

Prediction of Depression Severity Based on the Prosodic and Semantic Features With Bidirectional LSTM and Time Distributed CNN

Kaining Mao^{ID}, *Student Member, IEEE*, Wei Zhang, Deborah Baofeng Wang, Ang Li, Rongqi Jiao, Yanhui Zhu, Bin Wu, Tiansheng Zheng, Lei Qian^{ID}, Wei Lyu, Minjie Ye, and Jie Chen^{ID}, *Fellow, IEEE*

Abstract—Depression is increasingly impacting individuals both physically and psychologically worldwide. It has become a global major public health problem and attracts attention from various research fields. Traditionally, the diagnosis of depression is formulated through semi-structured interviews and supplementary questionnaires, which makes the diagnosis heavily relying on physicians' experience and is subject to bias. However, since the pathogenic mechanism of depression is still under investigation, it is difficult for physicians to diagnose and treat, especially in the early clinical stage. As smart devices and artificial intelligence advance rapidly, understanding how depression associates with daily behaviors can be beneficial for the early stage depression diagnosis, which reduces labor costs and the likelihood of clinical mistakes as well as physicians bias. Furthermore, mental health monitoring and cloud-based remote diagnosis can be implemented through an automated depression diagnosis system. In this article, we propose an attention-based multimodality speech and text representation for depression prediction. Our model is trained to estimate the depression severity of participants using the Distress Analysis Interview Corpus-Wizard of Oz (DAIC-WOZ) dataset. For the audio modality, we use the collaborative voice analysis repository (COVAREP) features provided by the dataset and employ a Bidirectional Long Short-Term Memory Network (Bi-LSTM) followed by a Time-distributed Convolutional Neural Network (T-CNN). For the text modality, we use global vectors for word representation (GloVe) to perform word embeddings and the embeddings are fed into the Bi-LSTM network. Results show that both audio and text models perform well on the depression severity estimation task, with best sequence level F_1 score of 0.9870 and patient-level F_1 score of 0.9074 for the audio model over five classes (healthy, mild, moderate, moderately severe, and severe), as well as sequence level F_1 score of 0.9709 and patient-level F_1 score of 0.9245 for the text model over five classes. Results are similar for the multimodality fused model, with the highest F_1 score of 0.9580 on the patient-level depression detection task over five classes. Experiments show statistically significant improvements over previous works.

Index Terms—Artificial intelligence, depression, machine learning, mental health, natural language processing, neural network

1 INTRODUCTION

MENTAL health disorder, such as depression, is considered one of the major challenges facing global society. During the COVID-19 pandemic, the prevalence of

depression and anxiety is exacerbated in the general population [1], [2], [3], [4]. By 2030, depression will be the second major cause of disability worldwide and thus it can impose a heavy healthcare burden globally [5]. It is estimated that the average cost of treating depression in 2010 is 24,000 € per patient and the total cost can be as high as € 92 billion in Europe [6]. In the United States, depression causes an estimated loss of \$44 billion, due to the absence or low working efficiency [7]. According to a report from the World Health Organization (WHO), over 264 million people of all ages suffer from depression in 2017 [8]. Nearly 50% of people with depression worldwide have difficulty receiving therapy [9]. Suicide is one of the severe results of depression, and the WHO reports that the number of people who passed away due to suicide is over 800,000 every year [10]. The number of attempted suicide is more frequent, possibly no less than 20 times that of those who died by suicide [10]. Patients with depression are more apt to generate suicide thoughts [11], [12]. It is estimated that more than 50% of people who died by suicide meet clinical criteria of depression [13], [14]. However, often the symptoms of depression are not displayed directly. Many individuals often express their sadness and hopelessness but without depression, whereas patients are usually reluctant to report their conditions and receive treatment [15]. For instance, many people

- Kaining Mao, Wei Zhang, and Ang Li are with the Department of Electrical and Computer Engineering, University of Alberta, Edmonton, Alberta AB, T6G 2R3, Canada. E-mail: {mmao1, wzhang6, ang6}@ualberta.ca.
- Deborah Baofeng Wang, Tiansheng Zheng, Lei Qian, and Wei Lyu are with the Affiliated Kangning Hospital of Wenzhou Medical University, Wenzhou, Zhejiang 325015, China. E-mail: debbiebwang@msn.com, {sky.tsh, lw19760309}@163.com, leiqianwz@foxmail.com.
- Rongqi Jiao, Yanhui Zhu, and Bin Wu are with the Wenzhou Medical University, Wenzhou, Zhejiang 325015, China. E-mail: {lovejrj321, rebekahhui, wubin_1995}@163.com.
- Minjie Ye is with the Affiliated Kangning Hospital of Wenzhou Medical University, Wenzhou, Zhejiang 325015, China. E-mail: minjiejye@163.com.
- Jie Chen is with the Department of Electrical and Computer Engineering, University of Alberta, Edmonton, Alberta AB, T6G 2R3, Canada. E-mail: jc65@ualberta.ca.

Manuscript received 9 Dec. 2021; revised 4 Feb. 2022; accepted 21 Feb. 2022. Date of publication 24 Feb. 2022; date of current version 13 Sept. 2023. This research is also supported by the China Scholarship Council (CSC) No. 202000810031.

(Corresponding authors: Minjie Ye and Jie Chen.)

Recommended for acceptance by K. Qian.

Digital Object Identifier no. 10.1109/TAFFC.2022.3154332

with depression ignore or refuse to admit their emotional instability and physical health conditions. The reason is that depression is a stigmatized disease, resulting in the depressive population hiding or camouflaging their symptoms. Traditionally, a semi-structured clinical interview based on Diagnostic Statistical Manual (DSM) criteria is the standard protocol for depression diagnosis [16] with self-test questionnaires such as the Patient Health Questionnaire Depression Scale (PHQ) [17], Beck's Depression Inventory (BDI) [18] and Montgomery-Asberg Depression Rating Scale (MADRS) [19]. The PHQ-8 is an assessment form created to examine the existence of the core depression symptoms, such as fatigue and anxiety. The PHQ-8 scale shows high sensitivity and specificity for diagnosing depression and other mental disorders among patients with different languages and cultures [20]. These methods play a key role in diagnosing depression, but the results are subject to physicians' experience. Previous articles argued that these clinical criteria, such as DSM and BDI, are not reliable enough [21]. Diagnosis of depression is not the same as other medical conditions since gold standards for mental disorders do not exist currently, which raises the likelihood of misdiagnosis and finally leads to unexpected results [22], [23], [24]. However, most depressed people do not have access to qualified psychological treatment due to economic conditions (low-/mid-income population) or living constraints (in rural regions) [25]. Currently, Schuller *et al.* proved that infrastructure such as a high-speed network and smartphones with high-performance computational units can provide support for continuous monitoring of the psycho-emotional state for a long period [26]. Therefore, it will be beneficial to develop a low-cost screening technique that can be deployed in communities and operated by people without special training. Early-stage mental disorder screening is also crucial for policymakers and security agencies because someone with mental health disorder could behave adversely to other innocent people, such as massive shootings which are attributed to mental health disorders [27]. Recently, the development of Artificial Intelligence (AI) shows its great potential in healthcare [28]. In the cyber world, we live now, it is very common to share personal information, concerns through the Internet, especially after the rise of social media. This raises an opportunity since the contents on social media increase the likelihood to detect potential depression patients from a large population. The automated depression diagnosis has been studied from many different aspects, such as collecting and analyzing data, induction and representation of emotions, and predicting depression based on different modalities. In this paper, we propose a multimodality automated depression diagnosis system with prosodic and semantic features to predict the depression levels with the combination of Bi-LSTM and T-CNN models. To the best of our knowledge, it is the first time that time-distributed CNN is adopted to further extract the temporal information from the output of the LSTM encoder. Additionally, our proposed model does not have a strict limitation of input duration, regardless of the number of frames, as long as the number of features meets our specification, our model can always provide a patient-independent depression prediction. The prediction is based on a specific text or audio feature sequence. Given a specific

participant with an audio/text feature sequence of arbitrary length, our model provides a series of estimations of depression severity based on the audio/text feature. The set of predictions can be merged through a major voting algorithm so that the final output of our model is a patient-level depression severity prediction. This mitigates the problem that the audio/text feature sequences are required to be the same in length in previous articles. LSTM performs well in learning temporal information because of its recurrent structure. The bidirectional LSTM model is used to learn long-term bidirectional dependencies in the audio and text feature sequences because it has been proved to perform better than a unidirectional LSTM model. The convolutional neural network (CNN) is a popular network architecture to learn the spatial features of data. A time-distributed CNN architecture is obtained by having multiple CNN layers for Bi-LSTM output features at each timestep. Given the complementary advantage of CNN and LSTM, the hybrid model of LSTM and T-CNN works well in learning the spatiotemporal sequence. The best patient-independent F_1 score of the audio and text model is 0.9870 and 0.9709, respectively, on the test partition of the DAIC-WOZ dataset. The fused multimodality model achieved the best F_1 score of 0.9580 on the test partition of the DAIC-WOZ dataset.

2 RELATED WORK

2.1 Depression Severity Prediction Based on Text

Text, especially non-verbal cues which are not expressed directly in dialogue has gained remarkable popularity in depression prediction and sentiment analysis due to two reasons. Firstly, psychiatrists observe speech attributes such as less variation in speech production during an interview, which are commonly used biomarkers in depression diagnosis [24], [29], [30], [31], [32]. Secondly, the text transcription is an explicit signal to record, which makes interview transcripts one of the best candidates for distress analysis tasks [20], [33]. The prediction of depression severity is based on the hypothesis that mental disorder causes some accessible and observable differences affecting how verbal content is produced [34], [35]. Previously, a problem in searching for the relationship between depression and semantic features is the difficulty in collecting qualified and sufficient data. With the advancement of social network usage, a large amount of text data inflow gives an opportunity to researchers to analyze distress state from text [24], [30], [31], [36]. Coppersmith *et al.* seminally proposed to acquire a qualified dataset via social network platforms, which solves the problem of data insufficiency [37]. However, those colloquial languages, such as abbreviated words, popular slang, etc., make data preprocessing very difficult. Additionally, people are more likely to publish negative content on social media because they are anonymous. This indicates that although someone without any mental disorders is still likely to publish many negative posts for some periods. The effectiveness of social media posts should be fully investigated before being widely used in automated depression diagnosis because the quality of the training data affects the performance of the classifier. We also disagree that collecting data from the Internet is an effective strategy because those posts not related to depression are

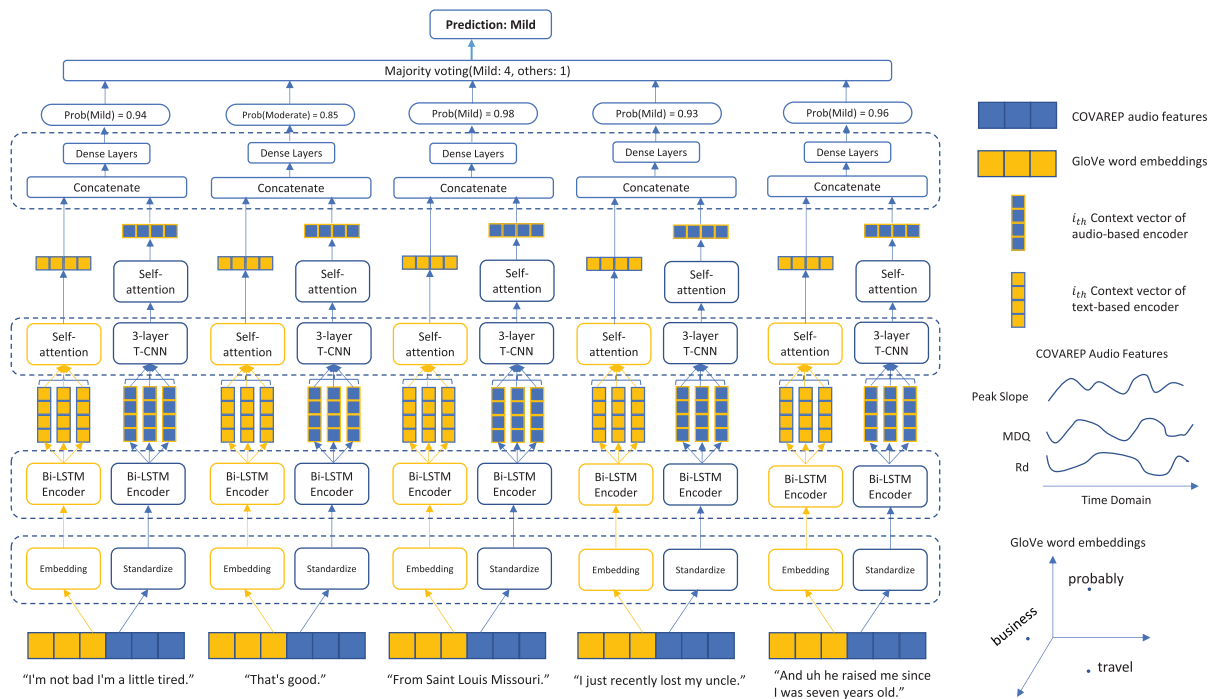


Fig. 1. Block diagram of our proposed multimodality depression level prediction algorithm given a specific example. Audio features are fed into the network through the input layer. After batch normalization, the input data is fed into the Bi-LSTM and time-distributed CNN block. In this proposed design, we have five time-distributed CNN blocks followed by a single-layer Bi-LSTM. The detailed architecture of each block is illustrated and explained in the remainder of this paper.

likely to be dominating factors compared with those depression-sensitive posts. To solve the problem of patients concealing their thoughts and emotional states, Scherer *et al.* proposed to collect the dialogue during the screening interview led by clinicians [38]. This data collection strategy has several advantages. Firstly, the questions are specifically designed by psychologists, and the efficiency of the question is better than user-generated content on social media. Before starting the interview, the participants are required to complete a questionnaire like PHQ-8 or BDI. After the interview, the clinician determines the depression severity of the participant based on the response of the patient during the interview and the questionnaire. Overall, these studies highlight the need for reliable corpora for speech-based depression prediction.

2.2 Depression Severity Prediction Based on Voice

In this section, we overview the application of prosodic and acoustic features in predicting depression. The relationship between depression and voice change has been well-studied [39], [40]. The earliest research on depressive voices can be traced back to the 1920 s. The father of modern psychiatry, Emil Kraepelin, characterized the voice to be depressive as “low voice, slowly, hesitatingly, monotonously, sometimes shuttering, whispering, try several times before they bring out a word, become mute in the middle of a sentence [41].” To train the audio model, the first thing is to extract audio features from the raw audio recordings. Feature extraction is the preprocessing technique that converts the original audio into more abstract, dense vectors. Cummins *et al.* pointed out several critical properties for a perfect feature to detect depression or other mental disorders [42]. The most

important property is that the feature should represent some recurring and noticeable effects caused by depression. The feature must also manifest large cross-label variability but small inner-label variability. Furthermore, the feature should be robust to environmental noise if it is intended to be used in the automated depression diagnosis system. Many previous works adopted a Support Vector Machine (SVM) and Gaussian Mixture Model (GMM) [43], [44]. They are two popular machine learning techniques and are robust to overfitting. Much of the available literature has attempted to use the combination of prosodic and glottal features to train the classifier [40], [45], [46], [47], [48]. As for the Mel-Cepstral Coefficients, it is reported that the low-order MFCCs perform much better than high-order MFCCs for emotion prediction or some para-linguistic analysis tasks [49]. Additionally, except for these low-level audio features, some researchers proposed to use pre-trained convolutional neural networks such as VGG-16 to extract high-level features in a frequency spectrogram [32]. However, the effectiveness of this deep frequency spectrum feature is questionable. Although the CNN model outperforms other traditional models in Computer Vision, the frequency spectrogram is different from other images. The CNN is spatial invariant because it applies a group of identical transformations to different regions of an image [50]. The frequency spectrogram consists of an X-axis denoting the frequency and the Y-axis as the intensity of the frequent component. The position of a component in the frequency spectrogram matters, but the components in ordinary images are less sensitive to the position. Regarding this concern, we do not adopt this deep frequency spectrum feature extraction method. Instead, our model utilizes the low-level audio features mentioned above. Together, these studies indicate that we should consider the

TABLE 1
The Showcase of a Participant's Transcript

Start time	Stop time	Speaker	Utterance
87.322	89.592	Ellie	So how are you doing today?
89.71	91.93	Participant	I'm not bad I'm a little tired but okay.
92.945	93.585	Ellie	That's good.
94.257	95.577	Ellie	Where are you from originally?
95.78	97.14	Participant	Uh from Saint Louis, Missouri.

combination of audio features as the input for training the depression prediction model.

3 DATASET AND METHODS

In this section, we briefly introduce the preliminary material we used for developing the audio model, text model, and multimodality model. We also discuss the dataset and framework for training and evaluating our proposed model.

3.1 Distress Analysis Interview Corpus-Wizard of Oz (DAIC-WOZ)

In this paper, we adopted the Distress Analysis Interview Corpus-Wizard-of-oz (DAIC-WOZ) dataset for training and testing [50]. The corpus consists of 189 recorded clinical interviews and transcripts as well as facial features from 189 subjects. The audio recordings were taken of semi-structured interviews between the participants and a virtual interviewer called Ellie, an animated role controlled by a human interviewer. The average audio duration of 189 subjects is 974 seconds. Subjects were solicited from the Greater Los Angeles Metropolitan region from two different populations. One was from civilians; the other was from veterans of the U.S armed forces. Subjects were characterized as depression, Post-Traumatic Stress Disorder (PTSD), and anxiety based on the self-report questionnaire during the data collection [50]. Only the interview recordings of the depression group were released for academic purposes. The gender distribution over all five groups as well as the dataset partition is shown in Table 3. In the training set, there are 44 female subjects (27 without significant depression symptoms, 17 with depression symptoms) and 63 male subjects (49 without significant depression symptoms, 14 with depression symptoms). In the validation set, there are 19 female subjects (12 without significant depression symptoms, 7 with depression symptoms) and 16 male subjects (11 without significant depression symptoms, 5 with depression symptoms). In the test set, there are 24 female subjects (17 without significant depression symptoms, 7 with depression symptoms) and 23 male subjects (16 without significant depression symptoms, 7 with depression symptoms). All interviews were transcribed verbatim into English. The interviews lasted from 5 to 20 minutes involving three phases: it started with neutral questions, which aimed to ensure subjects being able to calm down; the interview then proceeded into a targeted phase, and the questions asked by the interviewer were more related to the symptoms of depression and PTSD. Finally, the interview terminated with the annealing phase, which assisted the participants to get rid of the distressed state. The PHQ-8, ranging from 0 to 24, determines the severity of the

mental disorder. Subjects were divided into five groups: healthy ($\text{PHQ-8} < 5$), mild ($5 < \text{PHQ-8} < 10$), moderate ($10 < \text{PHQ-8} < 15$), moderately severe ($15 < \text{PHQ-8} < 20$), and severe ($\text{PHQ-8} > 20$) [51]. Table 1 shows a sample transcript in the DAIC-WOZ dataset, which contains four fields: beginning and end timestamp of the utterance, the speaker ID, and sentence content. Due to space limitation, Fig. 2 below illustrates the distribution of the first four audio features provided with the DAIC-WOZ dataset with the significant intra-subject variance. In the remaining part of this paper, the training, validation and test set are split by the instruction from the DAIC-WOZ dataset independently, which ensures all the subjects only appear in one of the above partitions.

3.2 Recurrent Neural Network and Long-Short Term Memory

A recurrent neural network (RNN) is a deep learning architecture that outputs a time sequence. The input of the neural network is transformed into hidden states at different time steps. Given an input vector x_t , the intermediate variables in the network are computed iteratively, from (h_1, z_1) , (h_2, z_2) to (h_t, z_t) , where h_t and z_t are the hidden state and output of the RNN cell, respectively. The traditional RNN performs well on some machine learning tasks, such as voice recognition [32]. However, the gradient vanishing/exploding problem during the backpropagation limits the depth of the RNN. To solve this problem, Hochreiter *et al.* proposed the LSTM, which stands for Long-Short Term Memory [52].

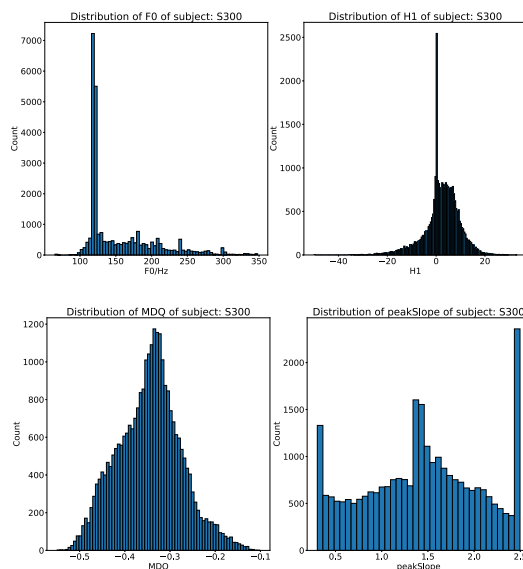


Fig. 2. The distribution of a subset of audio features within the dataset.

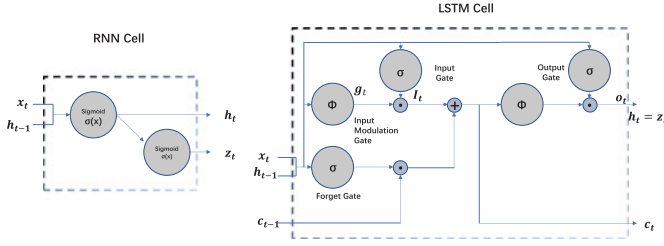


Fig. 3. A schematic of an RNN cell (left) and an LSTM cell (right) used in this paper.

LSTM can determine when to “forget” some previous information and update the hidden state during the training phase by combining different “gates” in the LSTM cell. The traditional RNN and the LSTM cell are illustrated in Fig. 3. Compared with the traditional RNN cell, the LSTM cell includes some special components such as the input gate, forget gate, output gate, input modulation gate, and the memory cell. The i_t and f_t determine whether the previous information should be memorized or forgotten. Similarly, the output gate determines how much information in the cell memory can be transferred to the hidden state. These gates enhance the performance of LSTM on time series-related tasks and make it possible to train a deeper network. The hidden state of the previous layer can be fed into the following layers to construct a deeper network, which improves the capability of LSTM to deal with more complicated time series. From the probabilistic perspective, automated depression diagnosis is to find a correct severity sequence y that maximizes the conditional probability of y given an input feature sequence (i.e., audio/text feature). Our proposed framework, based on an RNN encoder-decoder, learns to predict depression severity given a sequence of audio and text features. In the encoder neural network, an encoder reads and projects the input feature sequence $X = (x_1, x_2, \dots, x_T)$ into a context vector c , which is given by:

$$c = q(h_1, \dots, h_T) \quad (1)$$

$$h_t = f(x_t, h_{t-1}) \quad (2)$$

where h_t is a hidden state at time t , c is a vector computed from a sequence of hidden states. f and q are nonlinear functions. The decoder neural network is trained to predict the depression severity given the context vector and the input feature at time t . The probability of depression severity is given by:

$$p(y) = \prod_{t=1}^T p(y_t | \{x_1, \dots, x_{t-1}\}, c) \quad (3)$$

where $y = (y_1, \dots, y_T)$, and each term of the conditional probability is given by:

$$p(y_t | \{x_1, \dots, x_{t-1}\}, c) = g(x_{t-1}, s_t, c) \quad (4)$$

where g is a nonlinear function, s_t is the hidden state of the RNN.

3.3 Attention Mechanism

Most of the proposed deep learning-based depression prediction models are a member of a family of encoder-

decoders, with an encoder for high-level representation of original input audio or text features. The encoder network reads and encodes the variable-length input audio/text features into a fixed-length vector. A decoder then decodes the fixed-length vector and outputs a probability matrix from the encoded fixed-length vector. The cascade model, which consists of an encoder and decoder, is optimized simultaneously to maximize the probability of a correct depression severity given an original audio/text feature sequence.

A shortcoming of this encoder-decoder architecture is that the encoder network has to compress all the depression-sensitive information into a fixed-length vector. In the scenario of extremely long input sequences, this may make it challenging for the encoder network to encode necessary information into the fixed-length vector, especially during testing, when the length of the input sequence for testing is longer than the length of sequence for training. To overcome this drawback, we adopted an attention mechanism that allows the model to select a subset of encoded vectors adaptively while decoding the high-level representation. Each time the decoder makes an inference on depression severity, it goes through the encoded input sequence and works out the most depression-sensitive information. The most important feature of the attention mechanism is that it does not rely on a single fixed-length vector. The model can select a subset of encoded high-level representation adaptively during training, which frees the encoder network from compressing all necessary depression-related information, no matter how long the original sequence is, into a fixed-length vector. This improves the performance of our model, especially the performance coping with long sequences.

With the attention mechanism, we can compute the weighted context vector with RNN output hidden states. The depression conditional probability of time step t is given by:

$$p(y_t | \{x_1, \dots, x_{t-1}\}, c) = g(x_{t-1}, s_t, c_t) \quad (5)$$

Where s_t is the RNN hidden state for time t , which is given by:

$$s_t = f(s_{t-1}, x_{t-1}, c_t) \quad (6)$$

Unlike the traditional encoder-decoder framework, the depression conditional probability is not only conditioned on a uniform context vector c but a distinct vector c_t for each timestep. The context vector is given by a sequence of RNN hidden states, which are the output of the encoder neural network. A hidden state at time step t contains all information about the input feature sequence prior to time step t , with an emphasis on the part around the entries at time step t . The context vector is given by:

$$c_t = \sum_{j=1}^T \alpha_{ij} h_j \quad (7)$$

The coefficients α_{ij} for each hidden state is determined by:

$$\alpha_{ij} = \frac{\exp(e_{ij})}{\sum_{k=1}^T \exp(e_{ik})} \quad (8)$$

where e_{ij} is given by:

$$e_{ij} = a(s_{i-1}, h_j) \quad (9)$$

$a(x)$ is a score function that evaluates how well the inputs around the entries at time j and the output of RNN at time $(i - 1)$ match. The score function $a(x)$ is a distinct layer that is simultaneously trained with all other layers of the proposed model. The probability α_{ij} describes the importance of the hidden state h_j regarding the previous hidden state s_{i-1} during calculation of s_i . This allows the decoder itself to determine which part of the input sequence should be focused on. With the attention mechanism, we alleviate the burden of compressing the input sequence, regardless of its original length, to a fixed-length vector. Therefore, with the attention mechanism, the correlation in the context vector can be propagated through the network, which allows the decoder to selectively retrieve those depression-related hidden states.

3.4 Audio Features and Models

In this paper, the audio features are extracted by COVAREP [53], which can be divided into three categories: glottal flow features (NAQ, QQQ, H1-H2, PSP, MDQ, Peak slope, Rd), voice quality features (F_0 , VUV), and spectral features (MCEP, HMPDM, HMPDD). Normalized Amplitude Quotient (NAQ) quantifies the time-based feature of the speaker by amplitude-domain measurements calculated from the glottal flow and its first derivative [54], [55], Quasi Open Quotient (QQQ), which is a correlate of the open quotient (OQ) which involves the derivation of the quasi-open phase based on the amplitude of the glottal phase [56], [57], the amplitude difference of the first two harmonics of the differentiated glottal source spectrum (H1H2) [58], Parabolic Spectral Parameter (PSP), which is based on the quantification of the spectral decay of the speaker [58], and Maxima Dispersion Quotient (MDQ), which is designed to quantify the maxima dispersion as a result of phonation type moves towards a breathier phonation [59], [60]. Spectral features consist of Mel-Cepstral Coefficients (MCEP0-24), which is a representation of the short-term power spectrum of a sound [61], harmonic model and phase distortion mean (HMPDM0-24) and deviation (HMPDD0-12). Thus, there are 74 audio features in total. Each subject is represented in the COVAREP features, $X_i \in R^{T \times F}$ where T denotes the time dimension, which is proportional to the duration of the audio. Each 10 milliseconds frame of audio was transformed into an audio feature vector. F denotes the number of features COVAREP extracted for each frame. Among the 74 audio features, the entry "VUV" indicates whether the audio features are extracted from the audible or silent part of the original interview recording. Only those audio features where "VUV" is 1 can be the input to the following models. Among all the 189 subjects in the dataset, audio features are in an average of 35850 frames (rows) and a standard deviation of 15791 frames (rows). For each subject, we concatenated a constant number of audio feature frames into a set of successively retrieved audio feature sequences, which were used to represent this subject. The shape of the input tensor is thus (#samples, #frames, 73). The field "VUV" is always 1 in the input tensor so it is dropped, which results in the final input tensor shape as 73.

Audio models with different configurations for depression assessments are introduced as follows. The input to

these models is the previously mentioned audio feature sequences, the output of these models is the prediction of the depression severity given an audio feature sequence. The first audio model is a simple one that consists of the LSTM and fully connected layers. The LSTM served as a feature extractor and the following fully connected layers made the prediction based on the output of the LSTM. Then, we introduce our proposed model that consisted of the Bi-LSTM and T-CNN and they were evaluated for the prediction of depression severity.

3.4.1 Traditional LSTM-Based Model

Our first audio model comprises of single-layer Long-Short Term Memory (LSTM) network and fully connected layers. LSTM network was obtained using an LSTM layer containing 73 hidden units, connected to a fully connected layer. To avoid overfitting, the dropout was applied to the recurrent input signal on the LSTM units and between fully-connected layers with the dropout rate of 0.2. The time step is equal to the constant "#frames" and there were 73 features in each timestep. In this model, only the hidden state at the last time step was fed into the following fully connected layers, with 128 and 64 hidden units. The output of the fully connected layer was then fed into a batch normalization layer and flattened into a 1D tensor. The flattened tensor was fed into a fully connected layer with 5 hidden units, where the SoftMax activation function transformed the unnormalized output of each neuron into the probabilities of five severities. An Adam optimizer was adopted for the training, the initial learning rate was set to be 0.001, $\beta_1 = 0.9$, $\beta_2 = 0.999$ and the epsilon was 10^{-7} . A callback function monitored the validation loss and terminated the training if the validation loss did not decrease after five epochs. A loss function of cross-entropy was applied.

3.4.2 Hybrid of Bi-LSTM and T-CNN Model

Bidirectional LSTM is a variant of LSTM which consists of a forward layer on the original input sequence, and a backward layer on the reversed sequence. The Bi-LSTM outperforms the traditional LSTM because the forward and backward networks combine both forward and backward context information of the input sequence. Previous articles proposed to represent the input sequence by the last hidden state of the LSTM [61], [62]. However, depression assessment is a complicated task, which heavily relies on the relationship between the audio features at different time steps, thus it is insufficient to use the last hidden state for classification, otherwise, it leads to the loss of temporal information. To solve this issue, we utilized the T-CNN to learn potential temporal and spatial information in the output of the Bi-LSTM. The structure of T-CNN is illustrated in Fig. 5. In general, simple CNN only supports the 2D or 3D spatial tensors as the input. However, the output shape of the LSTM is (#samples, #frames, #LSTM neurons) given a unidirectional LSTM, and (#samples, #frames, 2*#LSTM neurons) given a bidirectional LSTM. The T-CNN convolves the LSTM output vector along its 3rd axis and the shape of the convolution result is (#samples, #frames, #output features, #kernels). Therefore, we expand the shape of the LSTM output vector by inserting one new axis so that it can be

processed by T-CNN. The T-CNN accepts a tensor with shape (#samples, #frames, $2 \times \text{\#LSTM neurons}$, 1) as the input, which denotes a time series of LSTM hidden states. Our proposed T-CNN block consisted of three layers, first, time-distributed convolution layer, then time-distributed pooling layer to downsample the feature maps; and finally batch normalization layer. There were five T-CNN blocks in total in our proposed design, the output of the last T-CNN block contained “#frame” samples, each sample is represented by 256 feature maps. Therefore, the last thing before the feature maps were fed into the following network was to downsample the output by the global average pooling layer, it slides along the time dimension of the feature vector and computes the mean value of each feature, which ensures that the relationship between each time step was taken into consideration. The output of the global average pooling layer was then fed into the following two linear layers. At last, the Softmax activation functions transformed neuron output into the probability of five severities. An Adam optimizer with a similar configuration in 3.4.1 was adopted for the training.

3.5 Text Features and Models

The input layer of the text model took tokenized transcripts of each subject. Among all the 189 subjects in the dataset, text transcripts are in an average of 80 rows and a standard deviation of 14 rows. The interviews were in colloquial speech, thus the first step was to rephrase these colloquial descriptions to written languages, otherwise, colloquial terms all became out-of-vocabulary words, which were represented by the token [UNK], and greatly diminished model performance. Semantic information is highly essential in depression diagnosis because psychologists also formulate diagnosis by text produced by the patients during the interview. To acquire the text features, we firstly removed stop words in the patients’ responses with Natural Language Toolkit (NLTK) and substituted some words and phrases such as “what’s,” “e-mail” with “what is” and “email,” this eliminates different expressions of the same word [63]. Next, we lemmatized the remaining words in the sentences, the WordNet lemmatizer removes the inflectional endings and returns the base form of a word. Then the remaining texts were tokenized into word lists and were used to build a vocabulary with 7373 words. Each word in the vocabulary was assigned an index, the word list was then represented by these indices. After we acquired the word list, the main issue was that each word list was different in length, which made it more difficult to batch process text data if they were different in length. Therefore, the sliding window technique was applied to generate sequences in the same length, which was the same length as the sliding window. Each window consists of a constant number of words while 20% words at the end were overlapping between two neighbouring time windows, which assigned higher weights to the words at the edge of the window so that the edge details were enhanced. The sliding window not only generated all training pairs but also performed data augmentation as well as directed the focus on a specific part of the sentence. Next, word sequences were encoded with the pre-trained 100D GloVe word embedding vector [64]. The word

embeddings were concatenated into a sentence embedding. For some short sentences, the size of the sliding window was greater than the length of the sentence, those short sentences were zero-padded to be the same length as the window. Therefore, the shape of the final input vector is (window size, 100). However, sentences shorter than 20% of window size were discarded.

3.5.1 Bi-LSTM Text Model

Our proposed text model consists of a single-layer Bi-LSTM network and fully connected layers. The text feature sequences mentioned above comprise the index of words in the vocabulary. Text feature sequences were preprocessed to map each word to word embedding space with a non-trainable embedding layer before being fed into the model, and the shape of the embedding layer is (vocabulary size + 1, 100). Next, a batch normalization layer and then the Bi-LSTM layer further captured the semantic information underlying the input word sequences. To avoid overfitting, the dropout was applied to the recurrent input signal on the LSTM units and between fully-connected layers with the dropout rate of 0.2, and the shape of the Bi-LSTM output was (batch size, 200) at each time step. We adopted the attention mechanism to allow the model to adaptively select those depression-sensitive hidden states. The attention vector was then fed into two linear layers with 256 and 128 hidden units, respectively. Finally, the last linear layer with 5 hidden units determined the probability of the five severities. An Adam optimizer with a similar configuration in 3.4.1 was adopted for the training. The cross-entropy loss calculated the distance between the output and the ground-truth label.

3.6 Fused Text-Audio Joint Model

Our final fused multimodality model was comprised of two subnetworks: text model and audio model, and followed by a shared late fusion neural network as Fig. 1 shows. The late fusion neural network concatenated the outputs of the text and audio model to integrate text and audio features. For any subject, we extracted a high-level representation that included both semantic and prosodic features through the previous recurrent neural network and convolutional neural network. This high-level representation could be used in the following assessment of mental disorders. The output of our proposed model was a scoring matrix that denoted the likelihood of the depression severity. As the timesteps of the audio and text model were different, the late fusion network had to deal with input of different sizes. To solve this issue, we first attempted to adopt a max-pooling method to downsample the output from audio and text models so that they were in the same shape. Moreover, an attention mechanism was exploited, which provided us insights into the ratio of the contribution of each modality towards the final prediction.

Regarding fusion, we designed a set of models to integrate different modalities. Firstly, we fused the text models with different window sizes with the audio model with constant configuration. Our text model could be divided into two categories, one is the unidirectional LSTM text model, the other is the bidirectional LSTM. Our proposed audio

TABLE 2
Result of the T-Test of the Control and Experiment Group

	Sentence length		Audio duration	
	Control	Experiment	Control	Experiment
p-value	8.7854±8.9475	7.3717±7.2975	951.3711±266.6010	997.8773±290.1901
	3.2397×10^{-14}		0.0952	

and text model was previously described in Section 3.5 and Section 3.4, respectively. The only difference was that the output size of the audio and text model was 32 instead of 5 since they acted as feature extractors rather than classifiers. Global max pooling was adopted to align the extracted audio and text features. In order to integrate text and audio modalities, the output of the text and audio model was concatenated into a tensor and passed through a fully connected layer with 5 units. Secondly, the other fused model was set up using a similar configuration to the first one. The difference was that the attention mechanism played its role in aligning the features from different modalities. The third one was all the same as the previous two models, except it was created with an attention mechanism not only during the feature alignment but also in the fusion of the high-level representations.

4 EXPERIMENTS AND RESULTS

In this section, the results of those models described in Section 3 are presented and discussed. We next assessed the effect of the hyperparameters for the proposed models. For the audio model, we compared the effect of architecture and timestep and investigated the potential long-term dependency of the audio features in severe patients. For the text model, we conducted experiments to investigate the effect of the hyperparameters such as the size of the window in preprocessing, the removal of stop words. Regarding the audio-text fused model, we mainly focused on the impact of fusion methods on the model performance. All the experiments were conducted on one RTX 2080Ti 11 GB GPU. The size of multimodality models was limited mainly by the amount of memory available on our GPU and the amount of time for training we can tolerate. Our single-modality model usually took between 3 to 5 hours to train, but the training of our proposed multimodality model always took around 20 hours. The results of our experiment provided an insight that our models could be improved by faster GPUs and larger datasets. The detailed results are discussed in the following parts.

4.1 The Statistics of Audio and Text Features

The pause time between responses is also longer than usual in the depressive population [41]. To verify whether the DAIC-WOZ dataset follows a similar pattern, we calculated the statistics of the raw interview recordings and the transcripts. The subjects were divided into two groups by PHQ-8 scale, the subjects were considered as normal or mild (control group) if their PHQ-8 is less or equal to 10, otherwise, they are considered as moderate or severe (experiment group). This threshold is given by a previous study on the efficacy of PHQ-8 on the diagnosis of major depressive disorder. It was reported that given the cutoff score of 10, the PHQ-8 exhibited a sensitivity of 58.3%, specificity of 83.1%

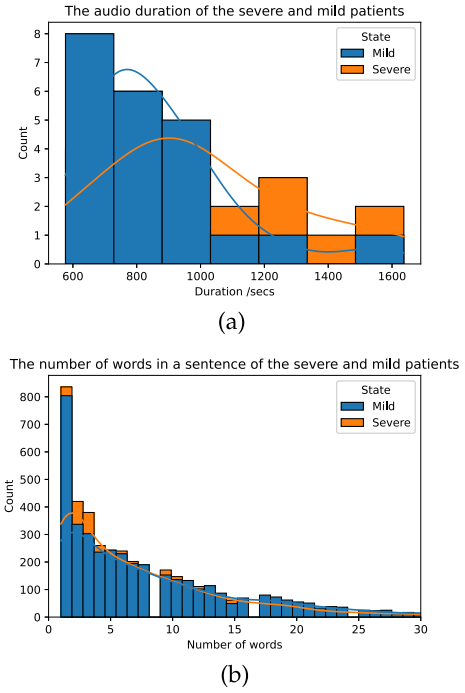


Fig. 4. The histograms of the audio duration and sentence length of control and experiment groups. (a) The audio duration of the control and experiment groups. (b) The sentence length of the control and experiment groups.

[65]. The two-sided T-test was applied to test if there was a significant difference in the audio duration between the control and experiment groups. The statistics of the two groups are listed in the Table 2. The histograms of the audio duration and sentence length of the control and experiment groups are illustrated in Fig. 4. The response duration of the control and experiment groups is on an average of 951.3711 ± 266.6010 and 997.8773 ± 290.1901 seconds, respectively. The two-tailed p-value is 0.0952. The sentence length of the control and experiment groups is on average of 8.7854 ± 8.9475 and 7.3717 ± 7.2975 in the number of words, respectively. The two-sided T-test was applied to test if there was a significant difference between the sentence length in the control and experiment groups. The two-tailed p-value is 3.2397×10^{-14} . The above results indicate no significant difference in the audio duration of the control and experiment groups. However, the sentence lengths for the control and experiment groups are significantly different. More responses in the experiment group consisted of less than 5 words. As the audio durations between the control and experiment groups have identical average values, we can conclude that there are more pauses in the conversations of the experiment group. This result is identical to other researchers' conclusions. Therefore, our dataset and criterion for depression are reasonable.

4.2 Results of the Audio Modality

As for the audio models, evaluation metrics accuracy, recall, precision, and f1 score used to evaluate models with different configurations are shown in Tables 4 and 5. The test set for evaluation is balanced by oversampling the minority class. Random forest was used as the baseline in evaluating the audio modality sequence-level prediction. Audio

TABLE 3
Gender Distribution Over All Groups and Dataset Partitions

		Training set				Validation set				Test set			
ID		303, 304, 305, 310, 312, 313, 315, 316, 317, 318, 319, 320, 321, 322, 324, 325, 326, 327, 328, 330, 333, 336, 338, 339, 340, 341, 343, 344, 345, 347, 348, 350, 351, 352, 353, 355, 356, 357, 358, 360, 362, 363, 364, 366, 368, 369, 370, 371, 372, 374, 375, 376, 379, 380, 383, 385, 386, 391, 392, 393, 397, 400, 401, 402, 409, 412, 414, 415, 416, 419, 423, 425, 426, 427, 428, 429, 430, 433, 434, 437, 441, 443, 444, 445, 446, 447, 448, 449, 454, 455, 456, 457, 459, 463, 464, 468, 471, 473, 474, 475, 478, 479, 485, 486, 487, 488, 491				302, 307, 331, 335, 346, 367, 377, 381, 382, 388, 389, 390, 395, 403, 404, 406, 413, 417, 418, 420, 422, 436, 439, 440, 451, 458, 472, 476, 477, 482, 483, 484, 489, 490, 492				300, 301, 306, 308, 309, 311, 314, 323, 329, 332, 334, 337, 349, 354, 359, 361, 365, 373, 378, 384, 387, 396, 399, 405, 407, 408, 410, 411, 421, 424, 431, 432, 435, 438, 442, 450, 452, 453, 461, 462, 465, 466, 467, 469, 470, 480, 481			
	Dataset profile for depression level classification												
	Female	Male	Female	Male	Female	Male	Female	Male					
#Healthy	7	9	2	3	3	6	2	2					
#Mild	12	25	7	6	9	9	11	11					
#Moderate	10	20	5	2	7	7	4	3					
#Moderately severe	10	5	2	2	1	4	4	4					
#Severe	5	4	3	3	4	3	3	3					
Dataset profile for depression detection													
#Subjects w/o significant symptom (PHQ-8≤10)	27	49	12	11	17		16						
#Subjects w/ significant symptom (PHQ-8>10)	17	14	7	5	7		7						

feature sequences for training and evaluating are non-stationary series, which are difficult to model and forecast. They were pre-processed by differencing to be made stationary. Differencing is the change from one audio feature sampling time to the next. The random forest model we used in this manuscript is an ensemble approach that fits a set of decision trees on different sub-sample of the dataset, and averaging the output of each decision tree to improve the prediction accuracy, as well as prevent the model from overfitting. In our article, 100 decision trees were trained on various sub-sample of the training set to construct the random forest model. Another baseline method, Madhavi *et al.* proposed a CNN consisting of 2 convolutional layers and two successive linear layers to extract high-level features from the frequency spectrogram of interview recordings. The output of CNN is fed into the following neural networks to predict an individual's depression level. They also evaluated their models on the DAIC-WOZ dataset. Moreover, Yang *et al.* proposed a similar but more complex model, they also adopted the combination of convolution neural networks and deep neural networks (i.e. multi-layer

perceptron model). Each subject was labelled by their depression-related symptoms, such as prior depression diagnosis, sleep disorder, present or not. Their proposed CNN consists of three convolution layers and the intermediate output of CNN is fed into the deep neural network to predict the presence of depression symptoms. These symptom labels are fed into another deep neural network for predicting depression severity. Their results on the DAIC-WOZ dataset are summarized in our comparative studies. For the LSTM with the fully connected layers model, it outperformed the baseline model machine learning model (i.e. decision tree) by 24% in terms of accuracy. In contrast, the Bi-LSTM with the fully connected layers model outperformed by 54% in terms of accuracy. For our proposed Bi-LSTM combined with the T-CNN model, we achieved 16% improvements over the best baseline model in terms of accuracy. From Tables 4 and 5, it can be concluded that the LSTM performed better on the depression level classification compared with the baseline machine learning models, such as the naïve Bayes model. Moreover, we observed that the network followed by the LSTM layer is critical for good performance. If the other configurations were fixed, Bi-LSTM with T-CNN outperformed other methods because the T-CNN learned more temporal and spatial information than others by capturing the correlation within all hidden states of the LSTM. We also investigated the influence of the value of the time step and concluded that our model performed best when the timestep was 16. Fig. 6(a) shows the receiver operating characteristic (ROC) curve when timestep=16. The micro-average AUC for our proposed model is 0.98, and the AUC for "mild" is smaller than any other, which indicates it is more challenging for the model to distinguish mild depression from the other levels correctly. This is likely to be attributed to the fact that mild patients behave very similarly to healthy people during the interview. Fig. 6(b) is the ROC when the time step is 32. The

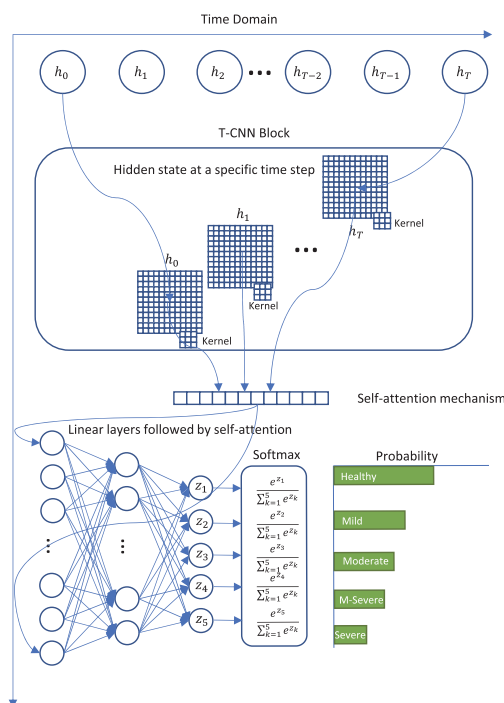


Fig. 5. The structure of the T-CNN model and the following linear neural network.

TABLE 4
Results of the Baseline Audio Models

Models	Test Baseline			
	Random Forest		Madhavi <i>et al.</i> [66]. Yang <i>et al.</i> [67]	
	Mean	St. dev	Mean	St. dev
Accuracy	0.3192	0.0085	0.7500	0.8273
Precision	0.3206	0.0064	0.7200	0.7930
Recall	0.3184	0.0040	0.7500	1.0000
F1 Score	0.3168	0.0076	0.7300	0.8850

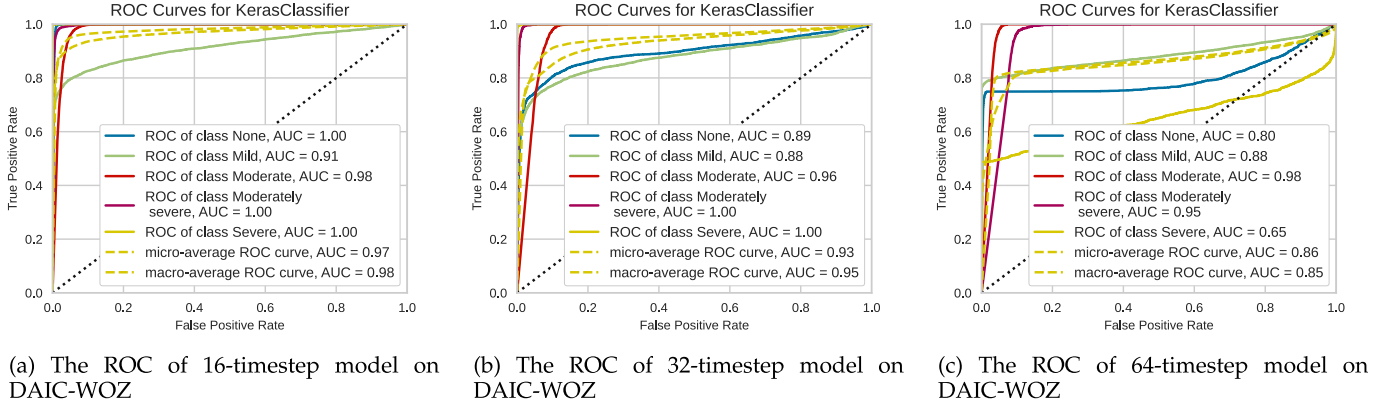


Fig. 6. The ROC of three different model configurations. (6(a)) The Bi-LSTM followed by T-CNN given the time step = 16. Micro-Average AUC: 0.97. The AUC of “Severe” is smaller than any other class, this indicates the detection of severe depression is more challenging than other depression levels. (6(b)) The Bi-LSTM followed by T-CNN given the time step = 32. Micro-Average AUC is 0.93. The micro-average AUC is smaller compared with that when the time step = 16. The longer sequence does not mean a better result because the noise introduced by the longer sequence can mislead the model. (6(c)) The Bi-LSTM followed by T-CNN given the time step = 32. Micro-Average AUC is 0.86, which is in line with our expectation that a longer input sequence makes it more challenging to predict the severity.

micro-average AUC for this model is 0.93. The performance of the model with 32-timesteps was worse than that of the model with 16-timesteps. This is likely due to the negative correlation between the signal-noise ratio of the input sequence and the length of the sequence. A longer input sequence contains more information to assess the emotional state, but as the sequence length grows, the increasing noise cannot be ignored and the bias of the model rises due to the noise. Another factor is the limitation of the memorization capability of LSTM. The longer the input sequence is, the more difficult it is for LSTM to memorize earlier information when processing the end of the sequence because the depth of the LSTM network is proportional to the timestep. Given a long sequence, the information cannot smoothly flow through the network, which results in diminished performance. The confusion matrix of the 32-time step model is illustrated in Fig. 7(b), which shows the performance of the model on the test partition of the DAIC-WOZ dataset. Comparing the models with different time steps, Fig. 7(a) shows the confusion matrix of the model with 16 timesteps, while Fig. 7(c) shows the confusion matrices of the model with 64 timesteps. Different timestep means the different sizes of the test set. To eliminate the influence of the size of the test set, we normalized the confusion matrix along each row. In terms of the normalized confusion matrix, the model with

16 timesteps performed the best, but from the entries on the second row of Fig. 7(c), the model with 64 timesteps was less likely to classify the mild patients incorrectly. The contribution of the model with a longer time step in the depression prediction should be further investigated to find the cut-off value of the time step that optimizes the trade-off between the computation cost (larger time step means more computation) and the misdiagnosed rate.

4.3 Results of the Text Modality

4.3.1 The Effect of Stop Words and Bidirectional Layer

In this experiment, we used NLTK to remove the stop words in English transcripts. Apart from the stop words, the other factor is the choice between LSTM and Bi-LSTM models. Compared with the unidirectional LSTM model, the bidirectional model converges faster, and the validation accuracy is higher. The following experiment demonstrates several advantages of the Bi-LSTM model over the traditional LSTM model on the depression level classification task. Four models were trained with the different configurations presented in Table 6. The test set for evaluation is balanced by oversampling the minority class.

From Table 6, we concluded that if the type of the LSTM was fixed (i.e., the two text models both consist of LSTM or

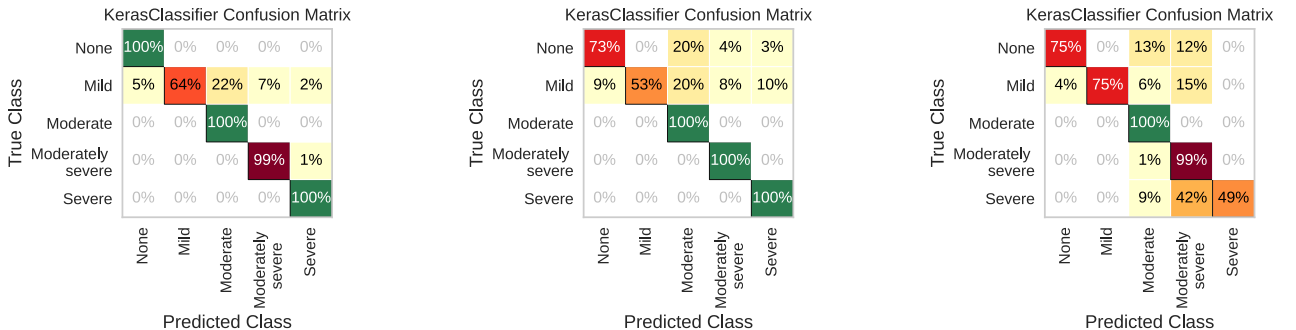


Fig. 7. Results of the proposed audio model on the DAIC-WOZ dataset.

TABLE 5
A Comparative Study of Different Proposed Audio Models

Models	Experimental Settings	Accuracy	F1	p-value
LSTM + FC	TS=16, HU=73, LHU=(128,64,5), Adam	0.5674 \pm 0.0034	0.5650 \pm 0.0042	< .01
Bi-LSTM + FC	TS=16, HU=73, LHU=(128,64,5), Adam	0.8717 \pm 0.0013	0.8818 \pm 0.0013	< .01
LSTM + T-CNN	TS=16, HU=73, #TCNNB=5, #KRNL=(64,64,64,128,256), KS=(3,3,3,3,9), Adam	0.8698 \pm 0.0897	0.8609 \pm 0.0988	< .01
Bi-LSTM + T-CNN	TS=16, HU=73, #TCNNB=5, #KRNL=(64,64,64,128,256), KS=(3,3,3,3,9), Adam	0.9871 \pm 0.0009	0.9870 \pm 0.0009	< .01

#TCNNB: Number of T-CNN blocks #KRNL: Number of conv kernels in each T-CNN block #KS: Kernel size

TABLE 6
A Comparative Study of the Proposed Text Models

Models	Experimental Settings	Accuracy	F1	Micro-average AUC
LSTM + FC	TS=64, HU=100, LHU=(256,128,5), Adam, Stopwords	0.9091	0.9094	0.9738
LSTM + FC	TS=64, HU=100, LHU=(256,128,5), Adam, No stopwords	0.9792	0.9754	0.9897
Bi-LSTM + FC	TS=64, HU=100, LHU=(256,128,5), Adam, Stopwords	0.9617	0.9610	0.9908
Bi-LSTM + FC	TS=64, HU=100, LHU=(256,128,5), Adam, No stopwords	0.9685	0.9709	0.9925

TS: Timestep; HU: #Hidden units in LSTM; LHU: #Hidden units in linear layers

Bi-LSTM network), the performance of the model without stop words was better. If the stop words were kept, the Bi-LSTM model still outperformed the traditional one. This result was in line with our expectation that Bi-LSTM was better in text classification because it learned more contextual information with the combination of the forward and backward networks.

4.3.2 The Effect of Window Size

Window size is another factor that influences the performance of the model. Intuitively, the longer the window, the more information it contains about the mental state of the subjects, which means our model can assess their emotions more accurately. However, if the window is too long, while making an inference, the impact of the noise cannot be ignored, which leads to significant performance degradation. Moreover, the memorization capability of LSTM is limited, which means the longer the sequence is, the more challenging for the LSTM to memorize and extract useful information. To demonstrate the relationship between the performance and the window size, we conducted experiments by changing the window size. As shown in Table 7, when the window size started to increase, the metrics increased firstly but began to decrease after the window size is greater than 64. This was in line with our expectation, the classifier gained a lot of information due to a larger window but started to degrade as the result of the noise in the large window and the reduced performance of LSTM. We concluded that the window size should be appropriately set to train the model with the best performance, in our experiment, the best window size is 64.

4.4 Results of the Fused Model

In this experiment, the audio and text models were jointly optimized so that we could verify whether our methods were still effective under multimodality configuration. We proposed three varieties of fusion models and merged these segment-wise predictions through major voting to obtain

the patient-level prediction. The configuration details of those fused models were described in Section 3.6. The metrics of each fusion model on the test partition were covered in Table 8. When experimenting with models made up of unidirectional LSTM, without an attention mechanism, the model with a window size of 32 performed better than others when classifying for a multi-class outcome in terms of the accuracy on the test set ($accuracy = 0.9209$). Theoretically, the models with Bi-LSTM should be better than a uni-LSTM one, however, with all other configurations fixed, except the Bi-LSTM model with a window size of 16, other Bi-LSTM models did not show significant improvement over the uni-LSTM one. Nevertheless, once the attention mechanism was introduced, the performance was boosted and the F_1 increased compared to the model without an attention mechanism, except the Bi-LSTM model with an attention mechanism and window size of 16. As we reported in the methodology section, the attention mechanism could be introduced during the multimodal feature aligning phase as well as the multimodality fusion phase. The attention mechanism during the fusion process weighed each modality and made it possible for the model to determine the contribution of each modality. From Table 8, we concluded that the highest sensitivity of 0.9941 was achieved by the model comprised of Bi-LSTM and two attention layers, with a window size of 32. Given that we expected to train an early-stage depression screening tool, we preferred higher sensitivity so that we would not miss

TABLE 7
A Comparative Study of the Text Model With Different Window Size

Window Size	Accuracy	Precision	Recall	F1 Score
16	0.8254	0.8318	0.8340	0.8141
32	0.8256	0.8371	0.8465	0.8260
64	0.8778	0.8779	0.8782	0.8705
128	0.8409	0.8599	0.8430	0.8304

TABLE 8
A Comparative Study of Our Proposed Patient-Level Methods and the State of the Art

Model	Experimental Settings	Accuracy	F1	Sensitivity	Specificity
UniLSTM as encoder	WIN=16, Stride=64	0.8604	0.8579	0.9844	0.8182
	WIN=32, Stride=64	0.9209	0.9188	0.9647	0.9777
	WIN=64, Stride=64	0.8674	0.8682	0.9705	0.9888
BiLSTM as encoder	WIN=16, Stride=64	0.9488	0.9500	0.9735	0.9444
	WIN=32, Stride=64	0.9186	0.9191	0.9852	0.9700
	WIN=64, Stride=64	0.8535	0.8546	0.9647	0.8778
BiLSTM as encoder	WIN=16, Stride=64, attention	0.8419	0.8427	0.9735	0.8222
	WIN=32, Stride=64, attention	0.9581	0.9580	0.9824	1.0000
	WIN=64, Stride=64, attention	0.9093	0.9086	0.9706	0.9889
UniLSTM as encoder	WIN=16, Stride=64, attention(aligning&fusion)	0.8977	0.8973	0.9559	0.9889
	WIN=32, Stride=64, attention(aligning&fusion)	0.9326	0.9315	0.9735	0.9889
	WIN=64, Stride=64, attention(aligning&fusion)	0.8581	0.8615	0.9412	0.8889
BiLSTM as encoder	WIN=16, Stride=64, attention(aligning&fusion)	0.8491	0.8439	0.9353	0.9000
	WIN=32, Stride=64, attention(aligning&fusion)	0.9047	0.9103	0.9941	0.9000
	WIN=64, Stride=64, attention(aligning&fusion)	0.6279	0.6560	0.7500	1.0000
Unimodality text model	WIN=16	*	BLSTM: 0.7929 ULSTM:0.8096	*	*
	WIN=32	*	BLSTM: 0.7964 ULSTM:0.7619	*	*
	WIN=64	*	BLSTM: 0.9245 ULSTM:0.9058	*	*
	WIN=128	*	BLSTM: 0.8266 ULSTM:0.7148	*	*
Unimodality audio model	BLSTM + FC	*	0.8819	*	*
	ULSTM + FC	*	0.7604	*	*
	BLSTM + TCNN	*	0.9074	*	*
	ULSTM + TCNN	*	0.8443	*	*
Srimadhur et al. [68]	End to end convolutional neural network	0.7464	0.7750	0.74	0.8
Alhanai et al. [62]	Combination of LSTM and CNN	*	0.77	0.83	*
Niu et al. [69]	Hierarchical context-aware graph attention model	*	0.92	0.92	*

those potential depression patients. The model with two attention layers led to results that outperformed the state of the art, Niu *et al.* by 8% in terms of sensitivity. In comparison with Alhanai *et al.* who adopted a similar method made up of CNN and LSTM, our proposed method was better by 17% in terms of sensitivity. This is not conclusive since the dataset for evaluation in their article was slightly different from ours. By conducting a student t-test between the F_1 of the best patient-level audio model with the result of p-value = 0.0099 (<0.01), as well as patient-level text model with the result of p-value = 0.0246 (<0.05), we could conclude that multimodality models statistically significant outperformed single modality models.

5 CONCLUSION

In this paper, a multimodality approach for automated depression detection was presented. Firstly, we performed the statistical test to investigate the difference between the audio and text features of severe and healthy subjects. We proved the pattern of severe depression patients was different from that of the healthy. Therefore, the audio feature sequence carried information that could be used to predict depression severity. Secondly, models that considered audio and text features individually were trained and evaluated at the patient-independent level. These unimodality models then acted as feature extractors and output features were combined by audio-text fused model. For the audio modality, at the patient-independent level, the model comprised of single-layer Bi-LSTM and five stacked T-CNN blocks achieved the best sequence level F_1 score of 0.9870 and patient-level F_1 score of 0.9074 with the test set. This result indicates that the Bi-LSTM provides a more reliable representation, from which the automated depression

detection model could benefit. Additionally, we evaluated the patient-independent audio models with different time-steps with the Area Under Curve (AUC) metric. We concluded that the 16-timestep model performed best and the micro-average AUC was higher than any other model. However, the 64-timestep model showed its strength in detecting the audio feature sequence from the mild patient, which met our expectation that the model should be able to distinguish mild patients so that clinical interference can be conducted in the early stage. Overall, the 16-timestep model outperformed the 32-timestep and 64-timestep models, which could be attributed to the relatively low signal-noise ratio of the shorter input sequences and the memorization limit of the LSTM. The new understanding assisted in our model selection and hyper-parameter configuration when we deployed this method in clinical settings. These findings provided the following insight for future research, our proposed unimodality model was patient-independent, and the prediction was based on a period of audio/text features. Therefore, compared with other models, our proposed model did not have limitations to the length of interview audio or transcript, which made it possible for people to monitor their mental state in daily use.

Moreover, for the text modality, the model consisted of Bi-LSTM and three fully connected layers achieved the best sequence level F_1 score of 0.9709 and patient-level F_1 score of 0.9245 on the test set. We conducted experiments to investigate the influence of the text model hyper-parameters, such as window size and stop words. We found the best window size is 64. In our experiment, we investigated the effect of stop words, the result indicated the text model performs better if the stop words were removed in advance. Currently, our patient-level prediction was carried out by a major voting algorithm, which yielded a patient-level

depression prediction model with satisfying performance. Our proposed multimodal method achieved the highest F_1 of 0.9580 on the patient-level depression detection task, which showed a significant improvement over the previous state-of-the-art. In the future, a study on how to represent the audio/text features during the whole interview should be carried out so that the model could make patient-level predictions based on a digest of text and audio features.

ACKNOWLEDGMENTS

We would like to acknowledge the funding support from the MITACS grant, Canada. The authors would also like to thank Jo-Sheen Yen for proofreading the manuscript.

REFERENCES

- [1] M. G. Mazza *et al.*, "Anxiety and depression in COVID-19 survivors: Role of inflammatory and clinical predictors," *Brain, Behavior, Immun.*, vol. 89, pp. 594–600, 2020.
- [2] N. Salari *et al.*, "Prevalence of stress, anxiety, depression among the general population during the COVID-19 pandemic: A systematic review and meta-analysis," *Globalization Health*, vol. 16, no. 1, pp. 1–11, 2020.
- [3] R. Barzilay *et al.*, "Resilience, COVID-19-related stress, anxiety and depression during the pandemic in a large population enriched for healthcare providers," *Transl. Psychiatry*, vol. 10, no. 1, pp. 1–8, 2020.
- [4] H. C. Nguyen *et al.*, "People with suspected COVID-19 symptoms were more likely depressed and had lower health-related quality of life: The potential benefit of health literacy," *J. Clin. Med.*, vol. 9, no. 4, p. 965, 2020.
- [5] C. D. Mathers and D. Loncar, "Projections of global mortality and burden of disease from 2002 to 2030," *PLoS Med.*, vol. 3, no. 11, 2006, Art. no. e442.
- [6] J. Olesen *et al.*, "The economic cost of brain disorders in Europe," *Eur. J. Neurol.*, vol. 19, no. 1, pp. 155–162, 2012.
- [7] W. F. Stewart, J. A. Ricci, E. Chee, S. R. Hahn, and D. Morganstein, "Cost of lost productive work time among us workers with depression," *Jama*, vol. 289, no. 23, pp. 3135–3144, 2003.
- [8] S. L. James *et al.*, "Global, regional, and national incidence, prevalence, and years lived with disability for 354 diseases and injuries for 195 countries and territories, 1990–2017: A systematic analysis for the global burden of disease study 2017," *Lancet*, vol. 392, no. 10159, pp. 1789–1858, 2018.
- [9] A. Rahman, A. Malik, S. Sikander, C. Roberts, and F. Creed, "Cognitive behaviour therapy-based intervention by community health workers for mothers with depression and their infants in rural Pakistan: A cluster-randomised controlled trial," *Lancet*, vol. 372, no. 9642, pp. 902–909, 2008.
- [10] World Health Organization, "Preventing suicide: A global imperative," Geneva, Switzerland: World Health Org., 2014.
- [11] K. Hawton, C. C. I Comabella, C. Haw, and K. Saunders, "Risk factors for suicide in individuals with depression: A systematic review," *J. Affect. Disord.*, vol. 147, no. 1–3, pp. 17–28, 2013.
- [12] J.-P. Lépine and M. Briley, "The increasing burden of depression," *Neuropsychiatric Dis. Treat.*, vol. 7, no. Suppl 1, p. 3, 2011.
- [13] T. E. Joiner Jr., J. S. Brown, and L. R. Wingate, "The psychology and neurobiology of suicidal behavior," *Annu. Rev. Psychol.*, vol. 56, pp. 287–314, 2005.
- [14] A. McGirr *et al.*, "An examination of dsm-iv depressive symptoms and risk for suicide completion in major depressive disorder: A psychological autopsy study," *J. Affect. Disord.*, vol. 97, no. 1–3, pp. 203–209, 2007.
- [15] S. Rodrigues *et al.*, "Impact of stigma on veteran treatment seeking for depression," *Amer. J. Psychiatr. Rehabil.*, vol. 17, no. 2, pp. 128–146, 2014.
- [16] M. B. First and M. Gibbon, "The structured clinical interview for DSM-IV axis I disorders (SCID-I) and the structured clinical interview for DSM-IV axis II disorders (SCID-II)," 2004.
- [17] K. Kroenke, R. L. Spitzer, and J. B. Williams, "The PHQ-9: Validity of a brief depression severity measure," *J. Gen. Intern. Med.*, vol. 16, no. 9, pp. 606–613, 2001.
- [18] A. T. Beck, C. H. Ward, M. Mendelson, J. Mock, and J. Erbaugh, "An inventory for measuring depression," *Arch. Gen. Psychiatry*, vol. 4, no. 6, pp. 561–571, 1961.
- [19] S. A. Montgomery and M. A. R. I. E. Åsberg, "A new depression scale designed to be sensitive to change," *British J. Psychiatry*, vol. 134, no. 4, pp. 382–389, 1977.
- [20] S. Rude, E.-M. Gortner, and J. Pennebaker, "Language use of depressed and depression-vulnerable college students," *Cogn. Emotion*, vol. 18, no. 8, pp. 1121–1133, 2004.
- [21] R. M. Bagby, A. G. Ryder, D. R. Schuller, and M. B. Marshall, "The hamilton depression rating scale: Has the gold standard become a lead weight?," *Amer. J. Psychiatry*, vol. 161, no. 12, pp. 2163–2177, 2004.
- [22] H. C. Kraemer, D. J. Kupfer, D. E. Clarke, W. E. Narrow, and D. A. Regier, "DSM-5: How reliable is reliable enough?," *Amer. J. Psychiatry*, vol. 169, no. 1, pp. 13–15, 2012.
- [23] S. Kapur, A. G. Phillips, and T. R. Insel, "Why has it taken so long for biological psychiatry to develop clinical tests and what to do about it?," *Mol. Psychiatry*, vol. 17, no. 12, pp. 1174–1179, 2012.
- [24] M. E. Aragón, A. P. López-Monroy, L. C. González-Gurrola, and M. Montes, "Detecting depression in social media using fine-grained emotions," in *Proc. Conf. North Amer. Chapter Assoc. Comput. Linguistics: Hum. Lang. Technol., Volume 1 (Long and Short Papers)*, 2019, pp. 1481–1486.
- [25] A. G. Reece, A. J. Reagan, K. L. Lix, P. S. Dodds, C. M. Danforth, and E. J. Langer, "Forecasting the onset and course of mental illness with twitter data," *Sci. Rep.*, vol. 7, no. 1, pp. 1–11, 2017.
- [26] J. Han *et al.*, "Deep learning for mobile mental health: Challenges and recent advances," *IEEE Signal Process. Mag.*, vol. 38, no. 6, pp. 96–105, 2021.
- [27] C. C. Joyal, J.-L. Dubreucq, C. Gendron, and F. Millaud, "Major mental disorders and violence: A critical update," *Curr. Psychiatry Rev.*, vol. 3, no. 1, pp. 33–50, 2007.
- [28] N. Cummins, F. Matcham, J. Klapper, and B. Schuller, "Artificial intelligence to aid the detection of mood disorders," in *Proc. Artif. Intell. Precis. Health*. New York, NY, USA: Elsevier, 2020, pp. 231–255.
- [29] M. Morales, S. Scherer, and R. Levitan, "A linguistically-informed fusion approach for multimodal depression detection," in *Proc. 5th Workshop Comput. Linguistics Clin. Psychol.: From Keyboard Clinic*, 2018, pp. 13–24.
- [30] I. Pirina and Ç. Çöltekin, "Identifying depression on reddit: The effect of training data," in *Proc. EMNLP Workshop SMM4H: The 3rd Social Media Mining Health Appl. Workshop Shared Task*, 2018, pp. 9–12.
- [31] A. Yates, A. Cohan, and N. Goharian, "Depression and self-harm risk assessment in online forums," 2017, *arXiv:1709.01848*.
- [32] A. Ray, S. Kumar, R. Reddy, P. Mukherjee, and R. Garg, "Multi-level attention network using text, audio and video for depression prediction," in *Proc. 9th Int. Audio/Visual Emotion Challenge Workshop*, 2019, pp. 81–88.
- [33] P. Resnik, A. Garron, and R. Resnik, "Using topic modeling to improve prediction of neuroticism and depression in college students," in *Proc. Conf. Empirical Methods Natural Lang. Process.*, 2013, pp. 1348–1353.
- [34] K. R. Scherer, "Vocal affect expression: A review and a model for future research," *Psychol. Bull.*, vol. 99, no. 2, p. 143, 1986.
- [35] K. Kroenke, R. L. Spitzer, J. B. Williams, and B. Löwe, "The patient health questionnaire somatic, anxiety, and depressive symptom scales: A systematic review," *Gen. Hosp. Psychiatry*, vol. 32, no. 4, pp. 345–359, 2010.
- [36] M. Hiraga, "Predicting depression for japanese blog text," in *Proc. ACL 2017, Student Res. Workshop*, 2017, pp. 107–113.
- [37] G. Coppersmith, M. Dredze, C. Harman, K. Hollingshead, and M. Mitchell, "Clipsych 2015 shared task: Depression and ptsd on twitter," in *Proc. 2nd Workshop Comput. Linguistics Clin. Psychology: From Linguistic Signal Clin. Reality*, 2015, pp. 31–39.
- [38] S. Scherer *et al.*, "Automatic audiovisual behavior descriptors for psychological disorder analysis," *Image Vis. Comput.*, vol. 32, no. 10, pp. 648–658, 2014.
- [39] M. Cannizzaro, B. Harel, N. Reilly, P. Chappell, and P. J. Snyder, "Voice acoustical measurement of the severity of major depression," *Brain Cogn.*, vol. 56, no. 1, pp. 30–35, 2004.
- [40] S. Scherer, G. Stratou, J. Gratch, and L.-P. Morency, "Investigating voice quality as a speaker-independent indicator of depression and ptsd," in *Proc. Interspeech*, 2013, pp. 847–851.

- [41] E. Kraepelin, "Manic depressive insanity and paranoia," *J. Nervous Ment. Dis.*, vol. 53, no. 4, p. 350, 1921.
- [42] N. Cummins, S. Scherer, J. Krajewski, S. Schnieder, J. Epps, and T. F. Quatieri, "A review of depression and suicide risk assessment using speech analysis," *Speech Commun.*, vol. 71, pp. 10–49, 2015.
- [43] J. R. Williamson, T. F. Quatieri, B. S. Helfer, R. Horwitz, B. Yu, and D. D. Mehta, "Vocal biomarkers of depression based on motor incoordination," in *Proc. 3rd ACM Int. Workshop Audio/Vis. Emotion Challenge*, 2013, pp. 41–48.
- [44] S. Alghowinem *et al.*, "A comparative study of different classifiers for detecting depression from spontaneous speech," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process.*, 2013, pp. 8022–8026.
- [45] N. Cummins, J. Epps, and E. Ambikairajah, "Spectro-temporal analysis of speech affected by depression and psychomotor retardation," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process.*, 2013, pp. 7542–7546.
- [46] E. Moore II, M. A. Clements, J. W. Peifer, and L. Weissner, "Critical analysis of the impact of glottal features in the classification of clinical depression in speech," *IEEE Trans. Biomed. Eng.*, vol. 55, no. 1, pp. 96–107, Jan. 2008.
- [47] J. F. Cohn *et al.*, "Detecting depression from facial actions and vocal prosody," in *Proc. IEEE 3rd Int. Conf. Affect. Comput. Intell. Interact. Workshops*, 2009, pp. 1–7.
- [48] A. C. Trevino, T. F. Quatieri, and N. Malyska, "Phonologically-based biomarkers for major depressive disorder," *EURASIP J. Adv. Signal Process.*, vol. 2011, no. 1, p. 42, 2011.
- [49] D. Bone, C.-C. Lee, and S. Narayanan, "Robust unsupervised arousal rating: A rule-based framework with knowledge-inspired vocal features," *IEEE Trans. Affect. Comput.*, vol. 5, no. 2, pp. 201–213, Apr.–Jun. 2014.
- [50] J. Gratch *et al.*, "The distress analysis interview corpus of human and computer interviews," in *Proc. 9th Int. Conf. Lang. Resour. Eval.*, 2014, pp. 3123–3128.
- [51] K. Kroenke, T. W. Strine, R. L. Spitzer, J. B. Williams, J. T. Berry, and A. H. Mokdad, "The phq-8 as a measure of current depression in the general population," *J. Affect. Disord.*, vol. 114, no. 1–3, pp. 163–173, 2009.
- [52] E. Crawford and J. Pineau, "Spatially invariant unsupervised object detection with convolutional neural networks," in *Proc. AAAI Conf. Artif. Intell.*, 2019, vol. 33, pp. 3412–3420.
- [53] G. Degottex, J. Kane, T. Drugman, T. Raitio, and S. Scherer, "Covarep—A collaborative voice analysis repository for speech technologies," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process.*, 2014, pp. 960–964.
- [54] P. Alku, T. Bäckström, and E. Vilkman, "Normalized amplitude quotient for parametrization of the glottal flow," *J. Acoustical Soc. Amer.*, vol. 112, no. 2, pp. 701–710, 2002.
- [55] T. Hacki, "Klassifizierung von glottiscysfunktionen mit hilfe der elektrolottographie [classification of glottal dysfunctions on the basis of electroglottography]," *Folia phoniatrica*, vol. 41, no. 1, pp. 43–48, 1989.
- [56] J. Kane, C. Gobl, S. Scherer, and L.-P. Morency, "A comparative study of glottal open quotient estimation techniques," *BDL*, vol. 178, no. 15.17, pp. 0–41, 2013.
- [57] E. B. Holmberg, R. E. Hillman, J. S. Perkell, P. C. Guio, and S. L. Goldman, "Comparisons among aerodynamic, electroglottographic, and acoustic spectral measures of female voice," *J. Speech, Lang., Hear. Res.*, vol. 38, no. 6, pp. 1212–1223, 1995.
- [58] P. Alku, H. Strik, and E. Vilkman, "Parabolic spectral parameter—a new method for quantification of the glottal flow," *Speech Commun.*, vol. 22, no. 1, pp. 67–79, 1997.
- [59] J. Kane and C. Gobl, "Wavelet maxima dispersion for breathy to tense voice discrimination," *IEEE Trans. Audio, Speech, Lang. Process.*, vol. 21, no. 6, pp. 1170–1179, Jun. 2013.
- [60] X. Huang, A. Acero, H.-W. Hon, and R. Reddy, *Spoken Language Processing: A Guide to Theory, Algorithm, and System Development*. Upper Saddle River, NJ, USA: Prentice hall PTR, 2001.
- [61] M. R. Makiuchi, T. Warnita, K. Uto, and K. Shinoda, "Multimodal fusion of bert-cnn and gated CNN representations for depression detection," in *Proc. 9th Int. Audio/Vis. Emotion Challenge Workshop*, 2019, pp. 55–63.
- [62] T. Al Hanai, M. M. Ghassemi, and J. R. Glass, "Detecting depression with audio/text sequence modeling of interviews," in *Proc. Interspeech*, 2018, pp. 1716–1720.
- [63] S. Bird, E. Klein, and E. Loper, *Natural Language Processing With Python: Analyzing Text With the Natural Language Toolkit*. Boston, MA, USA: O'Reilly Media, Inc., 2009.
- [64] J. Pennington, R. Socher, and C. D. Manning, "Glove: Global vectors for word representation," in *Proc. Conf. Empirical Methods Natural Lang. Process.*, 2014, pp. 1532–1543.
- [65] C. Shin, S.-H. Lee, K.-M. Han, H.-K. Yoon, and C. Han, "Comparison of the usefulness of the phq-8 and phq-9 for screening for major depressive disorder: Analysis of psychiatric outpatient data," *Psychiatry Investigation*, vol. 16, no. 4, p. 300, 2019.
- [66] I. Madhavi, S. Chamishka, R. Nawaratne, V. Nanayakkara, D. Alahakoon, and D. De Silva, "A deep learning approach for work related stress detection from audio streams in cyber physical environments," in *Proc. 25th IEEE Int. Conf. Emerg. Technol. Factory Automat.*, 2020, vol. 1, pp. 929–936.
- [67] L. Yang, H. Sahli, X. Xia, E. Pei, M. C. Oveneke, and D. Jiang, "Hybrid depression classification and estimation from audio video and text information," in *Proc. 7th Annu. Workshop Audio/Vis. Emotion Challenge*, 2017, pp. 45–51.
- [68] N. Srimadhur and S. Lalitha, "An end-to-end model for detection and assessment of depression levels using speech," *Procedia Comput. Sci.*, vol. 171, pp. 12–21, 2020.
- [69] M. Niu, K. Chen, Q. Chen, and L. Yang, "HCAG: A hierarchical context-aware graph attention model for depression detection," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process.*, 2021, pp. 4235–4239.



Kaining Mao is currently working toward the PhD degree with the Electrical Computer Engineering Department, University of Alberta. His work focuses specifically on multimodal machine learning (MML) and its impact on automated depression diagnosis (ADD).



Wei Zhang is currently working toward the PhD degree in electrical computer engineering. He is working on artificial intelligence field to detect tuberculosis and depression in different projects.



Ang Li is currently working toward the BSc degree in computer engineering with the University of Alberta. He is currently conducting the Dean Research Awards Program focuses on the application of AI in depression recognition. His research interests include machine learning and embedded system design.



Deborah Baofeng Wang received the MBA degree in management information system concentration and the PhD degree in educational psychology. Dr. Wang is responsible for the management of several international R&D projects with Wenzhou Kangning Hospital, focusing on the application of AI in mental health. With her MBA and PhD degrees, Dr. Wang has extensive experiences as a senior research analyst in the fields of mental health, public health, and education handling longitudinal and cross-sectional data with national (U.S.) and international levels.

Rongqi Jiao received the bachelor's degree in medicine. She is currently working toward the master's degree with Wenzhou Medical University. With the bachelor's degree and the opportunity to learn as a psychiatrist with Wenzhou Kangning Hospital, she has systematic academic training in medicine and hands-on research and clinical experiences in the area of mental health.

Yanhui Zhu is currently working toward the master's degree majoring in psychiatry and mental health with the School of Mental Health, Wenzhou Medical University. The primary focus of her research in the master's program has been on microexpression and mental illness. She is currently a clinical practitioner of mental health with the Second Affiliated Hospital of Wenzhou Medical University.

Bin Wu received the bachelor's degree in medicine. He is currently working toward the master's degree with Wenzhou Medical University. He is also a psychiatrist with Wenzhou Kangning Hospital. He has received systematic academic training in the area of mental health and has solid hands-on experience in clinical practice of mental health, particularly the application of AI in mental health.

Tiansheng Zheng received the master's degree in psychiatry. She is accredited as an associate chief physician of psychiatry. She is also the deputy director of Depression Center and Sleep Center, Affiliated Kangning Hospital, Wenzhou Medical University. Her research interests include cognitive function in patients with depression and insomnia disorders.

Lei Qian is currently working toward the master's degree in clinical psychology. He is accredited as an intermediate psychotherapist. He is also the deputy director of Higher Education Collaborations with the Affiliated Kangning Hospital, Wenzhou Medical University. His research interests include data analysis and mental health.



Wei Lyu accredited as a chief physician of psychiatry and intermediate psychotherapist. He is currently the vice president of Wenzhou Kangning Hospital and an adjunct professor with the School of Mental Health, Wenzhou Medical University. He has extensive practical experiences as a psychiatric practitioner. He is also actively involved in the field of mental health as a professional and also a leader. He is also a youth member with the Society of Psychiatrists, Zhejiang Medical Doctor Association.



Minjie Ye is currently the vice president and director of the Center for Child and Adolescent Mental Health Services of the Affiliated Kangning Hospital, Wenzhou Medical University. He is also the academic lead for Kangning Hospital's National Key Clinical Specialty (Psychiatry). In addition to his position as the associate dean of the School of Mental Health, Wenzhou Medical University, he is also heads the University's Child Mental Health Research Center. Dr. Ye is a member of the Society of Psychiatrists with Chinese

Medical Doctor Association and a member of Committee of Marriage and Family Psychology with the Chinese Association of Social Psychology.



Jie Chen (Fellow, IEEE) received the PhD degree from the University of Maryland at College Park, USA. He is currently a professor with the University of Alberta. He is a fellow of the Canadian Academy of Engineering and American Institute of Medical and Biological Engineering. He was the recipient of the Distinguished Alumni Award, the Department of Electrical and Computer Engineering, University of Maryland. He was also awarded the Killam Annual Professorship, which is among the highest honours given to a professor with Canadian Universities.

▷ For more information on this or any other computing topic, please visit our Digital Library at www.computer.org/csdl.