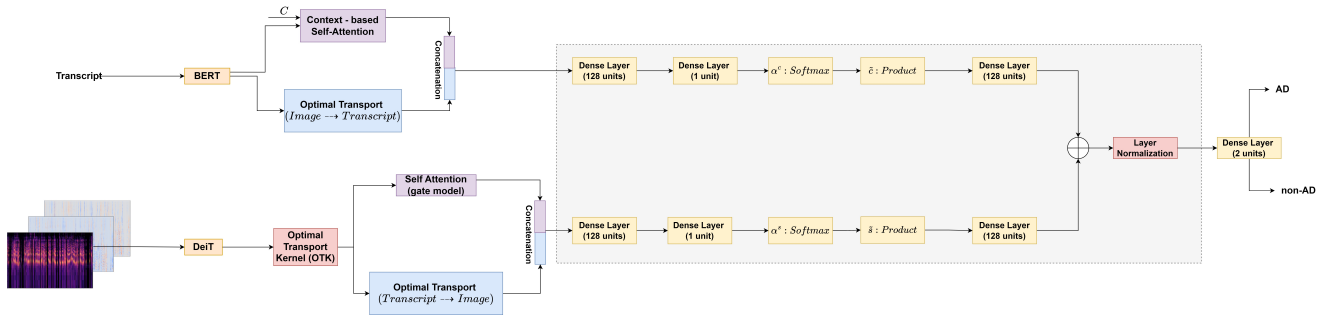
### 5.2. Experimental Setup

We divide the ADReSS Challenge train set into a train and a validation set (65%-35%). We use a batch size of 4. We train the introduced architectures five times and report the results on the ADReSS Challenge test set via mean  $\pm$  standard deviation. Similarly, we divide the ADReSSo Challenge train set into a train and a validation set (65%-35%). We train the introduced architectures five times and report the results on the ADReSSo Challenge test set via mean  $\pm$  standard deviation. We use *EarlyStopping*, where we stop training if the validation loss has stopped decreasing for eight consecutive epochs. Also, we apply *StepLR* with a *step\_size* of 4 and a *gamma* of 0.1. We set  $\alpha$  of Eq. 33 equal to 0.001. We set  $D = D_c = 768$ . We set  $D_k = D_q = 64$ . Regarding the global context strategy, we use

Context-Aware Attention Layers coupled with Optimal Transport Domain Adaptation and Multimodal Fusion methods for recognizing dementia from spontaneous speech



(a) Co-Attention. The shaded box corresponds to the co-attention mechanism. This method attends to the different representations simultaneously.



(b) Attention-based Fusion. The shaded box shows this fusion method. This method exploits two independent attentional reduction models. Features are fused through an add operation, while a layer normalization is used for stabilizing training.

**Figure 4:** Illustration of our Proposed Architectures. For the textual modality, we use BERT, while for the image modality, we use DeiT and exploit an Optimal Transport Kernel. Next, we use optimal transport domain adaptation methods for transporting between each pair of modalities. Also, we pass the textual representation through context-based self-attention layers, while the image representation is passed through a gated self-attention layer. Finally, methods for fusing the self- and cross-attention features are presented, namely Co-Attention and Attention-based Fusion. Each shaded box shows the fusion method used, namely Co-Attention and Attention-based Fusion.

one layer of the contextual self-attention mechanism. In terms of the deep-context strategy, we use three layers of the contextual self-attention mechanism. With regards to the deep-global context strategy, we use two layers of the contextual self-attention mechanism. We use the BERT base uncased version and the DeiT<sup>2</sup> model from the Transformers library [83]. For the optimal transport methods, we use the Python library Optimal Transport [20]. All the models have been created using the PyTorch library [57]. All experiments are conducted on a single Tesla P100-PCIE-16GB GPU.

### 5.3. Evaluation Metrics

#### 5.3.1. Performance Metrics

Accuracy (Acc.), Precision (Prec.), Recall (Rec.), F1-Score, and Specificity (Spec.) have been used for evaluating the results of the introduced architectures. These metrics have been computed by regarding the dementia class as the positive one. We report the average and standard deviation of these metrics over five runs.

<sup>2</sup>facebook/deit-base-distilled-patch16-224

#### 5.3.2. Calibration Metrics

We evaluate the calibration of our model using the metrics proposed by [50, 51, 23]. Specifically, we use the metrics mentioned below:

- **Expected Calibration Error (ECE).** The calibration error is the difference between the fraction of predictions in the bin that are correct (accuracy) and the mean of the probabilities in the bin (confidence). First, we divide the predictions into  $M$  equally spaced bins (size  $1/M$ ).

$$acc(B_m) = \frac{1}{|B_m|} \sum_{i \in B_m} 1(\hat{y}_i = y_i) \quad (35)$$

$$conf(B_m) = \frac{1}{|B_m|} \sum_{i \in B_m} \hat{p}_i \quad (36)$$

, where  $y_i$  and  $\hat{y}_i$  are the true and predicted labels for the sample  $i$  and  $\hat{p}_i$  is the confidence (predicted probability value) for sample  $i$ .

**Table 1**  
Baselines (ADReSS Challenge Dataset).

Reference/Architecture	Features/Methodology
<b>Baselines - Unimodal state-of-the-art approaches (only transcripts)</b>	
BERT [25]	Fine-tune a BERT model
<b>Baselines - Multimodal state-of-the-art approaches</b>	
Fusion Maj. (3-best) [13]	Majority Vote of the BoAW-MFCC-C125, ZFF, and bi-LSTM-Att
System 3: Phonemes and Audio [18]	Acoustic features (emobase, eGeMAPS, ComParE2016) along with feature selection techniques, transcription of the segmented text into phoneme written pronunciation using CMUDict
Fusion of System [58]	merged the x-vectors features set with the combination of linguistic feature sets (GMax/LSTM-RNNs/LSTM-RNNs-Pos) and trained a SVM classifier
Bimodal Network (Ensembled Output) [32]	For the acoustic modality, the authors use VGGish, while for the textual modality, the authors exploit GloVe, Transformer-XL, POS and HC features. Finally, the authors combine the results of the models via an ensemble approach.
GFI, NUW, Duration, Character 4-grams, Suffixes, POS tag, UD [43]	feature extraction, early fusion approach, train a Logistic Regression Classifier
Acoustic & Transcript [55]	For transcripts, the authors exploited BERT, while for speech, the authors used x-vector PCA-transformed coefficients.
Dual BERT [91]	concatenation of the representations obtained by BERT and Speech BERT
Model C [42]	The authors extracted emobase, eGeMAPS, ComParE features. For the text modality, the used GloVe embeddings and pos-tags. Finally, they trained a Neural network consisting of CNN, BiLSTM, Attention, GRU, and Dense layers.
Majority vote (NLP + Acoustic) [69]	The authors extracted a set of acoustic and linguistic features. After training shallow machine learning classifiers, they chose the three best-performing acoustic models along with the best-performing language model, and computed a final prediction by taking a linear weighted combination of the individual model predictions.
Audio + Text [70]	majority level approach of six models
LSTM with Gating (Acoustic + Lexical + Dis) [64]	Acoustic, Linguistic Features, Bi-LSTM, feed-forward highway layers with gating units
Ensemble [68]	A majority vote was taken between the predictions of the three individual models. Specifically, the authors extracted three sets of features, namely disfluency, acoustic, and interventions, and trained three deep neural networks.
BERT+ViT (log-Mel spectrogram) [28]	conversion of an audio file into an image of three channels, BERT for the text representation, Vision Transformer for the image representation, concatenation
BERT+ViT+Gated Multimodal Unit (log-Mel spectrogram) [28]	Gated Multimodal Unit to control the information flow of the different modalities.
BERT+ViT+Crossmodal Attention (log-Mel spectrogram) [28]	Similar to [75], the authors exploited a cross-attention mechanism.
BERT+ViT+Co-Attention [26]	The authors used a co-attention mechanism to fuse the representation matrices of the two modalities.
Multimodal BERT - eGeMAPS [26]	The authors injected acoustic information (eGeMAPS) into a BERT model.
Multimodal BERT - ViT [26]	The authors injected image information (via ViT) into a BERT model.
Multimodal BERT - eGeMAPS+ViT [26]	The authors injected both acoustic information (eGeMAPS) and image information (via ViT) into a BERT model.
BERT+ViT+Gated Self-Attention [26]	The authors concatenated the outputs of BERT and ViT and passed the resulting matrix through a self-attention layer incorporating a gate model for capturing the inter- and intra-modal interactions.
Transcript+Image+Acoustic [27]	The authors used a Tensor Fusion Layer for fusing the different modalities.
<b>Introduced Approaches without Label Smoothing</b>	
	Our proposed approaches described in Section 4 without label smoothing.

$$ECE = \sum_{m=1}^M \frac{|B_m|}{N} \left| acc(B_m) - conf(B_m) \right| \quad (37)$$

, where  $N$  is the total number of data points and  $B_m$  is the group of samples whose predicted probability values falls into the interval  $I_m = \left( \frac{m-1}{M}, \frac{m}{M} \right]$ .

Perfectly calibrated models have an ECE of 0.

**Table 2**  
Baselines (ADReSSo Challenge Dataset).

Reference/Architecture	Features/Methodology
<b>Baselines - Unimodal state-of-the-art approaches (only transcripts)</b>	
BERT	We exploit a BERT model, get the [CLS] token, and pass it through two dense layers consisting of 128 and 2 units respectively.
Model C: LR[Comp]+LR[DisFl]+Ernie+Bert (stacking) [60]	The authors employed model stacking to combine two logistic regression models (LR) using complexity and (dis)fluency features respectively, and the two pretrained language models, i.e. BERT and ERNIE.
Model 5: [53]	The authors concatenate the last three states of the BERT sequence classifier with the confidence score input. The confidence score input is generated by the ASR system.
Label Fusion selected models [71]	The authors extracted a set of handcrafted features, namely syntactic, readability, and lexical diversity, and a set of deep textual embeddings, including BERT and so on. Finally, the authors trained Logistic Regression and SVM classifiers.
Mp1 [92]	The authors add sentence-level pauses to ASR transcripts and exploit a BERT model.
<b>Baselines - Multimodal state-of-the-art approaches</b>	
LSTM w/ Gating (Words+Acoustic+Disf+Pse+WP) [65]	extraction of acoustic and language features, feed-forward highway layers with gating units
Global Fusion [54]	fusion of BERT (ASR) and acoustic models, namely x-vectors, x-vectors with 250ms frame-length, and encoder-decoder ASR embeddings (SB Enc/Dec).
Top-10 Avg. [9]	Average fusion of predicted class probabilities of the 10 best performing models
Attempt 1: [82]	The authors used acoustic features, linguistic features, and embedding features. For each type of feature, they exploited a deep neural network consisting of multihead attention layers, convolutional layers, and dilated convolutional layers. They used an attention layer for fusing the outputs of the different branches.
<b>Introduced Approaches without Label Smoothing</b>	
Our proposed approaches described in Section 4 without label smoothing.	

- **Adaptive Calibration Error (ACE).** Adaptive Calibration Error uses an adaptive scheme which spaces the bin intervals so that each contains an equal number of predictions.

$$ACE = \frac{1}{KR} \sum_{k=1}^K \sum_{r=1}^R |acc(r, k) - conf(r, k)| \quad (38)$$

, where  $acc(r, k)$  and  $conf(r, k)$  are the accuracy and confidence of adaptive calibration range  $r$  for class label  $k$ , respectively; and  $N$  is the total number of data points. Calibration range  $r$  defined by the  $[N/R]$ th index of the sorted and thresholded predictions.

## 6. Results

The results of our introduced models are reported in Tables 3 and 4. Specifically, Table 3 reports the results on the ADReSS Challenge dataset, while Table 4 reports the results on the ADReSSo Challenge dataset. Also, these tables present a comparison of our introduced approaches with existing research initiatives, which have proposed either unimodal or multimodal approaches. In order to compare models, we use the Almost Stochastic Order (ASO) test

[15, 17] of statistical significance implemented by [77]. We use  $confidence\_level = 0.95$  and  $num\_comparisons = 50$ . Generally, the ASO test determines whether a stochastic order [63] exists between two models or algorithms, i.e.,  $A$  and  $B$ . This method computes a score ( $\epsilon_{min}$ ) which represents how far the first is from being significantly better in respect to the second. When  $\epsilon_{min} = 0$ , then one can claim that  $A$  is truly stochastically dominant over  $B$ . When  $\epsilon_{min} < 0.5$ , one can claim that  $A$  is almost stochastically dominant over  $B$ . For  $\epsilon_{min} = 0.5$ , no order can be determined. ( $\dagger$ ) means that Attention-based Fusion (Deep Context) with label smoothing is stochastically dominant over the respective models. Similarly, in terms of the ADReSSo Challenge dataset, ( $\dagger$ ) means that Co-Attention (Deep Context) with label smoothing is stochastically dominant over the respective models. ( $\star$ ) denotes almost stochastic dominance of the Attention-based Fusion (Deep Context) with label smoothing over the respective approaches. Similarly, in terms of the ADReSSo Challenge dataset, ( $\star$ ) means that Co-Attention (Deep Context) with label smoothing is almost stochastically dominant over the respective models. Note that we cannot compare our approaches with all the existing research initiatives, since we do not have access to the multiple runs or the other approaches have not used multiple runs. In terms of the ECE and ACE calibration metrics, we use ASO for comparing

**Table 3**

Performance comparison among proposed models and state-of-the-art approaches on the ADReSS Challenge test set. Reported values are mean  $\pm$  standard deviation. Results are averaged across five runs. ( $\dagger$ ) means that Attention-based Fusion (Deep Context) with label smoothing is stochastically dominant over the respective models. ( $\star$ ) denotes almost stochastic dominance of the Attention-based Fusion (Deep Context) with label smoothing over the respective approaches.

Architecture	Prec. (%)	Rec. (%)	F1-score (%)	Acc. (%)	Spec. (%)	ECE	ACE
<b>Baselines - Unimodal state-of-the-art approaches (only transcripts)</b>							
BERT [25]	87.19 $\pm 3.25$	81.66 $\pm 5.00$	86.73 $\dagger$ $\pm 4.53$	87.50 $\dagger$ $\pm 4.37$	93.33 $\pm 5.65$		
<b>Baselines - Multimodal state-of-the-art approaches</b>							
Fusion Maj. (3-best) [13]	-	-	85.40	85.20	-		
System 3: Phonemes and Audio [18]	81.82	75.00	78.26	79.17	83.33		
Fusion of system [58]	94.12	66.67	78.05	81.25	95.83		
Bimodal Network (Ensembled Output) [32]	89.47	70.83	79.07	81.25	91.67		
GFI, NUW, Duration, Character 4-grams, Suffixes, POS tag, UD [43]	-	-	-	77.08	-		
Acoustic & Transcript [55]	70.00	88.00	78.00	75.00	63.00		
Dual BERT [91]	83.04 $\pm 3.97$	83.33 $\pm 5.89$	82.92 $\pm 1.86$	82.92 $\pm 1.56$	82.50 $\pm 5.53$		
Model C [42]	78.94	62.50	69.76	72.92	83.33		
Majority vote (NLP+Acoustic) [69]	-	-	-	83.00	-		
Audio + Text [70]	-	87.50	-	89.58	91.67		
LSTM with Gating (Acoustic + Lexical + Dis) [64]	81.82	75.00	78.26	79.17	83.33		
Ensemble [68]	83.00	83.00	83.00	83.00	-		
BERT+ViT [28]	90.73 $\pm 2.74$	80.83 $\pm 2.04$	85.47 $\dagger$ $\pm 1.70$	86.25 $\dagger$ $\pm 1.67$	91.67 $\pm 2.64$		
BERT+ViT+Gated Multimodal Unit [28] (log-Mel spectrogram)	80.92 $\pm 2.30$	91.67 $\pm 3.73$	85.92 $\dagger$ $\pm 2.37$	85.00 $\dagger$ $\pm 2.43$	78.33 $\pm 3.12$		
BERT+ViT+Crossmodal Attention [28] (log-Mel spectrogram)	86.13 $\pm 3.26$	91.67 $\pm 4.56$	88.69 $\star$ $\pm 2.12$	88.33 $\star$ $\pm 2.12$	85.00 $\pm 4.25$		
BERT+ViT+Co-Attention [26]	92.83 $\pm 6.39$	81.67 $\pm 2.04$	86.81 $\star$ $\pm 3.37$	87.50 $\star$ $\pm 3.49$	93.33 $\pm 6.24$		
Multimodal BERT - eGeMAPS [26]	74.51 $\pm 1.01$	87.50 $\pm 6.45$	80.35 $\dagger$ $\pm 2.77$	78.75 $\dagger$ $\pm 2.04$	70.00 $\pm 3.12$		
Multimodal BERT - ViT [26]	73.91 $\pm 2.40$	91.67 $\pm 2.64$	81.79 $\dagger$ $\pm 1.72$	79.58 $\dagger$ $\pm 2.04$	67.50 $\pm 4.08$		
Multimodal BERT - eGeMAPS+ViT [26]	76.57 $\pm 3.74$	89.17 $\pm 5.65$	82.28 $\dagger$ $\pm 3.49$	80.83 $\dagger$ $\pm 3.58$	72.50 $\pm 5.65$		
BERT+ViT+Gated Self-Attention [26]	90.87 $\pm 3.50$	89.17 $\pm 2.04$	89.94 $\star$ $\pm 1.36$	90.00 $\star$ $\pm 1.56$	90.83 $\pm 4.08$		
Transcript+Image+Acoustic [27]	90.88 $\pm 3.60$	80.83 $\pm 2.04$	85.48 $\dagger$ $\pm 0.76$	86.25 $\dagger$ $\pm 1.02$	91.66 $\pm 3.73$		
<b>Baselines - Introduced models (without label smoothing)</b>							
Co-Attention (Global Context)	89.62 $\pm 1.75$	85.83 $\pm 3.33$	87.63 $\dagger$ $\pm 1.80$	87.92 $\dagger$ $\pm 1.56$	90.00 $\pm 2.04$	0.1208 $\pm 0.2296$	0.1660 $\pm 0.0335$
Co-Attention (Deep Context)	88.25 $\pm 1.56$	87.50 $\pm 2.64$	87.85 $\star$ $\pm 1.66$	87.92 $\dagger$ $\pm 1.56$	88.33 $\pm 1.66$	0.1384 $\pm 0.0109$	0.1532 $\pm 0.0110$
Co-Attention (Deep-Global Context)	90.26 $\pm 1.70$	85.00 $\pm 4.25$	87.51 $\star$ $\pm 2.69$	87.92 $\star$ $\pm 2.43$	90.83 $\pm 1.66$	0.1355 $\pm 0.0183$	0.1648 $\pm 0.0119$
Attention-based Fusion (Global Context)	89.55 $\pm 7.31$	85.83 $\pm 6.24$	87.32 $\star$ $\pm 4.35$	87.50 $\star$ $\pm 4.37$	89.16 $\pm 8.58$	0.1256 $\pm 0.0291$	0.1279 $\pm 0.0277$
Attention-based Fusion (Deep Context)	91.06 $\pm 5.04$	89.16 $\pm 3.33$	89.95 $\star$ $\pm 1.91$	90.00 $\star$ $\pm 2.04$	90.83 $\pm 5.53$	0.0975 $\star$ $\pm 0.0188$	0.1046 $\star$ $\pm 0.0173$
Attention-based Fusion (Deep-Global Context)	90.45 $\pm 2.93$	85.83 $\pm 2.04$	88.04 $\star$ $\pm 1.65$	88.33 $\star$ $\pm 1.66$	90.83 $\pm 3.12$	0.1173 $\pm 0.0134$	0.1065 $\pm 0.0153$
<b>Introduced models (with label smoothing)</b>							
Co-Attention (Global Context)	88.65 $\pm 4.63$	88.33 $\pm 1.66$	88.39 $\star$ $\pm 1.76$	88.33 $\star$ $\pm 2.12$	88.33 $\pm 5.53$	0.1075 $\pm 0.0198$	0.1710 $\pm 0.0281$
Co-Attention (Deep Context)	93.57 $\pm 2.08$	84.16 $\pm 4.86$	88.53 $\star$ $\pm 2.79$	89.16 $\star$ $\pm 2.43$	94.16 $\pm 2.04$	0.1082 $\pm 0.0184$	0.1316 $\pm 0.0296$
Co-Attention (Deep-Global Context)	87.88 $\pm 3.73$	87.50 $\pm 6.97$	87.39 $\star$ $\pm 2.45$	87.50 $\dagger$ $\pm 1.86$	87.50 $\pm 4.56$	0.1176 $\pm 0.0167$	0.1568 $\pm 0.0306$
Attention-based Fusion (Global Context)	90.51 $\pm 3.40$	85.00 $\pm 4.25$	87.53 $\dagger$ $\pm 1.75$	87.92 $\dagger$ $\pm 1.56$	90.83 $\pm 4.08$	0.1094 $\pm 0.0086$	0.1168 $\pm 0.0099$
Attention-based Fusion (Deep Context)	93.08 $\pm 2.03$	89.17 $\pm 2.04$	91.06 $\pm 1.60$	91.25 $\pm 1.56$	93.33 $\pm 2.04$	0.0859 $\pm 0.0130$	0.0830 $\pm 0.0158$
Attention-based Fusion (Deep-Global Context)	89.87 $\pm 5.52$	83.33 $\pm 4.56$	86.20 $\dagger$ $\pm 0.90$	86.66 $\dagger$ $\pm 1.02$	90.00 $\pm 5.65$	0.1397 $\pm 0.0102$	0.1508 $\pm 0.0123$

our best performing model, namely Attention-based Fusion (Deep Context) or Co-Attention (Deep Context) with label smoothing, with the respective model without label smoothing.

## 6.1. ADReSS Challenge Dataset

Regarding our proposed models, one can observe that Attention-based Fusion (Deep Context) constitutes our best performing model outperforming all the other introduced models in all the evaluation metrics except Precision and

Specificity. Specifically, Attention-based Fusion (Deep Context) outperforms the other introduced models with label smoothing in Accuracy by 2.09-4.59%, in Recall by 0.84-5.84%, and in F1-score by 2.53-4.86%. Despite the fact that Attention-based Fusion (Deep Context) obtains a lower Precision score by other introduced models, it surpasses them in F1-score, which constitutes the weighted average of Precision and Recall. Although it achieves lower specificity scores by Co-Attention (Deep Context), it must be noted that in health related studies, F1-score is more important than Specificity, since high F1-score means that the model can detect better the AD patients, while high Specificity and low F1-score means that AD patients are misdiagnosed as non-AD ones. In addition, Co-Attention (Deep Context) constitutes our second best performing model attaining an Accuracy of 89.16%. It achieves the highest precision and specificity scores accounting for 93.57% and 94.16% respectively, while it achieves an F1-score of 88.53%. It outperforms all the introduced models, except Attention-based Fusion (Deep Context), in Accuracy by 0.83-2.50% and in F1-score by 0.14-2.33%. It outperforms all the models in Precision and Specificity by 0.49-5.69% and 0.83-6.66% respectively.

Next, we compare our introduced approaches with label smoothing with the ones without applying label smoothing. As one can easily observe, label smoothing leads to both performance improvement and better calibration of the proposed approaches. Specifically, we observe that Attention-based Fusion (Deep Context) with label smoothing obtains a higher Accuracy score than the one obtained by the respective model without label smoothing by 1.25%, Attention-based Fusion (Global Context) with label smoothing surpasses Attention-based Fusion (Global Context) without label smoothing in Accuracy by 0.42%, etc. In terms of the calibration metrics, namely ECE and ACE, one can observe that label smoothing leads to better calibrated models. For instance, Attention-based Fusion (Deep Context) with label smoothing obtains an ECE of 0.0859 and an ACE of 0.0830, which are significantly better than the ones obtained by Attention-based Fusion (Deep Context) without label smoothing by 0.0116 and 0.0216 respectively.

In comparison with the unimodal and multimodal state-of-the-art approaches, one can observe that our best performing model, namely Attention-based Fusion (Deep Context) with label smoothing, outperforms the research works in Accuracy by 1.25-18.33% and in F1-score by 1.12-21.30%. These differences in performance are attributable to the fact that our best performing model captures both the inter- and intra-modal interactions through the self-attention mechanisms and optimal transport domain adaptation methods, enhances the self-attention mechanism with contextual information, and applies label smoothing in contrast to the research initiatives. In addition, Co-Attention (Deep Context) outperforms the research works, except [26, 70], in Accuracy by 0.83-16.24%.

## 6.2. ADReSSo Challenge Dataset

As one can easily observe in Table 4, Co-Attention (Deep Context) with label smoothing constitutes our best performing model attaining an Accuracy of 85.35%, a Recall of 86.29%, and a F1-score of 85.27%. It surpasses the other introduced models (with label smoothing) in Accuracy by 1.41-4.22%, in Recall by 1.15-4.58%, and in F1-score by 1.92-4.32%. In addition, we observe that Co-Attention (Deep Context) with label smoothing achieves better performance than the one obtained by the respective model without label smoothing. Specifically, the Accuracy is improved by 2.25%, the Recall is improved by 2.29%, the F1-score presents a surge of 2.26%, the Precision is increased by 2.21%, and the Specificity is improved by 2.21%. In terms of the calibration metrics, we observe that the ECE is improved by 0.0171 (ASO test indicates almost stochastic dominance).

Comparing our introduced models with label smoothing with the ones without label smoothing, we observe that in most cases label smoothing contributes to both the performance improvement and better calibration. For instance, Co-Attention (Global Context) with label smoothing improves Accuracy by 1.12% compared with the respective model without label smoothing, while ECE and ACE are also improved by 0.0254 and 0.0387 respectively. Similarly, Attention-based Fusion (Deep Context) with label smoothing outperforms the respective model without label smoothing in F1-score and Accuracy by 2.27% and 2.53% respectively, while the ECE and ACE also present a decline of 0.0106 and 0.0077 respectively.

In comparison with the unimodal and multimodal baselines, we observe that our best performing model, namely Co-Attention (Deep Context) with label smoothing, outperforms these baselines in Accuracy by 0.84-5.35%. Also, it outperforms all the research works, except for [9], in F1-score by 0.34-6.74%. We observe that our best performing model attains a better performance than BERT (ASO test indicates stochastic dominance), verifying our initial hypothesis that both modalities, i.e., transcripts and audio files, contribute to a better performance. In addition, we observe that our second best performing model, namely Attention-based Fusion (Deep Context) outperforms some research works, except for [53, 71, 65, 54], in Accuracy by 0.84-3.94%.

## 7. Ablation Study

In this section, we run a series of ablation experiments using the ADReSS Challenge dataset to explore the effectiveness and robustness of the introduced architecture described in Section 4. We report the results of the ablation study in Tables 5 and 6.

First, we explore the effectiveness of the context-based self-attention. To do this, we remove the contextual information and exploit the conventional self-attention mechanism introduced by [78]. We observe that the Accuracy score drops from 91.25% to 87.08%, while the F1-score presents a decline of 4.60%. Also, we observe that the removal of

**Table 4**

Performance comparison among proposed models and state-of-the-art approaches on the ADReSSo Challenge test set. Reported values are mean  $\pm$  standard deviation. Results are averaged across five runs. ( $\dagger$ ) means that Co-Attention (Deep Context) with label smoothing is stochastically dominant over the respective models. ( $\star$ ) denotes almost stochastic dominance of the Co-Attention (Deep Context) with label smoothing over the respective approaches.

Architecture	Prec. (%)	Rec. (%)	F1-score (%)	Acc. (%)	Spec. (%)	ECE	ACE
<b>Baselines - Unimodal state-of-the-art approaches (only transcripts)</b>							
<i>BERT</i>	83.35 $\pm 0.86$	74.29 $\pm 2.55$	78.53 $\dagger$ $\pm 1.43$	80.00 $\dagger$ $\pm 1.05$	85.55 $\pm 1.11$	-	-
<i>Model C</i> : [60]	85.00	80.00	82.00	83.00	86.00	-	-
<i>Model 5</i> : [53]	81.58	88.57	84.93	84.51	80.56	-	-
<i>Label Fusion selected models</i> [71]	-	-	-	84.51	-	-	-
<i>Mp1</i> [92]	87.10	77.14	81.82	83.10	88.89	-	-
<b>Baselines - Multimodal state-of-the-art approaches</b>							
<i>LSTM w/ Gating (Words+Acoustic+Disf+Pse+WP)</i> [65]	-	-	-	84.00	-	-	-
<i>Global Fusion</i> [54]	92.00	74.00	83.00	84.51	94.00	-	-
<i>Top-10 Avg.</i> [9]	-	-	88.89	81.69	80.00	-	-
<i>Attempt 1</i> : [82]	75.00	91.67	82.50	80.28	68.57	-	-
<b>Baselines - Introduced models (without label smoothing)</b>							
<i>Co-Attention (Global Context)</i>	83.77 $\pm 4.59$	81.13 $\pm 9.13$	81.85 $\star$ $\pm 3.01$	82.54 $\star$ $\pm 1.69$	83.88 $\pm 6.66$	0.1536 $\pm 0.0311$	0.2017 $\pm 0.0214$
<i>Co-Attention (Deep Context)</i>	82.22 $\pm 1.79$	84.00 $\pm 4.28$	83.01 $\star$ $\pm 1.63$	83.10 $\star$ $\pm 1.26$	82.22 $\pm 2.83$	0.1349 $\star$ $\pm 0.0135$	0.1845 $\pm 0.0169$
<i>Co-Attention (Deep-Global Context)</i>	83.03 $\pm 2.07$	80.57 $\pm 2.79$	81.73 $\star$ $\pm 1.23$	82.25 $\dagger$ $\pm 1.13$	83.88 $\pm 2.72$	0.1414 $\pm 0.0091$	0.1948 $\pm 0.0265$
<i>Attention-based Fusion (Global Context)</i>	83.44 $\pm 1.16$	74.86 $\pm 2.14$	78.90 $\dagger$ $\pm 1.51$	80.28 $\dagger$ $\pm 1.26$	85.56 $\pm 1.11$	0.1633 $\pm 0.0207$	0.1825 $\pm 0.0140$
<i>Attention-based Fusion (Deep Context)</i>	81.52 $\pm 3.47$	81.14 $\pm 5.59$	81.08 $\dagger$ $\pm 1.58$	81.41 $\dagger$ $\pm 1.05$	81.66 $\pm 5.44$	0.1442 $\pm 0.0284$	0.1737 $\pm 0.0089$
<i>Attention-based Fusion (Deep-Global Context)</i>	79.58 $\pm 2.69$	85.71 $\pm 4.78$	82.38 $\star$ $\pm 1.59$	81.97 $\dagger$ $\pm 1.38$	78.33 $\pm 4.78$	0.1671 $\pm 0.0201$	0.1820 $\pm 0.0193$
<b>Introduced models (with label smoothing)</b>							
<i>Co-Attention (Global Context)</i>	84.77 $\pm 2.39$	81.71 $\pm 3.43$	83.12 $\star$ $\pm 0.95$	83.66 $\star$ $\pm 0.69$	85.55 $\pm 3.24$	0.1282 $\pm 0.0053$	0.1630 $\pm 0.0179$
<i>Co-Attention (Deep Context)</i>	84.43 $\pm 1.59$	86.29 $\pm 4.19$	85.27 $\pm 1.78$	85.35 $\pm 1.44$	84.43 $\pm 2.19$	0.1178 $\pm 0.0209$	0.1800 $\pm 0.0213$
<i>Co-Attention (Deep-Global Context)</i>	82.45 $\pm 0.99$	82.86 $\pm 4.78$	82.55 $\star$ $\pm 2.03$	82.82 $\star$ $\pm 1.38$	82.77 $\pm 2.08$	0.1443 $\pm 0.0046$	0.1749 $\pm 0.0082$
<i>Attention-based Fusion (Global Context)</i>	80.44 $\pm 1.65$	81.71 $\pm 4.98$	80.95 $\dagger$ $\pm 2.04$	81.13 $\dagger$ $\pm 1.44$	80.55 $\pm 3.04$	0.1540 $\pm 0.0195$	0.1920 $\pm 0.0215$
<i>Attention-based Fusion (Deep Context)</i>	85.10 $\pm 0.53$	81.71 $\pm 3.43$	83.35 $\star$ $\pm 2.04$	83.94 $\star$ $\pm 1.69$	86.11 $\pm 0.04$	0.1336 $\pm 0.0190$	0.1660 $\pm 0.0144$
<i>Attention-based Fusion (Deep-Global Context)</i>	81.45 $\pm 1.32$	85.14 $\pm 4.92$	83.18 $\star$ $\pm 2.42$	83.10 $\star$ $\pm 1.99$	81.11 $\pm 2.08$	0.1690 $\pm 0.0245$	0.1938 $\pm 0.0112$

contextual information yields to higher standard deviations of the performance metrics.

Next, we investigate the efficacy of the gate model, which is incorporated into the self-attention mechanism. To do this, we remove the gate model and exploit the conventional self-attention mechanism. We observe that Accuracy and F1-score present a decline of 2.50% and 2.84% respectively.

Moreover, we explore the effectiveness of the optimal transport domain adaptation method and the Optimal Transport Kernel. To do this, we remove these components from the introduced architecture. We observe that the Accuracy score is equal to 87.50%, which is lower by 3.75% than the one obtained by our best performing model. Also, this approach yields an F1-score accounting for 87.47%, which is lower by 3.59% than the one achieved by Attention-based Fusion (Deep Context).

Next, we explore the effectiveness of the Optimal Transport Kernel. To do this, we remove this component, exploit the average operation over the sequence length, and finally repeat the vector  $n$  times, so as to ensure that both the textual and image modalities have the same sequence length. As

one can observe, this method presents a decline in Accuracy score by 2.92%, while the F1-score is also reduced by 2.33%.

In addition, we explore the effectiveness of the fusion method. To prove this, we remove the fusion method, apply the average operation over  $C$  (Eq. 21) and  $S$  (Eq. 22) and concatenate these two representation vectors. We observe that the concatenation of features yields an Accuracy and F1-score of 87.50% and 87.65% respectively. This difference in performance can be justified by the fact that the concatenation operation does not capture the inherent correlations between the modalities.

Finally, we vary the layers of the context-based self-attention mechanism. The results of this ablation study are reported in Table 6. As the number of layers increases from 1 to 3, we observe that both the Accuracy and F1-score also increase. This justifies our initial hypothesis that stacking attention layers and fusing the outputs of different layers into one context vector, yields to better evaluation results, since the model captures both high-level and low-level syntactic and semantic information. However, we observe that the performance of our approach starts to present a decline by



**Table 5**

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

Architecture	Prec. (%)	Rec. (%)	F1-score (%)	Acc. (%)	Spec. (%)
<i>without contextual vector in self-attention</i>	91.34 $\pm 7.35$	83.33 $\pm 9.50$	86.46 $\pm 4.64$	87.08 $\pm 4.04$	90.83 $\pm 10.00$
<i>self-attention without gate model</i>	92.99 $\pm 4.28$	84.16 $\pm 3.12$	88.22 $\pm 0.88$	88.75 $\pm 1.02$	93.33 $\pm 4.25$
<i>without optimal transport and OTK</i>	87.60 $\pm 2.02$	87.50 $\pm 3.73$	87.47 $\pm 1.52$	87.50 $\pm 1.32$	87.50 $\pm 2.64$
<i>repeat vector instead of OTK</i>	86.08 $\pm 3.37$	91.66 $\pm 2.64$	88.73 $\pm 1.97$	88.33 $\pm 2.12$	85.00 $\pm 4.25$
<i>Concatenation - Without fusion</i>	87.23 $\pm 4.99$	88.33 $\pm 3.12$	87.65 $\pm 2.64$	87.50 $\pm 2.95$	86.66 $\pm 6.12$
<i>Proposed Framework</i>	<b>93.08</b> $\pm 2.03$	<b>89.17</b> $\pm 2.04$	<b>91.06</b> $\pm 1.60$	<b>91.25</b> $\pm 1.56$	<b>93.33</b> $\pm 2.04$

**Table 6**

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

Layers	Prec. (%)	Rec. (%)	F1-score (%)	Acc. (%)	Spec. (%)
1	90.37 $\pm 3.33$	83.33 $\pm 5.27$	86.52 $\pm 1.94$	87.08 $\pm 1.56$	90.83 $\pm 4.08$
2	88.09 $\pm 1.96$	91.66 $\pm 3.73$	89.77 $\pm 1.45$	89.58 $\pm 1.32$	87.50 $\pm 2.64$
<b>3</b> <b>(Our best performing model)</b>	<b>93.08</b> $\pm 2.03$	<b>89.17</b> $\pm 2.04$	<b>91.06</b> $\pm 1.60$	<b>91.25</b> $\pm 1.56$	<b>93.33</b> $\pm 2.04$
4	92.05 $\pm 3.70$	76.66 $\pm 5.65$	83.55 $\pm 4.28$	85.00 $\pm 3.58$	93.33 $\pm 3.33$
5	88.67 $\pm 4.20$	83.33 $\pm 3.73$	85.84 $\pm 2.86$	86.25 $\pm 2.83$	89.16 $\pm 4.25$

stacking four or five layers of context-based self-attention by applying the deep-context strategy. We assume that this decline in performance is attributable to the limited dataset used and consequently to the problem of overfitting.

## 8. Discussion

From the results obtained in this study, we found that:

- **Finding 1:** We proposed a context-based self-attention mechanism and exploited three approaches of adding contextual information to self-attention layers. Results on the ADReSS and ADReSSo Challenge datasets showed that the fusion of the outputs (low-level syntactic and semantic information) of different layers as a deep context vector yielded the highest evaluation results.
- **Finding 2:** We compared our proposed approaches with and without label smoothing. Findings suggested that label smoothing contributes to both the performance improvement and improvements in terms of the calibration metrics.
- **Finding 3:** We exploited two methods for fusing the self and cross-attention features. Findings of the experiments conducted on the ADReSS Challenge dataset suggested that the usage of two independent attentional reduction models, the add operation, and the layer normalization achieved better performance than the usage of a co-attention mechanism. On the other hand, results on the ADReSSo Challenge dataset

showed that the co-attention mechanism as a fusion method achieved the best evaluation results.

- **Finding 4:** Findings from a series of ablation studies showed the effectiveness and robustness of the introduced architecture.
- **Finding 5:** Our proposed models yielded competitive performances to the existing state-of-the-art approaches. We also used the Almost Stochastic Order test to test for statistical significance. This test does not make any assumptions about the distributions of the scores.
- **Finding 6:** We observed that in most cases the performance of the multi-modal models (baselines) was inferior to the transcript only BERT baseline. We hypothesize that this difference in performance is attributable to the fact that the multimodal approaches propose early and late fusion strategies or add / concatenate the representation vectors of different modalities during training. In this way, the inter-modal interactions cannot be captured effectively. This difference in performance justifies our initial motivation that more effective fusion methods must be explored for capturing the inter-modal interactions.

Our approaches entail some limitations, which are described below:

- **Hyperparameter Tuning:** In this study, we did not perform hyperparameter tuning due to the limited access to GPU resources. Hyperparameter tuning yields to an increase of the classification performance.

- **Explainability:** In this study, we did not apply explainability techniques, namely LIME, Integrated Gradients, and so on, for explaining the predictions of our introduced approaches.
- **Self-Supervised Learning:** Contrary to self-supervised learning methods, our approach relies heavily on plenty of training data.

## 9. Conclusion and Future Work

In this paper, we introduced some new approaches to detect AD patients from speech and transcripts, which capture the inter- and intra-modal interactions, enhance the conventional self-attention mechanism with contextual information, and deal with the problem of creating overconfident models by applying label smoothing. Our proposed architectures consist of BERT, DeiT, self-attention mechanism incorporating a gating model, context-based self-attention, optimal transport domain adaptation methods, and one new method for fusing the self and cross-attention features in the task of dementia detection from speech data. Furthermore, we designed extensive ablation experiments to explore the effectiveness of the components of the proposed architecture. Extensive experiments conducted on the ADReSS Challenge dataset demonstrate the efficacy of the proposed architectures reaching Accuracy and F1-score up to 91.25% and 91.06% respectively. Also, findings suggested that the label smoothing contributes to both the performance improvement and calibration of our model.

We evaluated our proposed approaches on the ADReSS Challenge dataset, which consists of a statistically balanced and acoustically enhanced set of recordings of spontaneous speech sessions. The proposed models can be embedded into an application, which will be capable of detecting AD patients with high accuracy. The creation of such an application appears to be very important especially in the era of covid-19, where there are restrictions for access to hospitals and medical centers.

In the future, we plan to exploit our methods in other health-related tasks, including the detection of Parkinson's disease. This will establish the generalizability of our introduced architectures. In addition, we aim to use automatic transcripts, since there are datasets, which do not include manual transcripts [40, 38]. Also, we plan to use the ADReSS-M Challenge dataset [38], where we will train our proposed approaches using automatic transcripts and audio files in the english language and test our introduced approach using automatic transcripts and audio files in the greek language. We also plan to employ more methods for measuring the uncertainty and calibration of our models. Applying self-supervised learning methods with limited speech data is one of our plans. In this way, our methods will not rely on labels and training data. Additionally, inspired by [36], we aim to exploit AlexNet instead of DeiT. Specifically, similar to [36], we plan to substitute the last several layers in AlexNet with extreme learning machine and optimize the extreme

learning machine with a chaotic bat algorithm. Finally, motivated by [37], we aim to apply some of the modules to our task, including the image-level representation learning algorithm, the universal neighboring-aware representation learning framework, and the graph random vector functional link.

## CRedit authorship contribution statement

**Loukas Ilias:** Conceptualization, Methodology, Software, Validation, Formal Analysis, Visualization, Investigation, Data Curation, Writing - Original Draft, Writing - Review & Editing. **Dimitris Askounis:** Supervision, Project Administration, Writing - Review & Editing.

## References

- [1] Al-Hameed, S., Benaissa, M., Christensen, H., 2017. Detecting and predicting alzheimer's disease severity in longitudinal acoustic data, in: Proceedings of the International Conference on Bioinformatics Research and Applications 2017, Association for Computing Machinery, New York, NY, USA. p. 57–61. URL: <https://doi.org/10.1145/3175587.3175589>, doi:10.1145/3175587.3175589.
- [2] Alzheimer's Society, 2021. *Dementia and language*. Available online at: <https://www.alzheimers.org.uk/about-dementia/symptoms-and-diagnosis/symptoms/dementia-and-language>.
- [3] Arevalo, J., Solorio, T., Montes-y Gomez, M., González, F.A., 2020. Gated multimodal networks. *Neural Computing and Applications*, 1–20.
- [4] Ba, J.L., Kiros, J.R., Hinton, G.E., 2016. Layer normalization. arXiv preprint arXiv:1607.06450.
- [5] Becker, J.T., Boiler, F., Lopez, O.L., Saxton, J., McGonigle, K.L., 1994. The Natural History of Alzheimer's Disease: Description of Study Cohort and Accuracy of Diagnosis. *Archives of Neurology* 51, 585–594. URL: <https://doi.org/10.1001/archneur.1994.00540180063015>, doi:10.1001/archneur.1994.00540180063015, arXiv:[https://jamanetwork.com/journals/jamaneurology/articlepdf/592905/archneur\\_](https://jamanetwork.com/journals/jamaneurology/articlepdf/592905/archneur_)
- [6] Bertini, F., Allevi, D., Lutero, G., Calzà, L., Montesi, D., 2022. An automatic alzheimer's disease classifier based on spontaneous spoken english. *Computer Speech & Language* 72, 101298. URL: <https://www.sciencedirect.com/science/article/pii/S0885230821000991>, doi:<https://doi.org/10.1016/j.csl.2021.101298>.
- [7] Cai, Y., Cai, H., Wan, X., 2019. Multi-modal sarcasm detection in twitter with hierarchical fusion model, in: Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics, pp. 2506–2515.
- [8] Chen, C., Han, D., Chang, C.C., 2022. Caan: Context-aware attention network for visual question answering. *Pattern Recognition* 132, 108980. URL: <https://www.sciencedirect.com/science/article/pii/S0031320322004605>, doi:<https://doi.org/10.1016/j.patcog.2022.108980>.
- [9] Chen, J., Ye, J., Tang, F., Zhou, J., 2021. Automatic Detection of Alzheimer's Disease Using Spontaneous Speech Only, in: Proc. Interspeech 2021, pp. 3830–3834. doi:10.21437/Interspeech.2021-2002.
- [10] Chollet, F., 2017. Xception: Deep learning with depthwise separable convolutions, in: 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pp. 1800–1807. doi:10.1109/CVPR.2017.195.
- [11] Courty, N., Flamary, R., Tuia, D., Rakotomamonjy, A., 2017. Optimal transport for domain adaptation. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 39, 1853–1865. doi:10.1109/TPAMI.2016.2615921.
- [12] Crowson, C.S., Atkinson, E.J., Therneau, T.M., 2016. Assessing calibration of prognostic risk scores. *Statistical Methods in Medical Research* 25, 1692–1706. URL: <https://doi.org/10.1093/oxfordjournals.smi.a017921>.

- //doi.org/10.1177/0962280213497434, doi:10.1177/0962280213497434, arXiv:https://doi.org/10.1177/0962280213497434, pMID: 23907781.
- [13] Cummins, N., Pan, Y., Ren, Z., Fritsch, J., Nallanthighal, V.S., Christensen, H., Blackburn, D., Schuller, B.W., Magimai-Doss, M., Strik, H., et al., 2020. A comparison of acoustic and linguistics methodologies for alzheimer's dementia recognition, in: Interspeech 2020, ISCA-International Speech Communication Association. pp. 2182–2186.
- [14] Dawid, A.P., 1982. The well-calibrated bayesian. *Journal of the American Statistical Association* 77, 605–610. URL: <https://www.tandfonline.com/doi/abs/10.1080/01621459.1982.10477856>, doi:10.1080/01621459.1982.10477856, arXiv:https://www.tandfonline.com/doi/pdf/10.1080/01621459.1982.10477856.
- [15] Del Barrio, E., Cuesta-Albertos, J.A., Matrán, C., 2018. An optimal transportation approach for assessing almost stochastic order, in: *The Mathematics of the Uncertain*. Springer, pp. 33–44.
- [16] Devlin, J., Chang, M.W., Lee, K., Toutanova, K., 2019. BERT: Pre-training of deep bidirectional transformers for language understanding, in: *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, Association for Computational Linguistics, Minneapolis, Minnesota. pp. 4171–4186. URL: <https://aclanthology.org/N19-1423>, doi:10.18653/v1/N19-1423.
- [17] Dror, R., Shlomov, S., Reichart, R., 2019. Deep dominance - how to properly compare deep neural models, in: Korhonen, A., Traum, D.R., Márquez, L. (Eds.), *Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019, Florence, Italy, July 28- August 2, 2019, Volume 1: Long Papers*, Association for Computational Linguistics. pp. 2773–2785. URL: <https://doi.org/10.18653/v1/p19-1266>, doi:10.18653/v1/p19-1266.
- [18] Edwards, E., Dognin, C., Bollepalli, B., Singh, M., 2020. Multiscale System for Alzheimer's Dementia Recognition Through Spontaneous Speech, in: *Proc. Interspeech 2020*, pp. 2197–2201. doi:10.21437/Interspeech.2020-2781.
- [19] Ferradans, S., Papadakis, N., Peyré, G., Aujol, J.F., 2014. Regularized discrete optimal transport. *SIAM Journal on Imaging Sciences* 7, 1853–1882.
- [20] Flamary, R., Courty, N., Gramfort, A., Alaya, M.Z., Boisbunon, A., Chambon, S., Chapel, L., Corenflos, A., Fatras, K., Fournier, N., Gautheron, L., Gayraud, N.T., Janati, H., Rakotomamonjy, A., Redko, I., Rolet, A., Schutz, A., Seguy, V., Sutherland, D.J., Tavenard, R., Tong, A., Vayer, T., 2021. Pot: Python optimal transport. *Journal of Machine Learning Research* 22, 1–8. URL: <http://jmlr.org/papers/v22/20-451.html>.
- [21] Freitag, M., Amiriparian, S., Pugachevskiy, S., Cummins, N., Schuller, B., 2017. audeep: Unsupervised learning of representations from audio with deep recurrent neural networks. *The Journal of Machine Learning Research* 18, 6340–6344.
- [22] Gu, Y., Yang, K., Fu, S., Chen, S., Li, X., Marsic, I., 2018. Hybrid attention based multimodal network for spoken language classification, in: *Proceedings of the 27th International Conference on Computational Linguistics*, Association for Computational Linguistics, Santa Fe, New Mexico, USA. pp. 2379–2390. URL: <https://aclanthology.org/C18-1201>.
- [23] Guo, C., Pleiss, G., Sun, Y., Weinberger, K.Q., 2017. On calibration of modern neural networks, in: Precup, D., Teh, Y.W. (Eds.), *Proceedings of the 34th International Conference on Machine Learning*, PMLR. pp. 1321–1330. URL: <https://proceedings.mlr.press/v70/guo17a.html>.
- [24] Haulcy, R., Glass, J., 2021. Classifying alzheimer's disease using audio and text-based representations of speech. *Frontiers in Psychology* 11, 3833. URL: <https://www.frontiersin.org/article/10.3389/fpsyg.2020.624137>, doi:10.3389/fpsyg.2020.624137.
- [25] Ilias, L., Askounis, D., 2022a. Explainable identification of dementia from transcripts using transformer networks. *IEEE Journal of Biomedical and Health Informatics* 26, 4153–4164. doi:10.1109/JBHI.2022.3172479.
- [26] Ilias, L., Askounis, D., 2022b. Multimodal deep learning models for detecting dementia from speech and transcripts. *Frontiers in Aging Neuroscience* 14. URL: <https://www.frontiersin.org/articles/10.3389/fnagi.2022.830943>, doi:10.3389/fnagi.2022.830943.
- [27] Ilias, L., Askounis, D., Psarras, J., 2022. A multimodal approach for dementia detection from spontaneous speech with tensor fusion layer, in: *2022 IEEE-EMBS International Conference on Biomedical and Health Informatics (BHI)*, pp. 1–5. doi:10.1109/BHI56158.2022.9926818.
- [28] Ilias, L., Askounis, D., Psarras, J., 2023. Detecting dementia from speech and transcripts using transformers. *Computer Speech & Language* 79, 101485. URL: <https://www.sciencedirect.com/science/article/pii/S0885230823000049>, doi:https://doi.org/10.1016/j.csl.2023.101485.
- [29] Jiang, X., Osl, M., Kim, J., Ohno-Machado, L., 2011. Calibrating predictive model estimates to support personalized medicine. *Journal of the American Medical Informatics Association* 19, 263–274. URL: <https://doi.org/10.1136/amiajnl-2011-000291>, doi:10.1136/amiajnl-2011-000291, arXiv:https://academic.oup.com/jamia/article-pdf/19/2/263/17374049/19-2-263.pdf.
- [30] Karlekar, S., Niu, T., Bansal, M., 2018. Detecting linguistic characteristics of Alzheimer's dementia by interpreting neural models, in: *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, Association for Computational Linguistics, New Orleans, Louisiana. pp. 701–707. URL: <https://aclanthology.org/N18-2110>, doi:10.18653/v1/N18-2110.
- [31] Khodabakhsh, A., Kuşxuoğlu, S., Demiroğlu, C., 2014. Natural language features for detection of alzheimer's disease in conversational speech, in: *IEEE-EMBS International Conference on Biomedical and Health Informatics (BHI)*, pp. 581–584. doi:10.1109/BHI.2014.6864431.
- [32] Koo, J., Lee, J.H., Pyo, J., Jo, Y., Lee, K., 2020. Exploiting Multi-Modal Features from Pre-Trained Networks for Alzheimer's Dementia Recognition, in: *Proc. Interspeech 2020*, pp. 2217–2221. doi:10.21437/Interspeech.2020-3153.
- [33] Kumar, K., Kim, C., Stern, R.M., 2011. Delta-spectral cepstral coefficients for robust speech recognition, in: *2011 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 4784–4787. doi:10.1109/ICASSP.2011.5947425.
- [34] Lee, J.L., Burkholder, R., Flinn, G.B., Coppess, E.R., 2016. Working with CHAT transcripts in Python. Technical Report TR-2016-02. Department of Computer Science, University of Chicago.
- [35] Lu, J., Yang, J., Batra, D., Parikh, D., 2016. Hierarchical question-attention co-attention for visual question answering, in: *Proceedings of the 30th International Conference on Neural Information Processing Systems*, Curran Associates Inc., Red Hook, NY, USA. p. 289–297.
- [36] Lu, S., Wang, S.H., Zhang, Y.D., 2021. Detection of abnormal brain in mri via improved alexnet and elm optimized by chaotic bat algorithm. *Neural Computing and Applications* 33, 10799–10811.
- [37] Lu, S., Zhu, Z., Gorriz, J.M., Wang, S.H., Zhang, Y.D., 2022. Nagnn: Classification of covid-19 based on neighboring aware representation from deep graph neural network. *International Journal of Intelligent Systems* 37, 1572–1598. URL: <https://onlinelibrary.wiley.com/doi/abs/10.1002/int.22686>, doi:https://doi.org/10.1002/int.22686, arXiv:https://onlinelibrary.wiley.com/doi/pdf/10.1002/int.22686.
- [38] Luz, S., Haider, F., Fromm, D., Lazarou, I., Kompatsiaris, I., MacWhinney, B., 2023. Multilingual alzheimer's dementia recognition through spontaneous speech: a signal processing grand challenge. arXiv preprint arXiv:2301.05562.
- [39] Luz, S., Haider, F., de la Fuente, S., Fromm, D., MacWhinney, B., 2020. Alzheimer's Dementia Recognition Through Spontaneous Speech: The ADReSS Challenge, in: *Proc. Interspeech 2020*, pp. 2172–2176. doi:10.21437/Interspeech.2020-2571.
- [40] Luz, S., Haider, F., de la Fuente, S., Fromm, D., MacWhinney, B., 2021. Detecting Cognitive Decline Using Speech Only: The ADReSSo Challenge, in: *Proc. Interspeech 2021*, pp. 3780–3784. doi:10.21437/Interspeech.2021-1220.

- [41] MacWhinney, B., 2000. The CHILDES Project: Tools for Analyzing Talk (third edition): Volume I: Transcription format and programs, Volume II: The database. Computational Linguistics 26, 657–657. URL: <https://doi.org/10.1162/coli.2000.26.4.657>, doi:10.1162/coli.2000.26.4.657, arXiv:<https://direct.mit.edu/coli/article-pdf/26/4/657/1797615/coli.2000.26.4.657>
- [42] Mahajan, P., Baths, V., 2021. Acoustic and language based deep learning approaches for alzheimer’s dementia detection from spontaneous speech. Frontiers in Aging Neuroscience 13, 20. URL: <https://www.frontiersin.org/article/10.3389/fnagi.2021.623607>, doi:10.3389/fnagi.2021.623607.
- [43] Martinc, M., Pollak, S., 2020. Tackling the ADReSS Challenge: A Multimodal Approach to the Automated Recognition of Alzheimer’s Dementia, in: Proc. Interspeech 2020, pp. 2157–2161. doi:10.21437/Interspeech.2020-2202.
- [44] McFee, B., Metsai, A., McVicar, M., Balke, S., Thomé, C., Raffel, C., Zalkow, F., Malek, A., Dana, Lee, K., Nieto, O., Ellis, D., Mason, J., Battenberg, E., Seyfarth, S., Yamamoto, R., viktorandreevichmorozov, Choi, K., Moore, J., Bittner, R., Hidaka, S., Wei, Z., nullmightybofo, Weiss, A., Hereñú, D., Stöter, F.R., Nickel, L., Friesch, P., Vollrath, M., Kim, T., 2022. librosa/librosa: 0.9.2. URL: <https://doi.org/10.5281/zenodo.6759664>, doi:10.5281/zenodo.6759664.
- [45] McFee, B., Raffel, C., Liang, D., Ellis, D.P., McVicar, M., Battenberg, E., Nieto, O., 2015. librosa: Audio and music signal analysis in python, in: Proceedings of the 14th python in science conference, pp. 18–25.
- [46] Mialon, G., Chen, D., d’Aspremont, A., Mairal, J., 2021. A trainable optimal transport embedding for feature aggregation and its relationship to attention, in: International Conference on Learning Representations.
- [47] Mittal, A., Sahoo, S., Datar, A., Kadiwala, J., Shalu, H., Mathew, J., 2021. Multi-modal detection of alzheimer’s disease from speech and text. arXiv:2012.00096.
- [48] Müller, R., Kornblith, S., Hinton, G.E., 2019. When does label smoothing help?, in: Wallach, H., Larochelle, H., Beygelzimer, A., d’Alché-Buc, F., Fox, E., Garnett, R. (Eds.), Advances in Neural Information Processing Systems, Curran Associates, Inc. URL: <https://proceedings.neurips.cc/paper/2019/file/f1748d6b0fd9d439f71450117eba2725-Paper.pdf>.
- [49] Murphy, A.H., Epstein, E.S., 1967. Verification of probabilistic predictions: A brief review. Journal of Applied Meteorology and Climatology 6, 748–755. URL: [https://journals.ametsoc.org/view/journals/apme/6/5/1520-0450\\_1967\\_006\\_0748\\_voppab\\_2\\_0\\_co\\_2.xml](https://journals.ametsoc.org/view/journals/apme/6/5/1520-0450_1967_006_0748_voppab_2_0_co_2.xml), doi:10.1175/1520-0450(1967)006<0748:VOPPAB>2.0.CO;2.
- [50] Naeini, M.P., Cooper, G., Hauskrecht, M., 2015. Obtaining well calibrated probabilities using bayesian binning, in: Twenty-Ninth AAAI Conference on Artificial Intelligence.
- [51] Nixon, J., Dusenberry, M.W., Zhang, L., Jerfel, G., Tran, D., 2019. Measuring calibration in deep learning, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops.
- [52] Pan, H., Lin, Z., Fu, P., Qi, Y., Wang, W., 2020. Modeling intra and inter-modality incongruity for multi-modal sarcasm detection, in: Findings of the Association for Computational Linguistics: EMNLP 2020, Association for Computational Linguistics, Online. pp. 1383–1392. URL: <https://aclanthology.org/2020.findings-emnlp.124>, doi:10.18653/v1/2020.findings-emnlp.124.
- [53] Pan, Y., Mirheidari, B., Harris, J.M., Thompson, J.C., Jones, M., Snowden, J.S., Blackburn, D., Christensen, H., 2021. Using the Outputs of Different Automatic Speech Recognition Paradigms for Acoustic- and BERT-Based Alzheimer’s Dementia Detection Through Spontaneous Speech, in: Proc. Interspeech 2021, pp. 3810–3814. doi:10.21437/Interspeech.2021-1519.
- [54] Pappagari, R., Cho, J., Joshi, S., Moro-Velázquez, L., Želasko, P., Villalba, J., Dehak, N., 2021. Automatic Detection and Assessment of Alzheimer Disease Using Speech and Language Technologies in Low-Resource Scenarios, in: Proc. Interspeech 2021, pp. 3825–3829. doi:10.21437/Interspeech.2021-1850.
- [55] Pappagari, R., Cho, J., Moro-Velázquez, L., Dehak, N., 2020. Using State of the Art Speaker Recognition and Natural Language Processing Technologies to Detect Alzheimer’s Disease and Assess its Severity, in: Proc. Interspeech 2020, pp. 2177–2181. doi:10.21437/Interspeech.2020-2587.
- [56] Park, D.S., Chan, W., Zhang, Y., Chiu, C.C., Zoph, B., Cubuk, E.D., Le, Q.V., 2019. SpecAugment: A Simple Data Augmentation Method for Automatic Speech Recognition, in: Proc. Interspeech 2019, pp. 2613–2617. doi:10.21437/Interspeech.2019-2680.
- [57] Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Kopf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J., Chintala, S., 2019. Pytorch: An imperative style, high-performance deep learning library, in: Wallach, H., Larochelle, H., Beygelzimer, A., d’Alché-Buc, F., Fox, E., Garnett, R. (Eds.), Advances in Neural Information Processing Systems 32. Curran Associates, Inc., pp. 8024–8035. URL: <http://papers.neurips.cc/paper/9015-pytorch-an-imperative-style-high-performance-deep-learning-library.pdf>.
- [58] Pospili, A., Rolland, T., Abad, A., 2020. The INESC-ID Multi-Modal System for the ADReSS 2020 Challenge, in: Proc. Interspeech 2020, pp. 2202–2206. doi:10.21437/Interspeech.2020-2833.
- [59] Pramanick, S., Roy, A., Patel, V.M., 2022. Multimodal learning using optimal transport for sarcasm and humor detection, in: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV), pp. 3930–3940.
- [60] Qiao, Y., Yin, X., Wiechmann, D., Kerz, E., 2021. Alzheimer’s Disease Detection from Spontaneous Speech Through Combining Linguistic Complexity and (Dis)Fluency Features with Pretrained Language Models, in: Proc. Interspeech 2021, pp. 3805–3809. doi:10.21437/Interspeech.2021-1415.
- [61] Radford, A., Kim, J.W., Xu, T., Brockman, G., McLeavey, C., Sutskever, I., 2022. Robust speech recognition via large-scale weak supervision. arXiv preprint arXiv:2212.04356.
- [62] Raghu, M., Blumer, K., Sayres, R., Obermeyer, Z., Kleinberg, B., Mullainathan, S., Kleinberg, J., 2019. Direct uncertainty prediction for medical second opinions, in: Chaudhuri, K., Salakhutdinov, R. (Eds.), Proceedings of the 36th International Conference on Machine Learning, PMLR. pp. 5281–5290. URL: <https://proceedings.mlr.press/v97/raghu19a.html>.
- [63] Reimers, N., Gurevych, I., 2018. Why comparing single performance scores does not allow to draw conclusions about machine learning approaches. arXiv preprint arXiv:1803.09578.
- [64] Rohanian, M., Hough, J., Purver, M., 2020. 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. doi:10.21437/Interspeech.2020-2721.
- [65] Rohanian, M., Hough, J., Purver, M., 2021. Alzheimer’s Dementia Recognition Using Acoustic, Lexical, Disfluency and Speech Pause Features Robust to Noisy Inputs, in: Proc. Interspeech 2021, pp. 3820–3824. doi:10.21437/Interspeech.2021-1633.
- [66] Sánchez Villegas, D., Aletras, N., 2021. Point-of-interest type prediction using text and images, in: Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing, Association for Computational Linguistics, Online and Punta Cana, Dominican Republic. pp. 7785–7797. URL: <https://aclanthology.org/2021.emnlp-main.614>, doi:10.18653/v1/2021.emnlp-main.614.
- [67] Sánchez Villegas, D., Mokaram, S., Aletras, N., 2021. Analyzing online political advertisements, in: Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021, Association for Computational Linguistics, Online. pp. 3669–3680. URL: <https://aclanthology.org/2021.findings-acl.321>, doi:10.18653/v1/2021.findings-acl.321.
- [68] Sarawgi, U., Zulfikar, W., Soliman, N., Maes, P., 2020. Multimodal Inductive Transfer Learning for Detection of Alzheimer’s Dementia and its Severity, in: Proc. Interspeech 2020, pp. 2212–2216. doi:10.21437/Interspeech.2020-3137.

- [69] Shah, Z., Sawalha, J., Tasnim, M., Qi, S.a., Stroulia, E., Greiner, R., 2021. Learning language and acoustic models for identifying alzheimer's dementia from speech. *Frontiers in Computer Science* 3, 4. URL: <https://www.frontiersin.org/article/10.3389/fcomp.2021.624659>, doi:10.3389/fcomp.2021.624659.
- [70] Syed, Z.S., Syed, M.S.S., Lech, M., Pirogova, E., 2021a. Automated recognition of alzheimer's dementia using bag-of-deep-features and model ensembling. *IEEE Access* 9, 88377–88390. doi:10.1109/ACCESS.2021.3090321.
- [71] Syed, Z.S., Syed, M.S.S., Lech, M., Pirogova, E., 2021b. Tackling the ADRESSO Challenge 2021: The MUET-RMIT System for Alzheimer's Dementia Recognition from Spontaneous Speech, in: *Proc. Interspeech 2021*, pp. 3815–3819. doi:10.21437/Interspeech.2021-1572.
- [72] Szegedy, C., Vanhoucke, V., Ioffe, S., Shlens, J., Wojna, Z., 2016. Rethinking the inception architecture for computer vision, in: *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 2818–2826. doi:10.1109/CVPR.2016.308.
- [73] Tan, M., Le, Q., 2019. Efficientnet: Rethinking model scaling for convolutional neural networks, in: *International Conference on Machine Learning*, PMLR, pp. 6105–6114.
- [74] Touvron, H., Cord, M., Douze, M., Massa, F., Sablayrolles, A., Jegou, H., 2021. Training data-efficient image transformers & distillation through attention, in: Meila, M., Zhang, T. (Eds.), *Proceedings of the 38th International Conference on Machine Learning*, PMLR, pp. 10347–10357. URL: <https://proceedings.mlr.press/v139/touvron21a.html>.
- [75] Tsai, Y.H.H., Bai, S., Liang, P.P., Kolter, J.Z., Morency, L.P., Salakhutdinov, R., 2019. Multimodal transformer for unaligned multimodal language sequences, in: *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, Association for Computational Linguistics, Florence, Italy, pp. 6558–6569. URL: <https://aclanthology.org/P19-1656>, doi:10.18653/v1/P19-1656.
- [76] Tu, Z., Liu, Y., Lu, Z., Liu, X., Li, H., 2017. Context gates for neural machine translation. *Transactions of the Association for Computational Linguistics* 5, 87–99. URL: <https://aclanthology.org/Q17-1007>, doi:10.1162/tac1\_a\_00048.
- [77] Ulmer, D., Hardmeier, C., Frellsen, J., 2022. deep-significance-easy and meaningful statistical significance testing in the age of neural networks. *arXiv preprint arXiv:2204.06815*.
- [78] Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I., 2017. Attention is all you need, in: *Proceedings of the 31st International Conference on Neural Information Processing Systems*, Curran Associates Inc., Red Hook, NY, USA, p. 6000–6010.
- [79] Villani, C., 2008. Optimal transport, old and new. notes for the 2005 saint-flour summer school. *Grundlehren der mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]*. Springer.
- [80] Voita, E., Serdyukov, P., Sennrich, R., Titov, I., 2018. Context-aware neural machine translation learns anaphora resolution, in: *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, Association for Computational Linguistics, Melbourne, Australia, pp. 1264–1274. URL: <https://aclanthology.org/P18-1117>, doi:10.18653/v1/P18-1117.
- [81] Wang, L., Tu, Z., Way, A., Liu, Q., 2017. Exploiting cross-sentence context for neural machine translation, in: *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, Association for Computational Linguistics, Copenhagen, Denmark, pp. 2826–2831. URL: <https://aclanthology.org/D17-1301>, doi:10.18653/v1/D17-1301.
- [82] Wang, N., Cao, Y., Hao, S., Shao, Z., Subbalakshmi, K., 2021. Modular Multi-Modal Attention Network for Alzheimer's Disease Detection Using Patient Audio and Language Data, in: *Proc. Interspeech 2021*, pp. 3835–3839. doi:10.21437/Interspeech.2021-2024.
- [83] Wolf, T., Debut, L., Sanh, V., Chaumond, J., Delangue, C., Moi, A., Cistac, P., Rault, T., Louf, R., Funtowicz, M., Davison, J., Shleifer, S., von Platen, P., Ma, C., Jernite, Y., Plu, J., Xu, C., Scao, T.L., Gugger, S., Drame, M., Lhoest, Q., Rush, A.M., 2020. Transformers: State-of-the-art natural language processing, in: *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, Association for Computational Linguistics, Online, pp. 38–45. URL: <https://www.aclweb.org/anthology/2020.emnlp-demos.6>.
- [84] World Health Organization, 2021. *Dementia*. Available online at: <https://www.who.int/news-room/fact-sheets/detail/dementia>.
- [85] Yang, B., Li, J., Wong, D.F., Chao, L.S., Wang, X., Tu, Z., 2019. Context-aware self-attention networks. *Proceedings of the AAAI Conference on Artificial Intelligence* 33, 387–394. URL: <https://ojs.aaai.org/index.php/AAAI/article/view/3809>, doi:10.1609/aaai.v33i01.3301387.
- [86] Yang, L., Wei, W., Li, S., Li, J., Shinozaki, T., 2022. Augmented Adversarial Self-Supervised Learning for Early-Stage Alzheimer's Speech Detection, in: *Proc. Interspeech 2022*, pp. 541–545. doi:10.21437/Interspeech.2022-943.
- [87] Yu, Z., Cui, Y., Yu, J., Tao, D., Tian, Q., 2019a. Multimodal unified attention networks for vision-and-language interactions. *arXiv preprint arXiv:1908.04107*.
- [88] Yu, Z., Yu, J., Cui, Y., Tao, D., Tian, Q., 2019b. Deep modular co-attention networks for visual question answering, in: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- [89] Zadeh, A., Chen, M., Poria, S., Cambria, E., Morency, L.P., 2017. Tensor fusion network for multimodal sentiment analysis, in: *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, Association for Computational Linguistics, Copenhagen, Denmark, pp. 1103–1114. URL: <https://aclanthology.org/D17-1115>, doi:10.18653/v1/D17-1115.
- [90] Zhang, B., Xiong, D., Su, J., Duan, H., 2017. A context-aware recurrent encoder for neural machine translation. *IEEE/ACM Transactions on Audio, Speech, and Language Processing* 25, 2424–2432. doi:10.1109/TASLP.2017.2751420.
- [91] Zhu, Y., Liang, X., Batsis, J.A., Roth, R.M., 2021a. Exploring deep transfer learning techniques for alzheimer's dementia detection. *Frontiers in Computer Science* 3, 22. URL: <https://www.frontiersin.org/article/10.3389/fcomp.2021.624683>, doi:10.3389/fcomp.2021.624683.
- [92] Zhu, Y., Obyat, A., Liang, X., Batsis, J.A., Roth, R.M., 2021b. WavBERT: Exploiting Semantic and Non-Semantic Speech Using Wav2vec and BERT for Dementia Detection, in: *Proc. Interspeech 2021*, pp. 3790–3794. doi:10.21437/Interspeech.2021-332.