

Sequential Fusion of Facial Appearance and Dynamics for Depression Recognition

Qian Chen¹, Iti Chaturvedi², Shaoxiong Ji³, and Erik Cambria^{*1}

¹*Nanyang Technological University, Singapore, 639798, Singapore*

²*James Cook University, Townsville, 4814, Australia*

³*Aalto University, Espoo, 02150, Finland*

**Corresponding author. Email: cambria@ntu.edu.sg*

Abstract

In mental health assessment, it is validated that nonverbal cues like facial expressions can be indicative of the depressive disorder. Recently, the multimodal fusion of facial appearance and dynamics based on convolutional neural networks has demonstrated encouraging performance in depression analysis. However, correlation and complementarity between different visual modalities have not been well studied in prior methods. In this paper, we propose a sequential fusion method for facial depression recognition. For mining the correlated and complementary depression patterns in multimodal learning, a chained-fusion mechanism is introduced to jointly learn facial appearance and dynamics in a unified framework. We show that such sequential fusion can provide a probabilistic perspective of the model correlation and complementarity between two different data modalities for improved depression recognition. Results on a benchmark dataset show the superiority of our method against several state-of-the-art alternatives.

Keywords— Depression recognition, facial representation, convolutional neural network, multimodal learning, sequential fusion.

1 Introduction

Major depressive disorder (MDD) is a psychological disorder that exhibits feelings of sadness, loss, or anger that may impact a person’s usual social activities. At a global level, over 300 million people of all ages suffer from different levels of depression, equivalent to 4.4% of the world’s population [37]. A depressive episode can be classified into a mild, moderate, or severe level, depending on the symptoms. Mild depression may bring difficulties in continuing with ordinary work and social activities. More seriously, the feeling of depression may occur comorbidity with self-mutilation [39], and depressed people are more likely to commit suicide than the general population [29]. Early detection of depressive or other mental disorders provides a possible way for mental intervention [23].

In clinical practice, the diagnosis procedure for MDD can usually be labor-intensive and highly relies on expertise observations. Due to the increasing number of people suffering from MDD around the world, methods for automated depression analysis appear to be urgent for objective and efficient diagnosis. Recently, automated depression diagnosis based on computer vision techniques has drawn increasing attention [46], and the significance of the verbal cues for depression analysis has been demonstrated in various depression detection/recognition tasks [53, 9, 8, 33, 55, 31]. Besides, visual cues like facial expression and facial dynamics have also proven to be effective in depression analysis [17, 52, 1, 60]. This paper investigates facial depression recognition, aiming to predict the depression level for a given face video based on the BDI-II metric [3].

While encouraging progress has been made over the past few years, automated depression analysis in videos remains challenging due to the following reasons:

On one hand, unlike those large-scale image datasets (e.g., ImageNet [12]) for visual recognition [26], the size of most existing depression datasets (e.g., AVEC 2014 [46]) is relatively small due to the privacy concerns. While representation learning based on convolutional neural network (CNN) has been proved to be more effective than hand-crafted descriptors in visual-based depression recognition [60], the lack of labeled data makes the model training with deep networks prone to over-fitting in practice.

On the other hand, many learning methods in the literature have been devoted to multimodal fusion of audio and/or video features for depression recognition [60, 33, 1], which have demonstrated boosted recognition

performance by exploiting the complementary information encoded in different modalities. However, essential correlation and diversity between different visual modalities have not been well investigated in previous visual-based methods, especially for multimodal fusion of visual cues with deep CNN architecture.

In this paper, we propose a multimodal deep learning approach for facial depression recognition to address these issues. Specifically, a sequential fusion of facial appearance and dynamics is introduced to facilitate such multimodal representation learning. Here, facial appearance and dynamics are adopted as the basis modalities in our multimodal fusion framework, as they have been validated to be effective in visual-based depression diagnosis [60]. Hence, a fusion between the two modalities is first operated on the blocks of a two-stream CNN. By such fusion of mid-level features in the CNN training, an initial interaction is conducted to optimize the complementary patterns. Then, the extracted feature, together with the predicted label from the first stream (e.g., RGB modality), is fed into the second stream (e.g., Optical Flow modality) to refine the final prediction. We show that such sequential fusion can provide a probabilistic perspective about model correlation and complementarity between two different data modalities for improved depression recognition. We conduct experiments on the benchmark dataset (AVEC2014), and the results show the superiority of our method against several state-of-the-art alternatives. The main contributions of this paper are:

- We proposed a sequential chained-fusion approach for depression recognition. With a probabilistic perspective, the proposed approach models the correlation and complementarity between facial appearance and facial dynamics at several network layers, such that the complementary and correlation information of different visual cues extracted from videos could be well exploited in model learning (Section 3).
- We evaluate our approach on the benchmark dataset and empirically show improvement over several state-of-the-art alternatives (Section 4).

The rest of this paper is organized as follows. We first give a brief review of related work in Section 2, and then detail the proposed sequential fusion method in Section 3. Experimental settings, results and discussions are presented in Section 4, and Section 5 concludes the paper.

2 Related work

This section briefly reviewed two related topics: 1) visual-based depression recognition and 2) multimodal learning with deep architectures.

Automated Depression Recognition. Depression analysis based on various behavioral signals has drawn increasing attention in the affective computing community. Such feasible signals include the vision- and speech-based cues of human communication. While several works in the literature focus on this research topic, we are interested in the visual-based approaches for depression recognition. In the AVEC 2013 competition[47], a facial descriptor named the local phase quantization (LPQ) [35] was used as a baseline for facial depression recognition, where the extracted LPQ features for each video frame are further employed to train a support vector regression (SVR). In [9], Cummins et al. used the pyramid of histogram of gradients (PHOG) [4] and the space-time interest points (STIPs) [27] for extraction of behavioural cues for depression analysis. Meng et al. [33] proposed to use motion history histogram (MHH) [34] feature to model motion in videos, and then use the partial least squares (PLS) [10] for training regression model. In [52], the motion cue is encoded by the LPQ-TOP feature extracted from sub-volumes of the cropped facial regions, by which the behavior pattern dictionary can be obtained based on sparse coding.

In the AVEC 2014 competition [46], the local motion descriptor LGBP-TOP [2] and the SVR were employed as the baseline video description and prediction model, respectively. In [38], various local motion features extracted from sub-volumes of the detected faces were used for training an SVR-based prediction model. Jan et al. [22] proposed a 1D MHH based on some local descriptors to train a PLS regressor for final prediction. In [24], the baseline LGBP-TOP combined with LPQ was used as the video descriptor for depression prediction.

The aforementioned depression recognition methods proposed to use hand-crafted descriptors, which are generally less effective to model and reveal high-level semantic cues. Recently, depression feature learning based on deep CNNs achieves considerable progress. For example, Zhu et al. [60] proposed jointly learning the facial appearance and dynamics based on a two-stream CNN, in which two different features are fused at a fully-connected (FC) layer. Improved performance reported in their experiments indicated the efficacy of such a simple fusion manner. Most recently, Uddin et al. [45] introduced a new two-stream network for deep spatiotemporal feature learning, in which spatial information is extracted by a ResNet network, and they used a volume local directional number (VLDN) based feature descriptor to model facial motions. Zhou et al. [58] proposed a deep network named DepressNet to learn facial depression features with visual explanation, such that facial salient regions with different depression levels can be detected by using the generated activation mapping. Later, Zhou et al. [59] proposed to jointly learn the feature embedding and label distribution to address the issue of deep representation learning on a limited amount of labeled depression data, and the improvement by such learning scheme was reported in their experiments in comparison several state-of-the-art alternatives. Besides deep CNNs, there are several methods for image feature learning such as binary descriptor [13].

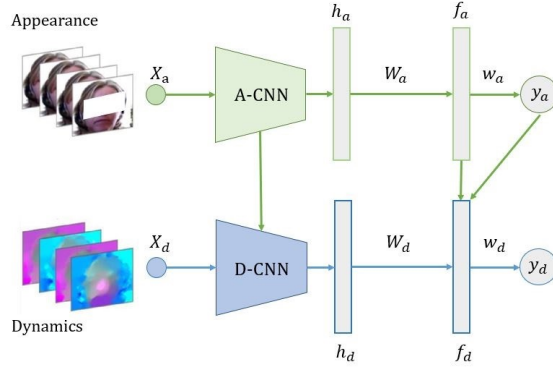


Figure 1: Overview of our proposed sequential fusion approach for facial depression recognition. The (mid- and high-level) features extracted from the first stream (RGB modality) together with the predicted label are fed into the second stream (dynamics modality) for refinement of the final prediction.

Multimodal Deep Representation Learning. In the multimodal setting, visual data consists of multiple input modalities [6, 7, 48, 5], and each one may have a different representation and structure. Intuitively, useful representations could be learned from such multimodal data by fusing them into a joint representation to characterize the real-world semantic concept that the visual data corresponds to. In practice, however, it is much more difficult to model and discover the nonlinear correlation and diversity across modalities than those among features in the same modality. Recently, a good number of multimodal deep learning methods have been proposed to better exploit useful information from different modalities for more robust visual analysis [48, 49, 14, 28, 60, 61, 45][54, 21, 30, 16, 15]. In [43], Srivastava et al. proposed a multimodal learning method with a deep Boltzmann machine (DBM) to jointly learn multimodal feature representations. They approached this by adding a concatenated layer that connects DBMs from different modalities. In [54, 21], multi-metric learning methods were applied to exploited multi features of human faces. Yan et. al. [54] proposed a jointly learning method to capture multiple features from multiple distance metrics. In [21], Hu et. al. further studied multi-view metric learning in a sharable and individual way, which shown the superiority of sharing features from different view. In [14], a two-stream CNN with an additional multimodal fusion layer was proposed for RGB-D object recognition. Motivated by the observation that the data from different modalities may contain modal-specific patterns as well as common patterns, Wang et al. [48] proposed a shareable and specific multimodal feature learning framework for RGB-D object recognition. Li et. al. [30] proposed a global-local framework to extract pose, appearance and motion features for RGB-D gesture recognition. By imposing the representation learning of associations between different modalities, Zolfaghari et al. [61] designed a chained multi-stream network to fully exploit the pose, motion, and appearance cues for action classification and detection. Feichtenhofer et. al. [16] proposed a two-stream CNN architecture to fuse a spatial and temporal network at a convolutional layer instead of at the softmax layer, which boosted performance on action recognition problem with a substantial saving in parameters. To further boost performance of action classification and detection, Feichtenhofer et. al. [15] presented SlowFast Networks for video recognition which was inspired by biological studies on the retinal ganglion cells in the primate visual system.

Our proposed method is related to the multimodal deep learning methods for visual recognition, which aimed to fuse multiple responses from the CNNs trained with different modalities of visual data. However, our work differs from them in the fusion manner. Our work aims at deep fusion between features from different modalities in a sequential manner, while only a simple fusion at the FC layer was performed in [45] as well as [60]. Most similar to ours is the work [61] which also performed fusion in a chained manner; however, the fusion of mid-level features in the CNN training was not involved in their approach, and hence our sequential fusion represents a more *deep* fusion mechanism that can be more suitable for multimodal representation learning with limited labeled depression data, as indicated in our experiments.

3 Our Approach

Predicting the severity of facial depression is a process of learning spatio-temporal features related to human emotion categorization [50] from face videos. The facial appearance of a subject is one of the important visual patterns for depression recognition. At the same time, facial dynamics characterized by optical flow captures the temporal variations of appearance across frames. As such, we propose a sequential fusion approach to investigate the correlation and complementarity between two different data modalities for depression analysis. As shown in Fig. 1, we use a two-stream network architecture, where the encoders for each stream have the

same backbone structure (can be any off-the-shelf CNNs).

Fundamentally, depression estimation can be viewed as a regression task, and hence we employ the mean square error (MSE) as the loss function. Mathematically, the loss L is defined as:

$$L = \frac{1}{2N} \sum_{i=1}^N \|y_i - \hat{y}_i\|^2, \quad (1)$$

where N is the number of the samples, \hat{y}_i is the output prediction of the second stream of our network for the i th sample, and y_i is its ground-truth label.

3.1 Appearance-CNN

Convolutional Neural Networks (CNNs) have been proved with powerful capability on image classification tasks over large-scale image data. Conversely, CNN is not a proper option to capture features from the dataset with a limited size. For depression estimation, available datasets are usually with limited data and subjects. To handle this issue, we employ the pre-training and fine-tuning strategies to train the facial appearance CNN (A-CNN).

Due to time restrictions or computational restraints, it's not always possible to build a deep model from scratch which is the reason why we use pre-trained model. To achieve facial representations, we train two pre-trained deep networks (e.g., GoogLeNet [44] and ResNet-50 [18]) over CASIA-WebFace database [56], which is a public face recognition database containing 494,414 facial images from 10,575 subjects.

After the pre-training step, we can obtain the general facial features through the pre-trained model while those features are not relevant to facial depression. Hence, depression data are fed into the pre-trained model for fine-tuning, such that the final model is capable of accurate depression estimation.

3.2 Dynamics-CNN

Along with facial appearance, facial dynamics is also an indispensable component in our proposed model. The dynamics-CNN (D-CNN) is built upon the same backbone as the A-CNN with the optical flow data. Unlike the static RGB data, facial dynamics model the motion patterns inherent in faces that can be highly indicative for visual depression analysis.

We compute the optical flow with the duality-based approach [57], which is a decent method with sufficiently fast speed. To feed the optical flow data into the D-CNN, we transform the optical flow signal into a 3-dimensional image data (x, y, m) , where x , y and m represent the x -component, y -component and the magnitude of the flow, respectively. For better observation and calculation, we multiply each image by 16 and convert it to the closest integer between 0 and 255 [51].

Similarly, we employ the same architecture, configurations, and loss used in the appearance CNN to train the dynamics CNN. The difference lies in that the input to the dynamics stream is the *flow* image computed from video frames, and there is no pre-training step in D-CNN.

3.3 Sequential Fusion

Since Simonyan et al. [40] proposed a simple fusion framework with two separate convolutional networks on raw images and optical flow, respectively, multi-stream architectures entered the public consciousness and became a popular approach to handle multimedia tasks.

The baseline fusions are illustrated in Fig. 2, where the solution (a) applies the normal concatenation to form the input features of a set of fully-connected layers, which combines features from two streams directly, and the two CNNs are independent. In such a case, the learning manners of the two streams exert a slight influence on each other. Furthermore, as the input to the dynamics, CNN is the optical flow between several sequential frames; the changes between continuous frames are minor and difficult to distinguish. In order to achieve a compact and discriminative multimodal deep representation, a proper fusion method is needed to gather features from facial appearance and dynamics.

To integrate the two individual deep networks mentioned above, we introduce a sequential fusion approach based on a two-stream architecture. By making use of the features from both modalities, a Markov chain is established to integrate the two streams, which may refine the depression prediction sequentially. Considering different modality as the main input, a refined prediction is achieved by combining the hidden features (e.g., high-level features) and the predictions from the previous stream (see Fig. 2(b)-(c)). However, such fusion is made only after the FC layers, which means the *mid-level features* before the FC layers have not been exploited for better fusion. With the limited promotion of the feature fusion, the refinement of the predictions is improved, though the improvement is not significant.

Unlike the baseline fusion schemes shown in Fig. 2, our proposed sequential fusion (shown in Fig. 1) perform feature fusion not only after the FC layer, but also on the blocks of CNNs, by which both high-level and

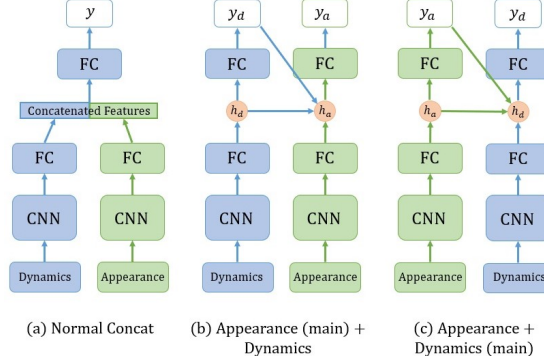


Figure 2: Different fusion baselines for facial depression recognition. (a) Normally concatenated features from the dynamics and appearance streams; (b) The features extracted from the dynamics stream together with the predicted label are fed to the appearance stream for final prediction; (c) The features extracted from the appearance stream together with the predicted label are fed to the dynamics stream for final prediction.

mid-level features of the CNNs can be well exploited to model the correlation and complementarity between different data modalities. In what follows, we give a probabilistic interpretation for such a fusion mechanism.

For different input modalities, we assume that the depression predictions are conditionally independent. Consequently, we can factorize the joint probability into the conditional probabilities according to the conditional independence property. In a Markov chain, we predict the outputs $Y = \{y_1, y_2, \dots, y_S\}$ on the given sequence of inputs $I = \{I_1, I_2, \dots, I_S\}$ with $P(y|I)$ maximized. Due to the Markov property, we have

$$P(y|I) = P(y_1|I) \prod_{s=2}^S P(y_s|I, y_1, \dots, y_{s-1}) \quad (2)$$

To model the likelihood in (2), the hidden feature and the prediction probability are respectively denoted by

$$h_s = f([h_{s-1}, \text{CNN}(I_s), (y_1, y_2, \dots, y_{s-1})]), \quad (3)$$

$$P(y_s|I, y_{<s}) = \mathcal{N}(w_s^T h_s; \bar{y}_s, \sigma^2)$$

where $s \in \{1, 2, \dots, S\}$, f is an activation function, h_{s-1} denotes the hidden feature from the previous stream, y_s denotes the prediction of the s th stream, w_s denotes the regression coefficient vector, \bar{y}_s is the ground-truth label, and σ is a certain standard deviation of the Gaussian \mathcal{N} . Here, $\text{CNN}(\cdot)$ denotes the convolutionary part and the first FC layer of the network, which can be any off-the-shelf backbones (e.g., VGG and ResNet).

In the proposed approach, the prediction of the dynamics stream is made conditioned on the appearance stream as well as the input dynamics data, which means that the final prediction is effected not only by the input of the current stream but also by the features and predictions from previous streams (see Fig.1). When the input modalities is $I = \{I_a, I_d\}$, we have:

$$h_a = \text{CNN}(I_a), \quad (4)$$

$$P(y_a|I) = \mathcal{N}(w_a^T h_a; \bar{y}_a, \sigma^2)$$

and

$$h_d = f([h_a, \text{CNN}(I_d), y_a]), \quad (5)$$

$$P(y_d|I, y_a) = \mathcal{N}(w_d^T h_d; \bar{y}_d, \sigma^2)$$

As known to all, maximization of $P(y_d|I, y_a)$ is equivalent to the minimization of the MSE loss defined in Eq. (1). Hence, our sequential fusion mechanism has a clear probabilistic interpretation.

It should be noticed that depressed patients are usually slow to initiate actions with stiff facial expressions. For depression analysis, motion features could be more discriminative than appearance cues. It is observed in our experiment that using the dynamics stream as the main stream to fuse the appearance stream can perform better prediction, and it also facilitates the training of the two-stream network. In the inference stage, only the prediction of the mainstream is used as the final depression prediction.

4 Experiments

To validate the effectiveness of our depression recognition approach, we conduct experiments on the AVEC 2014 benchmark dataset and compare its performance with several state-of-the-art algorithms as well as the

baselines. In what follows, a description of the dataset, data pre-processing, and experimental setting are first presented. Then, we present the results and analysis.

4.1 Dataset

We conduct the experiments on a database of the Audio/Visual Emotion Challenge (AVEC) 2014 [46], which is the most widely-used depression sub-challenge database for depression recognition. In AVEC dataset, to evaluate the severity, a depression level score is measured by a self-reported 21 multiple-choice inventory–Beck Depression Index (BDI) [36]. The BDI scores range from 0 to 63, where the lower score represents more mild symptoms. The score evaluated in [0, 13], [14, 19], [20, 28] and [29, 63] indicate minimal, mild, moderate and severe depression, respectively. For each video clip, there are 3~5 annotators predicting the BDI score.

In AVEC 2014, there are 84 subjects, and each subject needs to perform two different tasks named “Northwind” and “FreeForm” according to the instructions. All subjects in the two tasks speak German. There are 150 videos for each task, and the recordings were equally split into three partitions: training, development, and test set. Each partition includes 50 videos and has similar distributions in terms of gender, age, and depression levels for the partitions. All videos are recorded by webcam in a human-computer interaction scenario, and each video is approximately 2-minute length on average. There are at least three annotators per clip, and most clips are annotated by 5.

4.2 Data Pre-processing

As the raw data have a certain degree of noisy and redundant information which is irrelevant to depressive expressions. To extract meaningful information from noise, it is necessary to apply multiple pre-processing steps on the raw data before feeding it to the model. To avoid the waste of computing resources and speed up the training, subsampling is performed on the video frames with an interval of 10 frames which is determined experimentally.

To deal with the raw data, face detection and landmark localization of each subject in the video are implemented by Dlib [25]. After that, the facial region of an image size of 256×256 is cropped and aligned according to the eye locations.

After the above steps, we compute optical flow over a sequence of facial regions extracted from each video clip. As the input to dynamics stream, optical flow is computed between two frames which can capture facial motions known as face “flow images”. A “flow image” has three components (x, y, m) , where the first two channels are x , and y flow values and the third channel is the magnitude of the optical flow normalized between 0 and 255 with a median of 128.

4.3 Experimental Setting

We use two popular network architectures, i.e., VGG-11 [41] and ResNet-50 [18], as the backbone of our two-stream network to train the appearance and dynamic CNNs. In our experiments, the appearance CNN is pre-trained on a large-scale face dataset CASIA WebFace [56], while the dynamics CNN is trained from scratch. The MSE is adopted as the loss function for our depression regression.

The total number of training iterations for the appearance and dynamics CNNs are 400,000 and 600,000, respectively. We set the initial learning rate to 0.001, and decrease the learning rate by polynomial decay with power equals to 0.5, and set the momentum to 0.9 with a weight decay of 0.0002. For the joint training, the total number of iterations is 200,000 with an initial learning rate of 0.0001.

We use the mean absolute error (MAE) and root mean square error (RMSE) to measure the overall recognition performance. They are defined by $MAE = \frac{1}{N} \sum_{i=1}^N |y_i - \hat{y}_i|$ and $RMSE = \sqrt{\frac{1}{N} \sum_{i=1}^N (y_i - \hat{y}_i)^2}$, where N is the number of data samples, y_i and \hat{y}_i are the ground truth and the prediction for the i th sample.

4.4 Experimental Results

4.4.1 Performance of Individual Models

We first investigate the impact of different fusion models (see Fig. 2) in our fusion framework for depression analysis. Six baselines are defined for evaluation of the performance of individual models: 1) Appearance CNN, 2) Dynamics CNN, 3) Normal Concat: fusion by normal concatenation on the FC layer of the two-stream CNN, 4) Appearance (main)+Dynamics: Dynamics stream is fed to Appearance stream for sequential fusion, 5) Dynamics (main)+Appearance: Appearance stream is fed to Dynamics stream for sequential fusion, 6) DeepFusion: Appearance stream is fed to Dynamics stream for a sequential fusion of both mid-level and high-level features.

As shown in Table 1, our proposed DeepFusion achieves the MAE/RMSE of 6.16/8.13, which consistently outperforms the other baselines. Specifically, two sequential fusion models (dynamics CNN as the mainstream)

Table 1: Ablation study of our approach on AVEC 2014

Model	MAE	RMSE
Appearance CNN	6.60	8.88
Dynamics CNN	8.64	10.71
Normal Concat	6.72	8.68
Appearance(main)+Dynamics	6.71	8.58
Appearance+Dynamics(main)	6.41	8.70
DeepFusion (proposed)	6.16	8.13

Table 2: Depression prediction results with different backbones on AVEC 2014

Model	MSE	RMSE
Appearance CNN(VGG)	8.19	10.34
Dynamics CNN(VGG)	9.54	11.49
Normal Concat(VGG)	9.54	11.50
DeepFusion(VGG)	7.54	9.79
Appearance CNN(ResNet-50)	6.60	8.88
Dynamics CNN(ResNet-50)	8.64	10.71
Normal Concat(ResNet-50)	6.72	8.68
DeepFusion(ResNet-50)	6.16	8.13

perform better than other individuals, indicating the efficacy of such fusion mechanism. On the other hand, we can see that appearance CNN as the mainstream for sequential fusion may not be suitable for depression analysis, as the motion cues may play a more important role in features fusion for final prediction. Finally, the proposed DeepFusion performs better than Dynamics (main)+Appearance, indicating that integration of mid-level features in the sequential fusion framework can further boost the overall prediction performance.

In addition, we investigate the impact of different backbones in our proposed approach. As shown in Fig. 2, prediction models built upon ResNet-50 [18] consistently perform better than those built upon VGG-11 [41]. Also, our proposed DeepFusion achieves the best performance on both backbones in terms of MAE and RMSE.

4.4.2 Comparison with Previous Methods

We compared our proposed approach with several state-of-the-art depression recognition methods, and the results are presented in Table 3. For a fair comparison, the compared eight methods are based on the visual modality (e.g., face videos). Specifically, six of them are also deep learning-based solutions, and the other two are shallow learning models with hand-crafted video descriptors.

In [46], the baseline model for AVEC 2014 employed the epsilon-SVR with intersection kernel [32] trained using LGBP-TOP features. On the AVEC 2014 dataset, our approach beats the baseline approach by dropping the MAE by 2.70 and RMSE by 2.83. In [20], an SVR trained with the dynamic feature descriptors MRLBP-TOP and DPFV achieved the MAE/RMSE of 7.21/9.01. Also, our approach outperforms this shallow learning-based solution by a significant margin.

When compared to other seven deep learning-based works [15], [11, 45, 58, 59, 42, 19], our approach achieves the second-best performance on the AVEC 2014 dataset. In [15], the SlowFast networks transferred from action recognition model achieved the MAE/RMSE of 6.78/8.40. In [11], the combination of the global and local Convolutional 3D networks achieved the MAE/RMSE of 6.59/8.31. In [45], the MAE/RMSE of 6.86/8.78 was achieved by utilizing Bi-LSTM, whose input is the output of a deep CNN and TMP. In a very recent work [19], both local and global attention CNN are introduced for depression recognition and reduced the MAE/RMSE to 6.51/8.30. In [58], deep depression representation with visual explanation achieved the MAE/RMSE of 6.60/8.88, and later in [59], the deep metric learning-based solution achieved the MAE/RMSE of 6.59/8.30. Our approach consistently outperforms the aforementioned deep learning-based methods in terms of MAE and RMSE. The best-performed solution among the compared methods is the spectral representation of behavior primitives [42], which achieved the MAE/RMSE of 5.95/7.15.

5 Conclusion

In this paper, we have proposed a deep multimodal learning method for the representation fusion of facial appearance and dynamics. To model the correlated and complementary depression patterns in multimodal learning, a chained-fusion mechanism is introduced to jointly learn facial appearance and dynamics in a unified

Table 3: Comparison with previous methods on AVEC 2014

Methods	MAE	RMSE
LGBP-TOP+SVR [46] (2014)	8.86	10.86
MRLBP-TOP+DPFV+SVR [20] (2018)	7.21	9.01
SlowFast Networks [15] (2019)	6.78	8.40
C3D(Global+Local) [11] (2019)	6.59	8.31
VLDN+Bi-LSTM+TMP [45] (2020)	6.86	8.78
DepressNet [58] (2020)	6.60	8.88
DJ-LDML [59] (2020)	6.59	8.30
Spectral-Representation [42] (2020)	5.95	7.15
DLGA-CNN [19] (2021)	6.51	8.30
DeepFusion (proposed)	<u>6.16</u>	<u>8.13</u>

framework. We have shown that such sequential fusion provides a clear probabilistic perspective of the model correlation and complementarity between two different data modalities for improved depression recognition. Experimental results on a benchmark dataset demonstrated the efficacy of our method when compared to several state-of-the-art alternatives. In future work, investigation of the private-share model for multimodal depression representation learning appears to be an interesting topic.

References

- [1] Sharifa Alghowinem, Roland Goecke, Michael Wagner, Julien Epps, Matthew Hyett, Gordon Parker, and Michael Breakspear. Multimodal depression detection: fusion analysis of paralinguistic, head pose and eye gaze behaviors. *IEEE Transactions on Affective Computing*, 9(4):478–490, 2016.
- [2] Timur R Almaev and Michel F Valstar. Local gabor binary patterns from three orthogonal planes for automatic facial expression recognition. In *2013 Humaine association conference on affective computing and intelligent interaction*, pages 356–361. IEEE, 2013.
- [3] Aaron T Beck, Robert A Steer, Roberta Ball, and William F Ranieri. Comparison of beck depression inventories-ia and-ii in psychiatric outpatients. *Journal of personality assessment*, 67(3):588–597, 1996.
- [4] Anna Bosch, Andrew Zisserman, and Xavier Munoz. Representing shape with a spatial pyramid kernel. In *Proceedings of the 6th ACM international conference on Image and video retrieval*, pages 401–408, 2007.
- [5] Erik Cambria, Newton Howard, Jane Hsu, and Amir Hussain. Sentic blending: Scalable multimodal fusion for continuous interpretation of semantics and sentics. In *IEEE SSCI*, pages 108–117, Singapore, 2013.
- [6] Iti Chaturvedi, Ranjan Satapathy, Sandro Cavallari, and Erik Cambria. Fuzzy commonsense reasoning for multimodal sentiment analysis. *Pattern Recognition Letters*, 125(264-270), 2019.
- [7] Pawas Chhokra, Anurag Chowdhury, Gaurav Goswami, Mayank Vatsa, and Richa Singh. Unconstrained kinect video face database. *Information Fusion*, 44:113–125, 2018.
- [8] Jeffrey F Cohn, Tomas Simon Kruez, Iain Matthews, Ying Yang, Minh Hoai Nguyen, Margara Tejera Padilla, Feng Zhou, and Fernando De la Torre. Detecting depression from facial actions and vocal prosody. In *2009 3rd International Conference on Affective Computing and Intelligent Interaction and Workshops*, pages 1–7. IEEE, 2009.
- [9] Nicholas Cummins, Jyoti Joshi, Abhinav Dhall, Vidhyasaharan Sethu, Roland Goecke, and Julien Epps. Diagnosis of depression by behavioural signals: a multimodal approach. In *Proceedings of the 3rd ACM international workshop on Audio/visual emotion challenge*, pages 11–20, 2013.
- [10] Sijmen De Jong. Simpls: an alternative approach to partial least squares regression. *Chemometrics and intelligent laboratory systems*, 18(3):251–263, 1993.
- [11] Wheidima Carneiro de Melo, Eric Granger, and Abdenour Hadid. Combining global and local convolutional 3d networks for detecting depression from facial expressions. In *2019 14th IEEE International Conference on Automatic Face & Gesture Recognition (FG 2019)*, pages 1–8. IEEE, 2019.

- [12] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009.
- [13] Yueqi Duan, Jiwen Lu, Ziwei Wang, Jianjiang Feng, and Jie Zhou. Learning deep binary descriptor with multi-quantization. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1183–1192, 2017.
- [14] Andreas Eitel, Jost Tobias Springenberg, Luciano Spinello, Martin Riedmiller, and Wolfram Burgard. Multimodal deep learning for robust rgb-d object recognition. In *2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 681–687. IEEE, 2015.
- [15] Christoph Feichtenhofer, Haoqi Fan, Jitendra Malik, and Kaiming He. Slowfast networks for video recognition. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6202–6211, 2019.
- [16] Christoph Feichtenhofer, Axel Pinz, and Andrew Zisserman. Convolutional two-stream network fusion for video action recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1933–1941, 2016.
- [17] Jeffrey M Girard, Jeffrey F Cohn, Mohammad H Mahoor, S Mohammad Mavadati, Zakia Hammal, and Dean P Rosenwald. Nonverbal social withdrawal in depression: Evidence from manual and automatic analyses. *Image and vision computing*, 32(10):641–647, 2014.
- [18] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [19] Lang He, Jonathan Cheung-Wai Chan, and Zhongmin Wang. Automatic depression recognition using cnn with attention mechanism from videos. *Neurocomputing*, 422:165–175, 2021.
- [20] Lang He, Dongmei Jiang, and Hichem Sahli. Automatic depression analysis using dynamic facial appearance descriptor and dirichlet process fisher encoding. *IEEE Transactions on Multimedia*, 21(6):1476–1486, 2018.
- [21] Junlin Hu, Jiwen Lu, and Yap-Peng Tan. Sharable and individual multi-view metric learning. *IEEE transactions on pattern analysis and machine intelligence*, 40(9):2281–2288, 2017.
- [22] Asim Jan, Hongying Meng, Yona Falinie A Gaus, Fan Zhang, and Saeed Turabzadeh. Automatic depression scale prediction using facial expression dynamics and regression. In *Proceedings of the 4th International Workshop on Audio/Visual Emotion Challenge*, pages 73–80, 2014.
- [23] Shaoxiong Ji, Xue Li, Zi Huang, and Erik Cambria. Suicidal ideation and mental disorder detection with attentive relation networks. *Neural Computing and Applications*, 2021.
- [24] Heysem Kaya, Fazilet Çilli, and Albert Ali Salah. Ensemble cca for continuous emotion prediction. In *Proceedings of the 4th International Workshop on Audio/Visual Emotion Challenge*, pages 19–26, 2014.
- [25] Davis E King. Dlib-ml: A machine learning toolkit. *The Journal of Machine Learning Research*, 10:1755–1758, 2009.
- [26] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. *Advances in neural information processing systems*, 25:1097–1105, 2012.
- [27] Ivan Laptev, Marcin Marszalek, Cordelia Schmid, and Benjamin Rozenfeld. Learning realistic human actions from movies. In *2008 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–8. IEEE, 2008.
- [28] Ian Lenz, Honglak Lee, and Ashutosh Saxena. Deep learning for detecting robotic grasps. *The International Journal of Robotics Research*, 34(4-5):705–724, 2015.
- [29] Jean-Pierre Lépine and Mike Briley. The increasing burden of depression. *Neuropsychiatric disease and treatment*, 7(Suppl 1):3, 2011.
- [30] Benchao Li, Wanhua Li, Yongyi Tang, Jian-Fang Hu, and Wei-Shi Zheng. Gl-pam rgb-d gesture recognition. In *2018 25th IEEE International Conference on Image Processing (ICIP)*, pages 3109–3113. IEEE, 2018.

- [31] Xingchen Ma, Hongyu Yang, Qiang Chen, Di Huang, and Yunhong Wang. Depaudionet: An efficient deep model for audio based depression classification. In *Proceedings of the 6th international workshop on audio/visual emotion challenge*, pages 35–42, 2016.
- [32] Subhansu Maji, Alexander C Berg, and Jitendra Malik. Classification using intersection kernel support vector machines is efficient. In *2008 IEEE conference on computer vision and pattern recognition*, pages 1–8. IEEE, 2008.
- [33] Hongying Meng, Di Huang, Heng Wang, Hongyu Yang, Mohammed Ai-Shuraifi, and Yunhong Wang. Depression recognition based on dynamic facial and vocal expression features using partial least square regression. In *Proceedings of the 3rd ACM international workshop on Audio/visual emotion challenge*, pages 21–30, 2013.
- [34] Hongying Meng and Nick Pears. Descriptive temporal template features for visual motion recognition. *Pattern Recognition Letters*, 30(12):1049–1058, 2009.
- [35] Timo Ojala, Matti Pietikainen, and Topi Maenpaa. Multiresolution gray-scale and rotation invariant texture classification with local binary patterns. *IEEE Transactions on pattern analysis and machine intelligence*, 24(7):971–987, 2002.
- [36] Jason T Olin, Lon S Schneider, Elaine M Eaton, Mary F Zemansky, and Vicki E Pollock. The geriatric depression scale and the beck depression inventory as screening instruments in an older adult outpatient population. *Psychological Assessment*, 4(2):190, 1992.
- [37] World Health Organization et al. Depression and other common mental disorders: global health estimates. Technical report, World Health Organization, 2017.
- [38] Humberto Pérez Espinosa, Hugo Jair Escalante, Luis Villaseñor-Pineda, Manuel Montes-y Gómez, David Pinto-Avedaño, and Veronica Reyez-Meza. Fusing affective dimensions and audio-visual features from segmented video for depression recognition: Inaoe-buap’s participation at avec’14 challenge. In *Proceedings of the 4th International Workshop on Audio/Visual Emotion Challenge*, pages 49–55, 2014.
- [39] Shana Ross and Nancy Heath. A study of the frequency of self-mutilation in a community sample of adolescents. *Journal of youth and Adolescence*, 31(1):67–77, 2002.
- [40] Karen Simonyan and Andrew Zisserman. Two-stream convolutional networks for action recognition in videos. *Advances in neural information processing systems*, 27:568–576, 2014.
- [41] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- [42] Siyang Song, Shashank Jaiswal, Linlin Shen, and Michel Valstar. Spectral representation of behaviour primitives for depression analysis. *IEEE Transactions on Affective Computing*, 2020.
- [43] Nitish Srivastava, Ruslan Salakhutdinov, et al. Multimodal learning with deep boltzmann machines. In *NIPS*, volume 1, page 2. Citeseer, 2012.
- [44] Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. Going deeper with convolutions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1–9, 2015.
- [45] Md Azher Uddin, Joolekha Bibi Joolee, and Young-Koo Lee. Depression level prediction using deep spatiotemporal features and multilayer bi-lstm. *IEEE Transactions on Affective Computing*, 2020.
- [46] Michel Valstar, Björn Schuller, Kirsty Smith, Timur Almaev, Florian Eyben, Jarek Krajewski, Roddy Cowie, and Maja Pantic. Avec 2014: 3d dimensional affect and depression recognition challenge. In *Proceedings of the 4th international workshop on audio/visual emotion challenge*, pages 3–10, 2014.
- [47] Michel Valstar, Björn Schuller, Kirsty Smith, Florian Eyben, Bihan Jiang, Sanjay Bilakhia, Sebastian Schnieder, Roddy Cowie, and Maja Pantic. Avec 2013: the continuous audio/visual emotion and depression recognition challenge. In *Proceedings of the 3rd ACM international workshop on Audio/visual emotion challenge*, pages 3–10, 2013.
- [48] Anran Wang, Jianfei Cai, Jiwen Lu, and Tat-Jen Cham. Mmss: Multi-modal sharable and specific feature learning for rgb-d object recognition. In *Proceedings of the IEEE international conference on computer vision*, pages 1125–1133, 2015.

- [49] Anran Wang, Jiwen Lu, Jianfei Cai, Tat-Jen Cham, and Gang Wang. Large-margin multi-modal deep learning for rgb-d object recognition. *IEEE Transactions on Multimedia*, 17(11):1887–1898, 2015.
- [50] Zhaoxia Wang, Sengbeng Ho, and Erik Cambria. A review of emotion sensing: Categorization models and algorithms. *Multimedia Tools and Applications*, 79:35553–35582, 2020.
- [51] Philippe Weinzaepfel, Zaid Harchaoui, and Cordelia Schmid. Learning to track for spatio-temporal action localization. In *Proceedings of the IEEE international conference on computer vision*, pages 3164–3172, 2015.
- [52] Lingyun Wen, Xin Li, Guodong Guo, and Yu Zhu. Automated depression diagnosis based on facial dynamic analysis and sparse coding. *IEEE Transactions on Information Forensics and Security*, 10(7):1432–1441, 2015.
- [53] James R Williamson, Thomas F Quatieri, Brian S Helfer, Rachelle Horwitz, Bea Yu, and Daryush D Mehta. Vocal biomarkers of depression based on motor incoordination. In *Proceedings of the 3rd ACM international workshop on Audio/visual emotion challenge*, pages 41–48, 2013.
- [54] Haibin Yan, Jiwen Lu, Weihong Deng, and Xiuzhuang Zhou. Discriminative multimetric learning for kinship verification. *IEEE Transactions on Information forensics and security*, 9(7):1169–1178, 2014.
- [55] Ying Yang, Catherine Fairbairn, and Jeffrey F Cohn. Detecting depression severity from vocal prosody. *IEEE transactions on affective computing*, 4(2):142–150, 2012.
- [56] Dong Yi, Zhen Lei, Shengcai Liao, and Stan Z Li. Learning face representation from scratch. *arXiv preprint arXiv:1411.7923*, 2014.
- [57] Christopher Zach, Thomas Pock, and Horst Bischof. A duality based approach for realtime tv-l 1 optical flow. In *Joint pattern recognition symposium*, pages 214–223. Springer, 2007.
- [58] X. Zhou, K. Jin, Y. Shang, and G. Guo. Visually interpretable representation learning for depression recognition from facial images. *IEEE Transactions on Affective Computing*, 11(3):542–552, 2020.
- [59] X. Zhou, Z. Wei, M. Xu, S. Qu, and G. Guo. Facial depression recognition by deep joint label distribution and metric learning. *IEEE Transactions on Affective Computing*, pages 1–1, 2020.
- [60] Yu Zhu, Yuanyuan Shang, Zhuhong Shao, and Guodong Guo. Automated depression diagnosis based on deep networks to encode facial appearance and dynamics. *IEEE Transactions on Affective Computing*, 9(4):578–584, 2017.
- [61] Mohammadreza Zolfaghari, Gabriel L Oliveira, Nima Sedaghat, and Thomas Brox. Chained multi-stream networks exploiting pose, motion, and appearance for action classification and detection. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2904–2913, 2017.