Interpreting Depression from Question-wise Long-term Video Recording of SDS Evaluation

Wanqing Xie[†], Lizhong Liang[†], Yao Lu, Chen Wang, Jihong Shen, Hui Luo*, Xiaofeng Liu*, *Member, IEEE*

Abstract—Self-Rating Depression Scale (SDS) questionnaire has frequently been used for efficient depression preliminary screening. However, the uncontrollable self-administered measure can be easily affected by insouciantly or deceptively answering, and producing the different results with the clinician-administered Hamilton Depression Rating Scale (HDRS) and the final diagnosis. Clinically, facial expression (FE) and actions play a vital role in clinician-administered evaluation, while FE and action are underexplored for self-administered evaluations. In this work, we collect a novel dataset of 200 subjects to evidence the validity of self-rating questionnaires with their corresponding question-wise video recording. To automatically interpret depression from the SDS evaluation and the paired video, we propose an end-to-end hierarchical framework for the long-term variable-length video, which is also conditioned on the questionnaire results and the answering time. Specifically, we resort to a hierarchical model which utilizes a 3D CNN for local temporal pattern exploration and a redundancy-aware self-attention (RAS) scheme for question-wise global feature aggregation. Targeting for the redundant long-term FE video processing, our RAS is able to effectively exploit the correlations of each video clip within a question set to emphasize the discriminative information and eliminate the redundancy based on feature pair-wise affinity. Then, the question-wise video feature is concatenated with the questionnaire scores for final depression detection. Our thorough evaluations also show the validity of fusing SDS evaluation and its video recording, and the superiority of our framework to the conventional state-of-the-art temporal modeling methods.

Index Terms—Depression detection, Self-Rating Depression Scale, Facial expression video, Sparse self-attention

I. INTRODUCTION

DEPRESSION, a.k.a. major depressive disorder, is a common and serious mental health disorder, while it can be treated [1]. Effective early diagnosis can be beneficial

W. Xie is with the College of the Mathematical Sciences, Harbin Engineering University, Harbin, China.

L. Liang is with the School of Computer Science and Engineering, Sun Yat-sen University, Guangzhou 510006, China, and the Affiliated Hospital of Guangdong Medical University, Zhanjiang, Guangdong, 524023, China.

Y. Lu is with the School of Computer Science and Engineering, Sun Yat-sen University, Guangzhou 510006, China.

C. Wang is with the College of the Mathematical Sciences, Harbin Engineering University, Harbin, China.

J. Shen is with the College of the Mathematical Sciences, Harbin Engineering University, Harbin, China.

H. Luo is with the Southern Marine Science and Engineering Guangdong Laboratory (Zhanjiang), Zhanjiang, Guangdong, 524023 and the Marine Biomedical Research Institute of Guangdong, Zhanjiang, 524023, China.

X. Liu is with the Dept. of Neurology, Beth Israel Deaconess Medical Center, Harvard Medical School, Boston, MA, USA, and Suzhou Fanhan Information Technology Co., Ltd, Suzhou, China. (xliu61@mgh.harvard.edu)

†W. Xie and L. Liang contribute equally. *H. Luo and X. Liu are the corresponding authors.

Fig. 1. The illustration of collecting SDS questionnaires and the corresponding synchronized face videos. The depression can be better detected with both SDS score and facial expression videos.

for intervention therapy. The earlier that treatment can begin, the more effective it is [2]. However, the comprehensive clinical interview for final diagnosis, i.e., clinical golden standard [3], can be costly for the large population screening.

Depression is even more prevalent during the COVID-19 pandemic [4], [5], while the in-person clinician interview can be inconvenient or even prohibitive.

The Self-Rating Depression Scale (SDS) [6] is a widely adopted self administrative fast screening questionnaire with twenty questions, which involves affective, psychological, and somatic symptoms related to depression[1]. Each question is framed in terms of positive and negative statements, and be scored on a Likert scale ranging from 1 to 4. The final result is the sum of each question. We note that the larger score usually indicates the subject is more likely to be a depression patient. Conventionally, we set 50 as the threshold of normal or depression [7]. The SDS ratings have indicative depression level ranges that may help health assessment and testing [8].

The ratings include suggestive depression level ranges that may help therapeutic and scientific research, but the SDS outcome can vary from the clinical interview for verifying a depression diagnosis [9]. A reason can be the uncontrollable self-administered measure can be easily affected by insouciantly or deceptively answering [10], and producing the different results with the clinician-administered interview, e.g., Hamilton Depression Rating Scale (HDRS) [11].

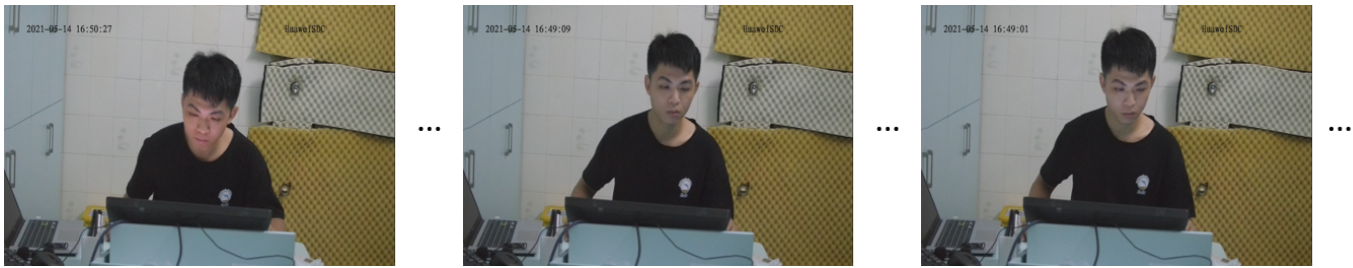[1]https://www.who.int/substance_abuse/research_tools/en/english_zung.pdf

Fig. 2. Illustration of a subject answering a question in SDS. A Software-Defined Camera is adopted with the parameters of 25 fps and 3840×2160 resolution.

Clinically, facial expression (FE) [12] and the actions [13], [14] can play an important role in clinician-administered evaluation, while FE and actions are underexplored for self-administered evaluation. Actually, expression and actions can be expressive features for many psychiatry analyses [15], [16]. Based on this insight, in this work, we collect a novel dataset of 200 subjects to evidence the validity of self-rating questionnaires with video recording. To provide a more fine-grained connection between the questionnaire and the video, we adopt a Software-Defined Camera (SDC) system which synchronized with the questionnaire software to record the video start from the question showing and end on the score is chosen. For each subject, there are 20 question and video pairs. To extract the region of interests (ROI), the face detector is applied, and the face box is extended 100% to incorporate the hand action, e.g., head-scratching, chin-touching. Moreover, the answering time may also affect the depression diagnosis [17].

The video provides additional information for analysis, while it also introduces several challenges for automatical analysis. Firstly, the video of each question can be quite long (e.g., 525 frames in our dataset), and the useful information can be sparse in a long-term sequence. 2) Secondly, the length of each video varies from 50 to 525 frames, which is depends on the question and the participants. 3) Moreover, a practical information fusion scheme is necessary to explore both the SDS evaluation and its question-wise video recording.

To automatically interpret depression from the SDS evaluation and its corresponding FE and action video recording, a redundancy-aware conditional self-attention framework is proposed. Specifically, we resort to a hierarchical model which utilizes a 3D convolutional neural network (CNN) [18] for local temporal pattern exploration and a self-attention scheme for question-wise global feature aggregation. We factorize the question-wise video into the fix-length clips according to the time of human facial expression developing. With the fix-length input, the 3D CNN is able to extract the local temporal cues efficiently. Then, the clip-wise representation is feed forwarded to a parametric redundancy-aware self-attention (RAS) scheme to eliminate the uninformative signals and extract the representative question-wise feature.

Targeting for the redundant long-term FE and action video processing, our RAS is able to effectively exploit the correlations of each video clip in a question to emphasize the discriminative information and eliminate the redundancy. It traverses the clips within a question to produce a refined representation based on pair-wise feature affinity [19]. Our residual attention term is able to prioritize discriminative clips while ignoring inferior ones for explicit redundancy reduction. In addition, the temporal sequence is explicitly considered with a Gaussian similarity kernel.

Then, the question-wise video feature is concatenated with the questionnaire scores for final depression detection. Our thorough evaluations also show the validity of fusing SDS evaluation and its video recording, and the superiority of our framework to the conventional state-of-the-art temporal modeling methods.

The main contribution of this work can be summarized as follows:

• To the best of our knowledge, this is the first attempt at exploring depression with both SDS evaluation and its corresponding question-wise face video recording. An elaborate synchronized system is designed, and the final clinician interview results are collected.

• A practical hierarchical conditional self-attention framework is proposed to explore the long-term variable-length video with 3D CNN for local temporal modeling, redundancy-aware RAS for global attention modeling, and SDS-score conditional question-wise fusion.

• Our parametric redundancy-aware self-attention (RAS) scheme explicitly emphasizes the discriminative clips and reduces the redundancy based on feature pair-wise affinity, and is aware of the temporal sequence with a Gaussian similarity kernel.

• The systematic and thorough comparisons with the previous temporal modeling methods provide further insights into the potential benefits of our framework. We note that the proposed framework can potentially be generalized to other classification tasks using both questionnaire and video modalities.

II. RELATED WORK

Deep learning for depression detection

In recent decays, there are numerous works dedicated to collect better quality, and larger quantities of depression datasets [13]. There has been a long history of using SDS evaluation for self-report [7]. In addition, the facial expression [20], eye movement [14] and body action [21] can be the important modality for depression detection [13]. However, publicly available SDS and its video recording datasets appropriate for incorporating machine learning methods are missing. To the best of our

TABLE I
THE STATISTICS OF COLLECTED 200 SUBJECTS. M INDICATES MALE,
WHILE F INDICATES FEMALE.

Final diagnosis	SDS result	Number
Normal	Normal	86 (M:42;F:44)
Normal	Depression	20 (M:13;F:7)
Depression	Normal	20 (M:11;F:9)
Depression	Depression	74 (M:38;F:36)

knowledge, this is the first attempt to automatically explore both the SDS evaluation and the corresponding question-wise video.

Moreover, the "ground truth" of many datasets are the self report (e.g., DAIC-WOZ [22] and AVEC [23]), which is highly unreliable [13]. The clinician interview has usually been used for final diagnosis [24], which can be costly for large-scale labeling. In this work, all subjects have taken a more comprehensive clinical interview to collect the golden standard label of depression. The scale of our collection, i.e., 200 subjects, is able to support the automatic analysis with the deep learning system.

We note that many modalities can be used for depression detection. For example, [25] targets to the conversation with the Patient Health Questionnaire (PHQ)-8 metric [26]. [27] propose to fuse the spoken language and 3D facial key points in DAIC-WOZ dataset [22]. The audio is used in [28] with a self-supervised embedding, and the phoneme feature is used in [29]. These works try to mimic the clinician-based interview. More recently, the electroencephalography and paralinguistic behaviors are fused with the classifier ensemble for depression detection [30]. [31] explores the EEG signal via kernel-target alignment. The fNIRS can also be used for diagnosis [32]. These modalities are able to provide more accurate features, while they are not salable for efficient screening as SDS.

Temporal modeling

Temporal modeling [33] can be the essential part of the video-based classification and analysis tasks, e.g., video facial expression [34], [35], [36] and action recognition [37]. The recursive neural network (RNN) [38] is widely used for temporal patterns modeling, which makes dependent sequential processing. However, it is notorious for the long-term forgetting and hard to train [39]. Therefore, its performance can be largely affected if the input is too long. The bi-directional long-short term memory (Bi-LSTM) has been proposed to alleviate the difficulties [40], [41], which is still relatively slow in both training and inference. More recently, the 3D CNN [42] is proposed to explore the spatial-temporal patterns in a unified manner. It can be fast processed and has demonstrated good performances in many tasks. Nevertheless, the input to the 3D CNN should be fixed [43], limiting its application to the video with variable length, e.g., different questions in our dataset. In this work, we propose to factorize a variable-length video into several fixed-length clips to utilize the 3D CNN to better balance performance and efficiency.

Self-attention and non-local scheme

Start from the machine translation in [44], the attention scheme has demonstrated great potential in many applications. It has been the core block of many successful systems [45].
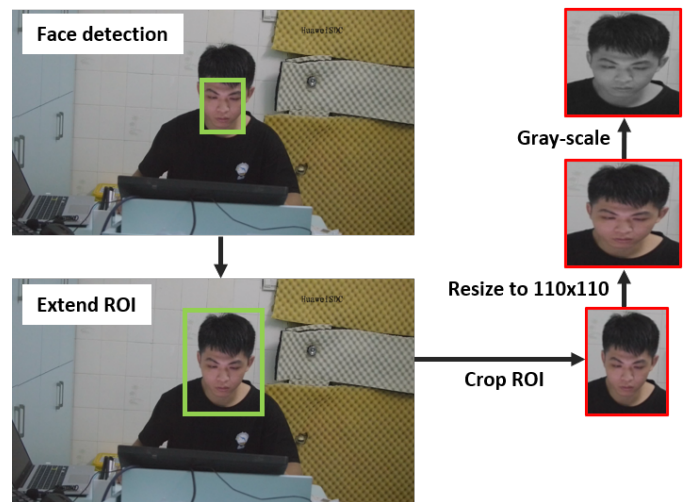

Fig. 3. Illustration of the pre-processing flow, which consists of face detection, ROI extension, ROI crop, resize, and gray-scale processing, sequentially.

Conventionally, it computes the adjusted output at a position with the weighted sum of all positions in that sentence. A similar philosophy has also been inherited in the non-local algorithms [46], which focused on the image denoising task. Pair-wise relationships were also modeled using interaction networks [47], [48], [49], [50]. Moreover, [19] proposes a link between self-attention and the wider category of non-local filtering activities. [51] proposes to learn various time scales of temporal dependencies between video frames. Inspired by these methods, we further adapt this idea to the variable length long-term SDS video analysis.

III. METHODOLOGY

A. Data Acquisition and Preprocessing

We collect the Self-Rating Depression Scale (SDS) questionnaires and the corresponding face video from 200 subjects. The study protocol was approved by the Ethics Committee of the Affiliated Hospital of Guangdong Medical University (No. PJ2021-026).

Each participant is instructed to sit in a quiet consulting room along and fill the self-report questionnaire following the instruction in the software interface to avoid being affected by the others. Moreover, a Huawei Software-Defined Camera (SDC) is hidden behind a one-way mirror, and the participants do not become aware of the camera in the evaluation. SDC adopts an open-ended software architecture that can flexibly integrate with the code of specific application[2]. In addition, the data is uploaded to the back-end server for processing. To connect each question with the video, the camera is synchronized with the questionnaire software to record the video start from the question showing and end on the score is chosen. The illustration of our data collection is shown in Fig. 1. In Fig. 2, we provide some frames of a subject for answering a question.
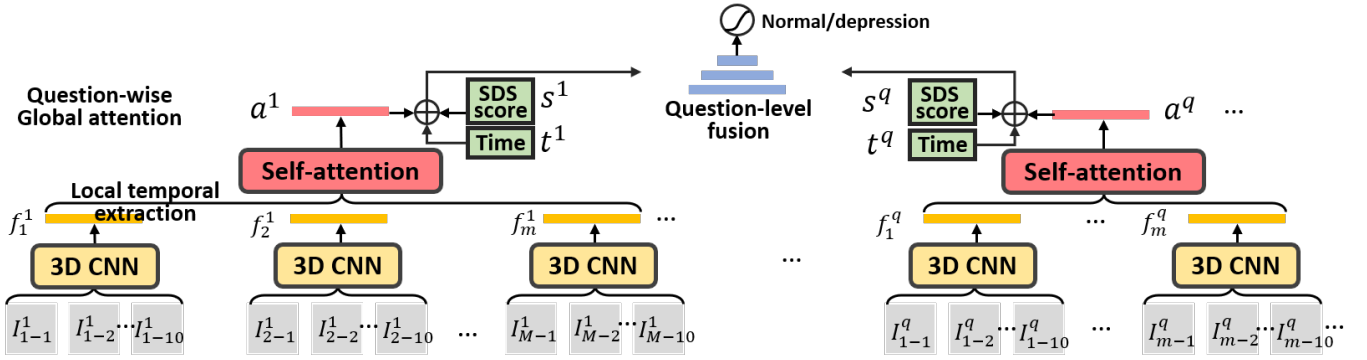
Fig. 4. Illustration of the proposed hierarchical conditional self-attention framework, comprising three hierarchically mounted modules for both self-rating questionnaires with video recording, i.e., 3D CNN-based local temporal pattern extraction module, redundancy-aware self-attention module, and question-wise conditional fusion module.

There are 20 questions for SDS evaluation, and they usually take about 10 minutes to complete[7]. Each participant takes a different time to finish different questions. In our collected dataset, the time of each question varies from 2s to 21s. The 25 fps videos with the resolution of 3840×2160 are collected for each question. Since the region of interest is the face of participants, we use the face detector [52] to crop the face region. To incorporate the hand action of head-scratching [53], chin-touching [54], etc., we extend the face box by 100% and resize the extended region of interests (ROI) in each frame to 110×110. Moreover, the image is followed by gray processing to reduce the size of the input. The pre-processing flow chart is given in Fig. 3.

After the self-report SDS evaluation, a more comprehensive clinical interview [3], including the clinician-administered Hamilton Depression Rating Scale (HDRS) assessment [11], SCL-90-R Symptom Checklist [55], and Self Rating Anxiety Scale (SAS)[56], is taken to help the clinician for confirming a diagnosis of depression. We use the final diagnosis result as our ground truth label.

Considering the self-administrated SDS can be uncontrollable, the result of SDS evaluation can be different from the final diagnosis [9]. In Tab. I, we provide the detailed statistics of the SDS and final results of our 200 subjects. We note that there are only two classes, i.e., normal and depression, in our dataset. We can see that there are about 10% of subjects have different results. Specifically, the SDS evaluation of 10 subjects is diagnosed as depression, while the other 10 subjects with the high SDS score are diagnosed as normal with the subsequent clinical interview.

B. Hierarchical conditional framework

The long-term video recording of SDS evaluation inherits rich emotional information, while it also poses several challenges for processing. First, the long video can be redundant, and only very few (i.e., sparse) frames may indicate the useful cues for depression detection. Second, the length of the video can be varied across different subjects and questions.

Considering that the human expression usually takes 200ms to 500ms, it can be reasonable to segment the video into several fixed-length short clips and explore the local temporal patterns
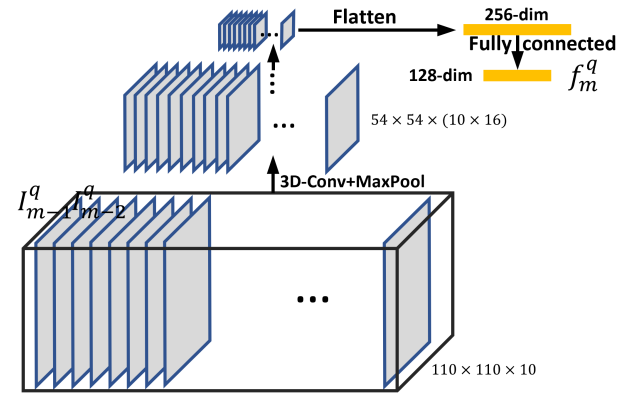

Fig. 5. Illustration of the 3D CNN module for local temporal pattern extraction.

within it. We empirically set each video clip to 10 successive frames in our task. With the fixed length, the 3D CNN can be an efficient module for fast processing [42].

To avoid splitting an expression, we set an overlap ratio to 0.5, that the first clip of the first question $\{I^1_{1-i}\}^{10}_{i=1}$ is from 1 to 10 frames, and the second clip $\{I^1_{2-i}\}^{10}_{i=1}$ is from 6 to 15 frames, respectively. We use the superscript to denote the question from 1 to 20, and the subscript denotes the images in each clip. Therefore, for a N-frame video, there can be $M = 2 \times (N/10) - 1$ clips. For a 16 seconds video with 25fps in our dataset, we have 79 clips, and only very few of them are not neutral expressions and do not contribute to complementary information. We extract the 128-dimensional feature of each clip and denote the clip-wise representation of the first question as $\{f^1_m\}^M_{m=1}$, where $m \in \{1, \cdots, M\}$ index the clips within a question. We note that M can be different for different questions, and we omit the question index for simple notations.

After the clip-wise representations are extracted, the global attention is applied on top of them to extract the 128-dimensional question-wise features $\{a^i\}^{20}_{i=1}$ for 20 questions. To adaptively learn the significance of each clip within a question, we resort to a redundancy-aware self-attention scheme.

Then, the question-wise features $\{a^i\}^{20}_{i=1}$ are concatenated with the corresponding tabular questionnaire results and the

TABLE II
THE DETAILED STRUCTURE OF OUR 3D CNN MODULE FOR LOCAL
TEMPORAL FEATURE EXTRACTION.

Input Size	Type	Filter Shape
$[110 \times 110 \times 10]$ $[108 \times 108 \times (10 \times 16)]$	3D-Conv MaxPool	[16 kernels of $3 \times 3 \times 3$] $(2{\times}2{\times}1)$-$(1{\times}1{\times}1)$
$[54 \times 54 \times (10 \times 16)]$ $[52 \times 52 \times (10 \times 32)]$	3D-Conv MaxPool	[32 kernels of $3 \times 3 \times 3$] $(2{\times}2{\times}1)$-$(1{\times}1{\times}1)$
$[26 \times 26 \times (10 \times 32)]$ $[24 \times 24 \times (10 \times 64)]$	3D-Conv MaxPool	[64 kernels of $3 \times 3 \times 3$] $(2{\times}2{\times}1)$-$(1{\times}1{\times}1)$
$[12 \times 12 \times (10 \times 64)]$ $[10 \times 10 \times (10 \times 128)]$	3D-Conv MaxPool	[128 kernels of $3 \times 3 \times 3$] $(2{\times}2{\times}2)$-$(1{\times}1{\times}1)$
$[5 \times 5 \times (5 \times 128)]$	3D-Conv	[256 kernels of $5 \times 5 \times 5$]
$[1 \times 1 \times (1 \times 256)]$	Flatten	N/A
256-dim	FC	128-dim

Fig. 6. Illustration of the proposed redundancy-aware self-attention module.

answering time. We note that the SDS has 4 scales, and we use four-dimensional one-hot vector to encode the choice of each problem. The answering time is also concatenated as a one-dimensional scalar. The concatenated feature of each question can be a 133 dimension vector, and all of the 20 questions are concatenated together to form a 2660-dimensional questionnaire and video fused feature. We use the fully connected layers with the sigmoid output unit for binary classification, i.e., normal or depression.

We note that all of the 3D CNN and self-attention modules are shared for all of the clips and questions. In the following subsections, we provide the detailed framework of our hierarchically constructed 3D CNN, redundancy-aware attention, and question-level fusion modules.

C. 3D CNN for local temporal exploration

3D CNN has demonstrated its effectiveness of fast temporal representation extraction for relatively short fixed-length video [18]. We apply the standard 3D convolution operation to model the relationships between the successive frames. We illustrate the basic framework of 3D CNN in Fig. 5. After a few 3D Convolutional and maxpooling layers, we get a 256-dimensional feature vector, which is sent to a fully connected layer and result in a 128-dimensional clip-wise representation f_m^i, where i and m index the 20 questions and M clips within this question, respectively.

Considering the height and width dimensions of the video clips (e.g., 110) is usually much larger than the frame-wise dimension (e.g., 10), the first three maxpooling layers only half the height and width dimensions and denote as $(2{\times}2{\times}1)$-$(1{\times}1{\times}1)$. In the 4-th maxpooling layer, we half all of the three dimensions and denote as $(2{\times}2{\times}2)$-$(1{\times}1{\times}1)$. The detailed network structure is given in Tab. II.

D. Redundancy-aware self-attention

To explore the correlations between each clip, we resort to the affinity of clip-wise feature vectors. We use i to index the M clips within a question, and the i-th vector is regarded as a probe vector. j index the other $M - 1$ clips other than the i-th clip. We note that different questions can have a different number of fix-length clips, while the 3D CNN is not applicable [57], [42].
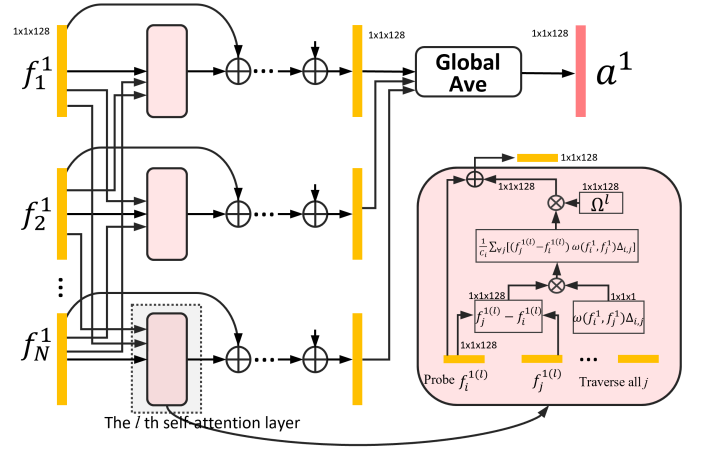
In our redundancy-aware self-attention module, we configure several self-attention blocks, which are indexed by $l \in \{1, 2 \cdots, L\}$. L is the total number of stacked sub-self attention blocks. In each self-attention block, we first traverse all of the clips to set the current clip as the probe. Then, we traverse the M-1 clips other than the probe to explore their correlations with the current probe clip.

Specifically, our self-attention block can be formulated as

$$f_i^{q(l)} = f_i^{q(l-1)} + \frac{\Omega^{(l)}}{C_i} \sum_{\forall j} \omega(f_i^{q(0)}, f_j^{q(0)})(f_j^{q(l-1)} - f_i^{q(l-1)})\Delta_{i,j}$$

$$C_i = \sum_{\forall j} \omega(f_i^{q(0)}, f_j^{q(0)})\Delta_{i,j}; l = 0, 1, \cdots, L \qquad (1)$$

where $\Omega^{(l)} \in \mathbb{R}^{1 \times 1 \times 128}$ is the weight vector to be learned and $f_i^{q(0)} = f_i^q$. The response is normalized by C_i. The pairwise affinity $\omega(\cdot, \cdot)$ is an scalar.

The operation of ω in Eq. (1) has many possible function candidates [19], [51]. We simply choose the embedded Gaussian given by

$$\omega(f_i^{q(0)}, f_j^{q(0)}) = e^{\psi(f_i^{q(0)})^T \phi(f_j^{q(0)})}, \qquad (2)$$

where $\psi(f_i^{q(0)}) = \Psi f_i^{q(0)}$ and $\phi(f_j^{q(0)}) = \Phi f_j^{q(0)}$ are two embedding functions, and $\Psi, \Phi \in \mathbb{R}^{128 \times 128}$ are the corresponding learnable mapping matrix [19].

The residual term, i.e., $f_j^{q(l-1)} - f_i^{q(l-1)}$, is the difference of the neighboring feature (i.e., $f_j^{q(l-1)}$) and the current probe feature $f_i^{q(l-1)}$. If $f_j^{q(l-1)}$ incorporates complementary information or more significant cues compared to the current probe feature $f_i^{q(l-1)}$, then our redundancy-aware attention scheme will eliminate the information from the inferior $f_i^{q(l-1)}$ and emphasis the more discriminative $f_j^{q(l-1)}$. Compared to the original non-local network which uses only $f_j^{q(l-1)}$ [19], our formulation can be more similar to the diffusion maps [58], graph Laplacian [59] and non-local image processing [60]. All of them are non-local analogues [61] of local diffusions, which are expected to be more stable than its original non-local counterpart [19] due to the nature of its inherit Hilbert-Schmidt
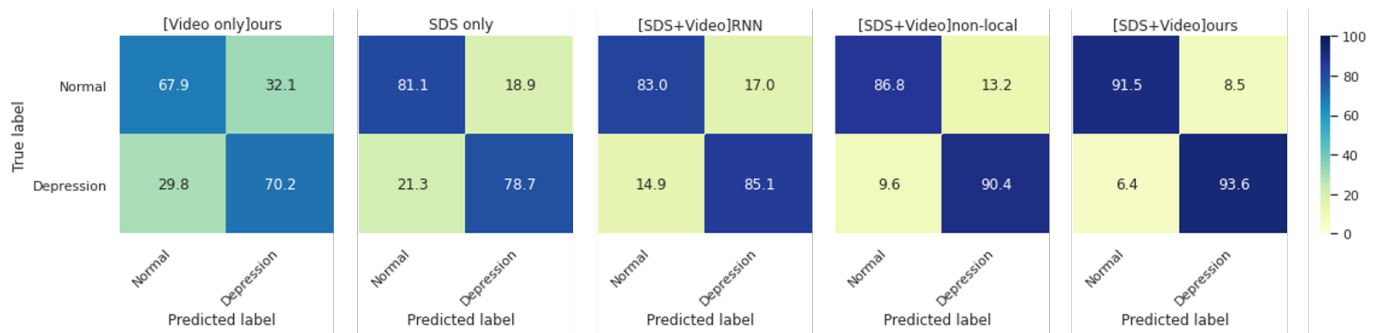
Fig. 7. The normalized confusion matrix of the different methods and modalities. From left to right, our hierarchical model using only video, sum operation with tabular SDS evaluation results, SDS and video with RNN, non-local and our hierarchical redundancy-aware self-attention.

operator [61]. We not that the the current probe feature will be added back in the last as the residual neural network. Therefore, the previous steps in the self-attention block is adjusting the information of each clip that need to be emphasised and translate to the later blocks.

Since the pair-wise residual term takes all possible clips into consideration and not following the sequential modeling, it does not suffer from long-term forgetting. Therefore, it can be an ideal choice for the attention modeling of many clips. In addition, the weighted average operation is able to take any number of inputs, which is fit for the different clip numbers in different questions. We note that the input and output of a self-attention block have the same size. Specifically, the M inputs with the size of $1 \times 1 \times 128$ will be processed to M outputs with the size of $1 \times 1 \times 128$.

Permutation-invariant is a special property of a self-attention scheme [39]. Since we use the sum operation in Eq. 1 to fuse the pair-wise residual terms. In the previous self-attention-based video analysis works, each frame is regarded as independent of each other and discards the sequential patterns [19]. Our video clips are inherently sequential data, and have an overlap with the neighboring clips. For exploiting the temporal patterns, the Gaussian kernel is proposed for sequential neighboring distance measure

$$\Delta_{i,j} = \exp(\frac{\parallel m_i - m_j \parallel_2^2}{\sigma}), \qquad (3)$$

where m_i, $m_j \in \mathbb{R}$ represent the position of i^{th} and j^{th} feature vectors in the video of a question, respectively. σ is a hyperparameter to control the shape of Gaussian Kernel.

After several concatenated self-attention blocks, the global pooling [19] is added on these M feature maps for element-wise averaging. The final output is a question-wise video feature $a^q \in \mathbb{R}^{1 \times 1 \times 128}$.

E. Question-level conditional fusion

The SDS questionnaire score of each question is denoted as $s^q \in \mathbb{R}^4$. Moreover, we empirically found that the answering time of each question can also be helpful for the diagnosis. Therefore, we also concatenate the video length of each question $t^q \in \mathbb{R}$. For a video that takes 3 seconds, we set its t^q to 3. The too short or too long answering maybe unreliable [17].

TABLE III
THE DETAILED FULLY-CONNECTED STRUCTURE OF OUR QUESTION-LEVEL FUSION.

Input Size	Type	Filter Shape
[128+4+1]×20	Concatenate	2660
2660	FC	1024 with ReLU
1024	FC	256 with ReLU
256	FC	1 with sigmoid

We concatenate a^q, s^q and t^q to form a 133 dimension feature vector for each question. Then, the 20 question-wise video and questionnaire features are concatenated to a 2660 dimensional feature for a final depression diagnosis. The fully connected layers are adopted, and the detailed network structure is given in Tab. III. The widely used Rectified Linear Unit (ReLU) [62] is used as the non-linear mapping function between each fully connected layer.

We use the sigmoid output unit $p = \frac{1}{1+e^{out}} \in (0,1)$ for binary classification, where out is the network prediction scalar in the last layer and will be normalized to a probability value, i.e., the likelihood of this subject is depression patient. We note that we indicate the normal subject with 0 (i.e., $y = 0$) and the depression subject with 1 (i.e., $y = 1$) according to the final clinician interview. To automatically train our model with backpropagation, we use the binary cross-entropy loss as the optimization objective.

$$\mathcal{L} = -y\log(p) - (1-y)\log(1-p), \qquad (4)$$

which has the zero loss if p matches its corresponding y. In addition, we simply set the threshold of binary prediction to $p = 0.5$ for testing.

IV. EXPERIMENTS

In this section, we compare the classification performance of our VoxelHop against 3D CNN-based classification. We also provide a systematic ablation study and sensitive analysis to demonstrate the effectiveness of the design choice of our framework.

A. Implementation Details and metrics

All the experiments were implemented using the widely adopted deep learning library Pytorch [63] on our server with

TABLE IV
COMPARISON OF THE BINARY CLASSIFICATION PERFORMANCE.

Methods	Accuracy	Sensitivity	Specificity
[SDS only]sum	0.800±0.000	0.787±0.000	0.811±0.000
[Video only]ours	0.690±0.006	0.660±0.009	0.720±0.007
[SDS+Video]RNN	0.840±0.007	0.816±0.007	0.846±0.006
[SDS+Video]non-local	0.885±0.005	0.859±0.003	0.893±0.007
[SDS+Video]ours	**0.925±0.004**	**0.907±0.006**	**0.924±0.005**

an NVIDIA V100 GPU, Xeon E5 v4 CPU with 128GB memory. Our model is trained with the Adam [64] optimizer with the hyper-parameters of β_1=0.9 and β_2=0.999. We used a batch size of 2 for our dataset. The networks of our framework and the compared methods are trained for 200 epochs for a fair comparison. We report the results of five random initialization, and provide the standard deviation ($\pm sd$) along with the average performance. We note that the 3D CNN and the redundancy-aware self-attention modules are shared for all of the questions, which can be processed parallelly. We empirically set $L = 5$. The training takes about 8 hours, while the average inference time of a subject only takes 1.3s. The threshold of binary classification testing is set to 0.5.

We adopt the five-fold cross-validation for the 200 subjects in our dataset. Specifically, we split the dataset into five subsets, and each has 40 subjects. We note that there are no overlap w.r.t. subjects between two folds. Then, we sequentially select a fold as our testing set (i.e., 40 participants), while the remaining four folds (i.e., 160 participants) are used for training.

For performance evaluation, we use the widely accepted binary classification metrics of accuracy, sensitivity (i.e., recall), specificity. More formally:

$$accuracy = \frac{TP + TN}{TP + TN + FP + FN}, \quad (5)$$

$$sensitivity = \frac{TP}{TP + FN}, \quad (6)$$

$$specificity = \frac{TN}{TN + FP}, \quad (7)$$

where TP, TN, FP, and FN indicate true positive, true negative, false positive, and false negative, respectively. We note that the positive and negative corresponding to the depression and normal, respectively.

In addition, by varying a threshold, the receiver-operating characteristic (ROC) curve plots the true positive rate (TPR) against the false positive rate (FPR). It demonstrates the diagnostic performance of a binary classification algorithm. The larger area under the curve (AUC), the better performance. We note that the TPR= $\frac{TN}{TN+FP}$ and FPR= $\frac{FP}{FP+TN}$.

B. Baselines and Comparison results

With our SDS and video dataset, there can be three choices of the modality, i.e., SDS only, video only, and both of them. With only the SDS evaluation result, we can simply add the score of 20 questions and using the threshold of 50 for normal and depression classification. According to the statistics in Tab. I, we can see that the SDS results can be different from the clinician diagnosis. We also tried only using the video
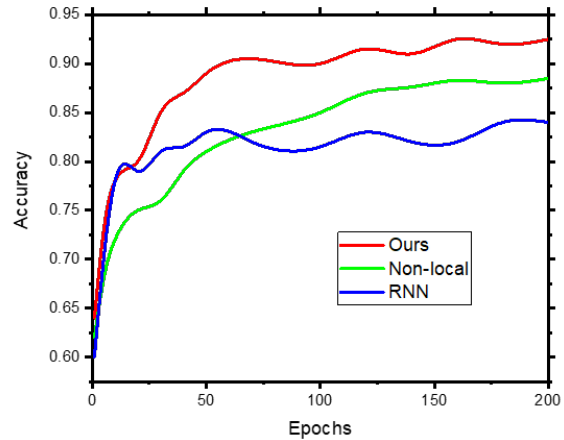

Fig. 8. The accuracy with respect to the number of epochs for different methods.

modality for classification, which did not concatenate s^q for the question-level fusion.

It is clear that using both SDS and video can outperform SDS only by a large margin, which evidenced the effectiveness of the additional video modality. The complementary information in the corresponding video is helpful to detect depression more accurately. It is also appealingly that even we only use the video, we are also able to predict the depression with an accuracy of 0.69, which is higher than the chance probability of 0.5.

To demonstrate the effectiveness of our model, we also applied two baseline methods, i.e., recurrent neural networks (RNN) [40], and non-local networks [19], for comparison. We note that this is also the first attempt to apply RNN and non-local attention for SDS and video analysis.

RNN is a typical choice for temporal modeling [65]. Specifically, we use the bi-directional LSTM [41] for our video feature extraction. Moreover, the non-local scheme [19] is recently proposed to address the long-term forgetting of RNN [66]. We use RNN or non-local to replace the 3D CNN and redundancy-aware self-attention in our framework to extract the 128-dimensional question-wise feature representation. The quantitative evaluation results are shown in Tab. IV. In addition, the accuracy of different epochs is plotted in Fig. 8. Our proposed framework achieves significantly better performance than the RNN and non-local counterparts.

C. Ablation study

We also provide the systematical ablation study for our framework modules.

• Our:color indicates using RGB frames as input, which without the gray-scale pre-processing. We note that we can simply modify the 3D CNN for multi-channel input, while the computation cost can be significantly increased.

• Our:non-local indicates using the conventional non-local [19] as the alternative of our redundancy-aware self-attention module. We note that the original non-local [19] is also first
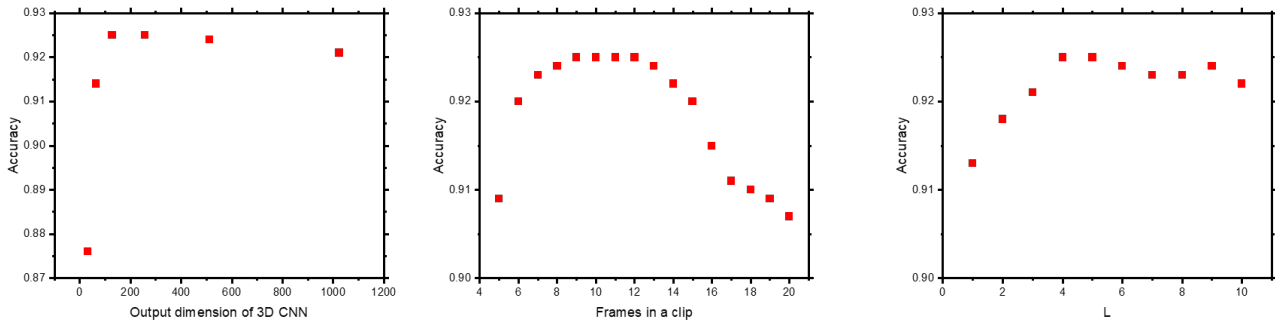
Fig. 9. Sensitivity study of using different output dimension of 3D CNN, length of each clip, and the number of self-attention blocks.

TABLE V
ABLATION STUDY OF THE DIFFERENT SETTINGS.

Methods	Accuracy	AUC
Ours	**0.925±0.004**	**0.918±0.005**
Ours:color	**0.925±0.006**	0.917±0.007
Ours:non-local	0.912±0.005	0.904±0.006
Ours:w/o Time	0.921±0.004	0.912±0.004
Ours:w/oΔ	0.922±0.006	0.913±0.006
Ours:f_j^{l-1}	0.920±0.005	0.911±0.005
Ours:$(\Psi\Phi)^L$	0.925±0.006	0.918±0.007
Ours+gender	0.924±0.007	0.918±0.006
Our:MLP	0.864±0.006	0.859±0.005
Our:SLF	0.905±0.005	0.896±0.007

TABLE VI
SENSITIVITY ANALYSIS OF TRAINING SAMPLES. WE USE DIFFERENT
TRAINING SAMPLES IN EACH CROSS-VALIDATION ROUND.

Methods	Accuracy	AUC
Ours (160 subjects)	**0.925±0.004**	**0.918±0.005**
Ours (120 subjects)	0.925±0.006	0.917±0.007
Ours (80 subjects)	0.912±0.005	0.904±0.006
Ours (40 subjects)	0.921±0.004	0.912±0.004

D. Sensitivity study

There are several hyperparameters in our framework. We provide a sensitivity analysis of these settings in this subsection and provide the analysis results in Fig. 9.

Specifically, we tested using different output dimensions of 3D CNN to explore the balance of the representative and computational cost. In Fig. 9, we can see that the output dimension of 3D CNN can be stable between 128 to 256 dimensions. The longer output feature can introduce significant additional costs for the subsequent redundancy-aware self-attention scheme. Moreover, since the subsequent redundancy-aware self-attention scheme does not change the length of the feature, the fully connected layers can be hard to process the longer inputs without enlarging its network structure.

The length of each clip can be related to the developing time of human expression and action in this task. The performance is not sensitive for a large range, e.g., 8 to 13 frames in a clip. The too-short clip may not able to incorporate an expression or action, while the longer clip can be hard to be effectively processed to extract useful information.

In addition, we can use different L to configure the number of redundancy-aware self-attention blocks. With five self-attention blocks, it is able to achieve the best performance in our task.

In Table VI, fix the 40 subjects in each cross-validation round, and reduce the training sample in each round to 40, 80, and 120 subjects. The performance can be better with more training data. We note that the difference between using 120 subjects or 160 subjects can be similar.

In Table VII, we compared the performance of using different head box sizes. The performance is not sensitive to the size within a relatively large range, while the 200% can be a good balance of the efficiency and performance.

introduced to depression detection in this paper, and is regarded as a baseline.

• Our:w/o time denotes the t^q is not concatenated for the question-level fusion.

• Our:f_j^{l-1} refers to using f_j^{l-1} instead of the difference term $f_j^{q(l-1)} - f_i^{q(l-1)}$ in Eq. (1). It does not explicitly consider the redundancy and lead to lower accuracy.

• Our:$(\Psi\Phi)^L$ indicates using the embedded Gaussian pairwise affinity for every block, which has similar performance but usually doubles the training time.

• Our:w/oΔ indicates excise $\Delta_{i,j}$ in Eq. (1) and does not taking the temporal sequence into consideration.

• Our+gender indicates that we add the 1/0 label of male/female along with the time.

• Our:MLP indicates that only using MLP to fuse the score and SDS time.

• Our:SLF indicates that using score-level fusion instead of feature level fusion.

The results are provided in Tab. V. The relatively inferior performance of the compared settings demonstrates the effectiveness of our choices. By adding the gender label, we do not achieve improvements w.r.t. the accuracy and AUC metrics. Our proposed method is able to explore the video information and achieve better performance than using MLP to fuse the score and SDS time. In addition, our feature-level fusion can outperforms the score-level fusion significantly.

TABLE VII
SENSITIVITY ANALYSIS OF TRAINING SAMPLES. WE USE DIFFERENT
TRAINING SAMPLES IN EACH CROSS-VALIDATION ROUND.

Methods	Accuracy	AUC
Ours (100%)	0.916±0.007	0.910±0.005
Ours (200%)	**0.925±0.004**	**0.918±0.007**
Ours (300%)	0.919±0.006	0.915±0.005
Ours (400%)	0.915±0.005	0.903±0.005

TABLE VIII
SENSITIVITY ANALYSIS OF TRAINING SAMPLES. WE USE DIFFERENT
TRAINING SAMPLES IN EACH CROSS-VALIDATION ROUND.

Methods	Accuracy	AUC
Ours (25fps)	**0.925±0.006**	**0.919±0.006**
Ours (20fps)	0.922±0.006	0.915±0.005
Ours (15fps)	0.913±0.007	0.910±0.004
Ours (10fps)	0.902±0.005	0.897±0.005

In Table VIII, we investigated the performance of using different fps of the video. We can see a significant performance drop w.r.t. both accuracy and AUC for the lower fps. Therefore, we chose the highest fps in our dataset. It can be promising to apply a higher frequency camera to capture the micro-expression information, while it can be costly in computation and memory.

V. DISCUSSIONS

A. Clinical prospects

Depression has been affecting more than 300 million people worldwide, while the early diagnosis can be immensely helpful for the treatment. The popularity of depression is even more significant in the COVID-19 season [4], which has long-term quarantine rules. A recent COVID-19 mental health survey [5] indicates that 23% of adults in Ireland reported suffering from depression[3]. However, the clinician interview can be prohibited or difficult considering the restrictions of avoiding infectious diseases, e.g., COVID-19. This further requires efficient self-administrative screening for depression detection.

The proposed framework has demonstrated good prediction accuracy for normal and depression, and has the potential for clinical practice in the future, especially for self-screening. With the Software-Defined Camera (SDC), we are able to transfer the questionnaire and its video to the back-end server for processing. Moreover, a similar protocol can be potentially applied to the smartphone APPs, which can easily take the face video of the user with the front camera. We note that the captured view can be different from our collected data and may result in the domain shift of the appearance. A possible solution to avoid the largen scale labeling of the mobile captured video is using the unsupervised domain adaptation to transfer the knowledge from our dataset to the unlabeled mobile dataset [67], [68], [69]. In addition, it is promising to apply a face pose invariant or robust feature extractor as [35], [70], [71], [72].

[3]https://www.maynoothuniversity.ie/newsevents/
covid19mentalhealthsurvey/maynoothuniversityandtrinitycollegefindshighratesanxiety

Therefore, the subsequent video-level aggregation modules can be shared across the datasets with different face poses. Actually, our pre-processing only crops a small area of the head, which is robust to the background changes. We are able to adjust the threshold for different applications with different sensitivity to the misclassification. Positive patients will then be referred to specialized clinics for a more comprehensive diagnosis.

The swift, automated deep learning system will partially substitute and support primary doctors' long-term clinical training, improving the primary diagnosis accuracy of depression in developing countries and laying the groundwork for early diagnosis and care of depression patients.

B. Limitations and future directions

Our system targets to explore the SDS evaluation and its video, while the clinician interview usually involves the round-based conversation. The spontaneous reaction and the speech (including text and phoneme) can provide more informative features. The interactive multi-round dialog system can be a promising direction.

In addition, we only collect the subject in China, especially in the Guangdong province, the population shift may affect the performance of our system. we will continuously collect more samples from different areas in the following study to achieve better performance.

Moreover, anxiety is also closely related to depression, while can have different treatment. In future work, we are also planning to incorporate anxiety into our diagnostic system.

VI. CONCLUSIONS

This study targets to automatically explore both the SDS evaluation and its question-wise video recording. By extending the face detector box, the facial expression, eye movement, and the actions of head-scratching and chin-touching are taken into account. A hierarchical end-to-end neural network framework is proposed to process the long-term variable-length video, which is also conditioned on the questionnaire results and the answering time. Based on the collected SDS and video recording dataset with an accurate clinician interview label, our model can make the diagnosis by fusing the information in tabular SDS results and video sequence. The 3D CNN module is able to efficiently explore the local temporal feature, while the novel redundancy-aware self-attention module can explicitly emphasizes the discriminative clips and reduces the redundancy based on feature pair-wise affinity. Our system exhibits appealing accuracy for depression detection, which can be promising for clinical practice in the future, especially in smartphones. The positive cases will then be referred to a specialized hospital for a final clinician diagnosis and care.

REFERENCES

[1] A. T. Beck, B. A. Alford, M. A. T. Beck, and P. D. B. A. Alford, *Depression*. University of Pennsylvania Press, 2014. 1

[2] A. T. Beck and B. A. Alford, *Depression: Causes and treatment*. University of Pennsylvania Press, 2009. 1

[3] E. Paykel, "The clinical interview for depression: development, reliability and validity," *Journal of affective disorders*, vol. 9, no. 1, pp. 85–96, 1985. 1, 4

[4] C. K. Ettman, S. M. Abdalla, G. H. Cohen, L. Sampson, P. M. Vivier, and S. Galea, "Prevalence of depression symptoms in us adults before and during the covid-19 pandemic," *JAMA network open*, vol. 3, no. 9, pp. e2 019 686–e2 019 686, 2020. 1, 9

[5] P. Hyland, M. Shevlin, O. McBride, J. Murphy, T. Karatzias, R. P. Bentall, A. Martinez, and F. Vallières, "Anxiety and depression in the republic of ireland during the covid-19 pandemic," *Acta Psychiatrica Scandinavica*, vol. 142, no. 3, pp. 249–256, 2020. 1, 9

[6] W. W. Zung, "A self-rating depression scale," *Archives of general psychiatry*, vol. 12, no. 1, pp. 63–70, 1965. 1

[7] W. W. Zung, C. B. Richards, and M. J. Short, "Self-rating depression scale in an outpatient clinic: further validation of the sds," *Archives of general psychiatry*, vol. 13, no. 6, pp. 508–515, 1965. 1, 2, 4

[8] J. T. Biggs, L. T. Wylie, and V. E. Ziegler, "Validity of the zung self-rating depression scale," *The British Journal of Psychiatry*, vol. 132, no. 4, pp. 381–385, 1978. 1

[9] J. B. Gabrys and K. Peters, "Reliability, discriminant and predictive validity of the zung self-rating depression scale," *Psychological Reports*, vol. 57, no. 3_suppl, pp. 1091–1096, 1985. 1, 4

[10] W. W. Zung, "Factors influencing the self-rating depression scale," *Archives of general psychiatry*, vol. 16, no. 5, pp. 543–547, 1967. 1

[11] J. B. Williams, "A structured interview guide for the hamilton depression rating scale," *Archives of general psychiatry*, vol. 45, no. 8, pp. 742–747, 1988. 1, 4

[12] C. A. Corneanu, M. O. Simón, J. F. Cohn, and S. E. Guerrero, "Survey on rgb, 3d, thermal, and multimodal approaches for facial expression recognition: History, trends, and affect-related applications," *IEEE transactions on pattern analysis and machine intelligence*, vol. 38, no. 8, pp. 1548–1568, 2016. 2

[13] A. Pampouchidou, P. G. Simos, K. Marias, F. Meriaudeau, F. Yang, M. Pediaditis, and M. Tsiknakis, "Automatic assessment of depression based on visual cues: A systematic review," *IEEE Transactions on Affective Computing*, vol. 10, no. 4, pp. 445–470, 2017. 2, 3

[14] J. Zhu, Z. Wang, T. Gong, S. Zeng, X. Li, B. Hu, J. Li, S. Sun, and L. Zhang, "An improved classification model for depression detection using eeg and eye tracking data," *IEEE transactions on nanobioscience*, vol. 19, no. 3, pp. 527–537, 2020. 2

[15] D. R. Rubinow and R. M. Post, "Impaired recognition of affect in facial expression in depressed patients," *Biological psychiatry*, vol. 31, no. 9, pp. 947–953, 1992. 2

[16] R. Krause, E. Steimer, C. Sänger-Alt, and G. Wagner, "Facial expression of schizophrenic patients and their interaction partners," *Psychiatry*, vol. 52, no. 1, pp. 1–12, 1989. 2

[17] P. M. Lewinsohn and G. E. Atwood, "Depression: A clinical-research approach." *Psychotherapy: Theory, Research & Practice*, vol. 6, no. 3, p. 166, 1969. 2, 6

[18] S. Ji, W. Xu, M. Yang, and K. Yu, "3D convolutional neural networks for human action recognition," *IEEE transactions on pattern analysis and machine intelligence*, vol. 35, no. 1, pp. 221–231, 2012. 2, 5

[19] X. Wang, R. Girshick, A. Gupta, and K. He, "Non-local neural networks," in *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 2, 3, 5, 6, 7

[20] S. Scherer, G. Stratou, M. Mahmoud, J. Boberg, J. Gratch, A. Rizzo, and L.-P. Morency, "Automatic behavior descriptors for psychological disorder analysis," in *2013 10th IEEE International Conference and Workshops on Automatic Face and Gesture Recognition (FG)*. IEEE, 2013, pp. 1–8. 2

[21] J. Joshi, "An automated framework for depression analysis," in *2013 Humaine Association Conference on Affective Computing and Intelligent Interaction*. IEEE, 2013, pp. 630–635. 2

[22] J. Gratch, R. Artstein, G. M. Lucas, G. Stratou, S. Scherer, A. Nazarian, R. Wood, J. Boberg, D. DeVault, S. Marsella *et al.*, "The distress analysis interview corpus of human and computer interviews." in *LREC*, 2014, pp. 3123–3128. 3

[23] M. Valstar, B. Schuller, K. Smith, F. Eyben, B. Jiang, S. Bilakhia, S. Schnieder, R. Cowie, and M. Pantic, "Avec 2013: the continuous audio/visual emotion and depression recognition challenge," in *Proceed-ings of the 3rd ACM international workshop on Audio/visual emotion challenge*, 2013, pp. 3–10. 3

[24] J. W. Williams Jr, P. H. Noël, J. A. Cordes, G. Ramirez, and M. Pignone, "Is this patient clinically depressed?" *Jama*, vol. 287, no. 9, pp. 1160–1170, 2002. 3

[25] H. Dinkel, M. Wu, and K. Yu, "Text-based depression detection on sparse data," *arXiv e-prints*, pp. arXiv–1904, 2019. 3

[26] 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," *Journal of affective disorders*, vol. 114, no. 1-3, pp. 163–173, 2009. 3

[27] A. Haque, M. Guo, A. S. Miner, and L. Fei-Fei, "Measuring depression symptom severity from spoken language and 3d facial expressions," *arXiv preprint arXiv:1811.08592*, 2018. 3

[28] H. Dinkel, P. Zhang, M. Wu, and K. Yu, "Depa: Self-supervised audio embedding for depression detection," *arXiv preprint arXiv:1910.13028*, 2019. 3

[29] M. Muzammel, H. Salam, Y. Hoffmann, M. Chetouani, and A. Othmani, "Audvowelconsnet: A phoneme-level based deep cnn architecture for clinical depression diagnosis," *Machine Learning with Applications*, vol. 2, p. 100005, 2020. 3

[30] X. Zhang, J. Shen, Z. ud Din, J. Liu, G. Wang, and B. Hu, "Multimodal depression detection: fusion of electroencephalography and paralinguistic behaviors using a novel strategy for classifier ensemble," *IEEE journal of biomedical and health informatics*, vol. 23, no. 6, pp. 2265–2275, 2019. 3

[31] J. Shen, X. Zhang, X. Huang, M. Wu, J. Gao, D. Lu, Z. Ding, and B. Hu, "An optimal channel selection for eeg-based depression detection via kernel-target alignment," *IEEE Journal of Biomedical and Health Informatics*, 2020. 3

[32] S. Zheng, C. Lei, T. Wang, C. Wu, J. Sun, and H. Peng, "Feature-level fusion for depression recognition based on fnirs data," in *2020 IEEE International Conference on Bioinformatics and Biomedicine (BIBM)*. IEEE, 2020, pp. 2906–2913. 3

[33] J. F. Roddick and M. Spiliopoulou, "A survey of temporal knowledge discovery paradigms and methods," *IEEE Transactions on Knowledge and data engineering*, vol. 14, no. 4, pp. 750–767, 2002. 3

[34] X. Liu, L. Jin, X. Han, J. Lu, J. You, and L. Kong, "Identity-aware facial expression recognition in compressed video," *ICPR*, 2020. 3

[35] X. Liu, L. Jin, X. Han, and J. You, "Mutual information regularized identity-aware facial expression recognition in compressed video," *Pattern Recognition*, p. 108105, 2021. 3, 9

[36] G. He, X. Liu, F. Fan, and J. You, "Classification-aware semi-supervised domain adaptation," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, 2020, pp. 964–965. 3

[37] Z. Zhou, Y. Huang, W. Wang, L. Wang, and T. Tan, "See the forest for the trees: Joint spatial and temporal recurrent neural networks for video-based person re-identification," in *Computer Vision and Pattern Recognition (CVPR), 2017 IEEE Conference on*. IEEE, 2017, pp. 6776–6785. 3

[38] J. Wang, X. Liu, F. Wang, L. Zheng, F. Gao, H. Zhang, X. Zhang, W. Xie, and B. Wang, "Automated interpretation of congenital heart disease from multi-view echocardiograms," *Medical Image Analysis*, vol. 69, p. 101942, 2021. 3

[39] X. Liu, Z. Guo, S. Li, P. Jia, J. You, and K. B.V.K, "Permutation-invariant feature restructuring for correlation-aware image set-based recognition," *ICCV 2019*. 3, 6

[40] M. Schuster and K. K. Paliwal, "Bidirectional recurrent neural networks," *IEEE Transactions on Signal Processing*, vol. 45, no. 11, pp. 2673–2681, 1997. 3, 7

[41] A. Ullah, J. Ahmad, K. Muhammad, M. Sajjad, and S. W. Baik, "Action recognition in video sequences using deep bi-directional lstm with cnn features," *IEEE Access*, vol. 6, pp. 1155–1166, 2018. 3, 7

[42] X. Liu, Z. Guo, J. You, and B. V. Kumar, "Dependency-aware attention control for image set-based face recognition," *IEEE Transactions on Information Forensics and Security*, vol. 15, pp. 1501–1512, 2019. 3, 4, 5

[43] J. Gao and R. Nevatia, "Revisiting temporal modeling for video-based person reid," *arXiv preprint arXiv:1805.02104*, 2018. 3

[44] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention is all you need," in *Advances in Neural Information Processing Systems*, 2017, pp. 5998–6008. 3

[45] P. Shaw, J. Uszkoreit, and A. Vaswani, "Self-attention with relative position representations," *arXiv preprint arXiv:1803.02155*, 2018. 3

[46] A. Buades, B. Coll, and J.-M. Morel, "A non-local algorithm for image denoising," in *Computer Vision and Pattern Recognition, 2005. CVPR 2005. IEEE Computer Society Conference on*, vol. 2.   IEEE, 2005, pp. 60–65. 3

[47] P. Battaglia, R. Pascanu, M. Lai, D. J. Rezende *et al.*, "Interaction networks for learning about objects, relations and physics," in *Advances in neural information processing systems*, 2016, pp. 4502–4510. 3

[48] Y. Hoshen, "Vain: Attentional multi-agent predictive modeling," in *Advances in Neural Information Processing Systems*, 2017, pp. 2701–2711. 3

[49] N. Watters, D. Zoran, T. Weber, P. Battaglia, R. Pascanu, and A. Tacchetti, "Visual interaction networks: Learning a physics simulator from video," in *Advances in Neural Information Processing Systems*, 2017, pp. 4539–4547. 3

[50] F. S. Y. Yang, L. Zhang, T. Xiang, P. H. Torr, and T. M. Hospedales, "Learning to compare: Relation network for few-shot learning," 2018. 3

[51] B. Zhou, A. Andonian, and A. Torralba, "Temporal relational reasoning in videos," *In ECCV*, 2018. 3, 5

[52] S. Zhang, X. Zhu, Z. Lei, H. Shi, X. Wang, and S. Z. Li, "Faceboxes: A cpu real-time face detector with high accuracy," in *2017 IEEE International Joint Conference on Biometrics (IJCB)*.   IEEE, 2017, pp. 1–9. 4

[53] L. L. Carpenter and L. H. Price, "Psychotic depression: what is it and how should we treat it?" *Harvard review of psychiatry*, vol. 8, no. 1, pp. 40–42, 2000. 4

[54] A. E. Kazdin, R. B. Sherick, K. Esveldt-Dawson, and M. D. Rancurello, "Nonverbal behavior and childhood depression," *Journal of the American Academy of Child Psychiatry*, vol. 24, no. 3, pp. 303–309, 1985. 4

[55] L. R. Derogatis and K. L. Savitz, "The scl-90-r, brief symptom inventory, and matching clinical rating scales." 1999. 4

[56] R. O. Jegede, "Psychometric attributes of the self-rating anxiety scale," *Psychological Reports*, vol. 40, no. 1, pp. 303–306, 1977. 4

[57] X. Liu, K. B.V.K, C. Yang, Q. Tang, and J. You, "Dependency-aware attention control for unconstrained face recognition with image sets," in *European Conference on Computer Vision*, 2018. 5

[58] Y. Tao, Q. Sun, Q. Du, and W. Liu, "Nonlocal neural networks, nonlocal diffusion and nonlocal modeling," *arXiv preprint arXiv:1806.00681*, 2018. 5

[59] F. R. Chung and F. C. Graham, *Spectral graph theory*.   American Mathematical Soc., 1997, no. 92. 5

[60] G. Gilboa and S. Osher, "Nonlocal linear image regularization and supervised segmentation," *Multiscale Modeling & Simulation*, vol. 6, no. 2, pp. 595–630, 2007. 5

[61] Q. Du, M. Gunzburger, R. B. Lehoucq, and K. Zhou, "Analysis and approximation of nonlocal diffusion problems with volume constraints," *SIAM review*, vol. 54, no. 4, pp. 667–696, 2012. 5, 6

[62] K. Hara, D. Saito, and H. Shouno, "Analysis of function of rectified linear unit used in deep learning," in *2015 International Joint Conference on Neural Networks (IJCNN)*, 2015, pp. 1–8. 6

[63] A. Paszke, S. Gross, S. Chintala, G. Chanan, E. Yang, Z. DeVito, Z. Lin, A. Desmaison, L. Antiga, and A. Lerer, "Automatic differentiation in pytorch," 2017. 6

[64] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014. 7

[65] R. Pascanu, C. Gulcehre, K. Cho, and Y. Bengio, "How to construct deep recurrent neural networks," 2014. 7

[66] R. Pascanu, T. Mikolov, and Y. Bengio, "On the difficulty of training recurrent neural networks," in *International conference on machine learning*.   PMLR, 2013, pp. 1310–1318. 7

[67] X. Liu, X. Liu, B. Hu, W. Ji, F. Xing, J. Lu, J. You, C.-C. J. Kuo, G. E. Fakhri, and J. Woo, "Subtype-aware unsupervised domain adaptation for medical diagnosis," *AAAI*, 2021. 9

[68] X. Liu, B. Hu, X. Liu, J. Lu, J. You, and L. Kong, "Energy-constrained self-training for unsupervised domain adaptation," in *2020 25th International Conference on Pattern Recognition (ICPR)*.   IEEE, 2021, pp. 7515–7520. 9

[69] G. He, X. Liu, F. Fan, and J. You, "Image2audio: Facilitating semi-supervised audio emotion recognition with facial expression image," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, 2020, pp. 912–913. 9

[70] X. Liu, Y. Chao, J. J. You, C.-C. J. Kuo, and B. Vijayakumar, "Mutual information regularized feature-level frankenstein for discriminative recognition," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2021. 9

[71] X. Liu, S. Li, L. Kong, W. Xie, P. Jia, J. You, and B. Kumar, "Feature-level frankenstein: Eliminating variations for discriminative recognition," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2019, pp. 637–646. 9

[72] X. Liu, "Disentanglement for discriminative visual recognition," *Recognition and Perception of Images: Fundamentals and Applications*, pp. 143–187, 2021. 9