

Improving state-of-the-art in Detecting Student Engagement with Resnet and TCN Hybrid Network

Ali Abedi and Shehroz S. Khan
KITE, Toronto Rehabilitation Institute
University Health Network
Toronto, Canada
ali.abedi@uhn.ca, shehroz.khan@uhn.ca

Abstract—Automatic detection of students’ engagement in online learning settings is a key element to improve the quality of learning and to deliver personalized learning materials to them. Varying levels of engagement exhibited by students in an online classroom is an affective behavior that takes place over space and time. Therefore, we formulate detecting levels of students’ engagement from videos as a spatio-temporal classification problem. In this paper, we present a novel end-to-end Residual Network (ResNet) and Temporal Convolutional Network (TCN) hybrid neural network architecture for students’ engagement level detection in videos. The 2D ResNet extracts spatial features from consecutive video frames, and the TCN analyzes the temporal changes in video frames to detect the level of engagement. The spatial and temporal arms of the hybrid network are jointly trained on raw video frames of a large publicly available students’ engagement detection dataset, DAiSEE. We compared our method with several competing students’ engagement detection methods on this dataset. The ResNet+TCN architecture outperforms all other studied methods, improves the state-of-the-art engagement level detection accuracy, and sets a new baseline for future research.

Keywords—engagement detection; spatio-temporal; Temporal Convolutional Network; Residual Neural Network

I. INTRODUCTION

With the widespread availability and adoption of internet services across major urban centers and universities, online education, telehealth, telemedicine, and telerehabilitation are becoming more ubiquitous and mainstream. In situations such as the COVID-19 pandemic with strict social distancing guidelines, these online services made it possible for students to complete their courses [1] and patients to receive the necessary care and health support [2]. Online services offer many advantages compared to the traditional in-person settings, in terms of being more accessible, economical, and customizable [1]. However, these online services also bring other types of challenges. For instance, in online classroom setting, students and tutor are behind a ‘virtual wall’, it becomes very difficult for the tutor to assess the students engagement in the class being taught [3]. This problem is further exacerbated if the group of students is large [4]. Therefore, from a tutor’s perspective, it is important to automatically detect the level of engagement

among students to provide them real-time feedback and take necessary actions to engage the students to maximize their learning objectives.

Various modalities have been utilized to automatically detect students’ engagement, including students’ images [5], videos [6], [7], audio [6], [8], and Electrocardiogram (ECG) [9]. Video cameras/webcams are mostly used in student learning environment; thus, they have been extensively used in assessing students’ engagement in online classroom setting. Video cameras and webcams offer a cheaper, ubiquitous and unobtrusive alternative to other sensing modalities. Therefore, majority of the recent works on student engagement detection are based on the visual data of students acquired by cameras and using computer-vision techniques [9], [10].

The computer-vision based approaches for student engagement are categorized into image-based and video-based approaches. The former approaches detect engagement based on single images [11], or single frames extracted from the videos [5]. A major limitation of this approach is that it only utilizes spatial information from single frames, whereas engagement detection is a spatio-temporal affective behaviour [12]. Another challenge with frame based approaches is that annotation is needed for each frame, which is an arduous task in practice. The latter approach is to detect students’ engagement from videos instead of using single frames. In this case, one label is needed after each video segment. Less annotations is required in this case; however, the classification problem is more challenging due to the coarse labeling.

In traditional settings, in the video-based approaches, handcrafted features, such as eye gaze and head pose, can be extracted and classification algorithms can be trained to detect the level of engagement [7], [13], [14], [15], [16]. More recently, end-to-end video-based approaches have been proposed to detect student engagement, in which consecutive raw frames of video are fed to variants of Convolutional Neural Networks (CNNs) to detect the level of engagement [7], [13], [14], [15], [16].

Sinatra et al. [17] defined students’ engagement from person-oriented perspective: engagement is the cognitive,

affective, and motivational *states* of the student at the moment of learning and are best captured with fine-grained physiological and behavioral measures (e.g., facial expressions, actions, body postures). D’Mello et al. [12] implied that students’ engagement is not stable over time, and should be assessed at fine-grained time scales ranging from seconds to a few minutes. For this reason, we consider students’ engagement as a spatio-temporal data (video) analysis problem. To detect the level of engagement from video data, in addition to the *state* of the student in each video frame, the changes in the *state* of the student over consecutive frames need to be analyzed. Therefore, for detecting students’ engagement in an end-to-end setting, utilizing sequential neural networks in conjunction with CNNs can be beneficial. With this motivation, we propose a novel end-to-end Residual Network (ResNet) [18] and Temporal Convolutional Network (TCN) [19] hybrid neural network architecture (Resnet+TCN) for detecting students’ engagement level from videos. The ResNet extracts spatial features from consecutive frames, and the TCN analyzes the temporal changes in consecutive frames to detect the level of engagement. We evaluate the performance of the Resnet+TCN architecture on a publicly available dataset, the Dataset for the Affective States in E-Environments (DAiSEE) [7]. The DAiSEE dataset contains 9,068 10-second videos captured from 112 students in online classroom setting, and annotated by the engagement level of students. Our results on the DAiSEE dataset outperformed other competing methods and improved the results on the state-of-the-art methods by 3.9% in engagement level detection accuracy.

II. LITERATURE REVIEW

Over the recent years, extensive research efforts have been devoted to study student engagement detection using computer vision and deep learning techniques [9], [10], [6]. We review recent works on the computer-vision based student’s engagement detection with a focus on the works that performed their experiments on the DAiSEE dataset (described in Section. IV).

Two types of approaches were found prominently in computer-vision based engagement detection problem, i.e., feature-based models and end-to-end models.

In the **feature-based** engagement detection approaches, firstly, multi-modal handcrafted features are extracted from video/image and then fed to a classifier or regressor to detect the level of engagement in video/image [10], [6], [20], [21]. Wu et al. in [20] proposed a feature-based approach for student’s engagement level detection in EmotiW dataset [6]. They extracted facial and upper-body features from videos and classified the features using a combination of Long Short-Term Memory (LSTM) and Gated Recurrent Unit (GRU) to detect the level of engagement. Zhu et al. in [21] proposed an attention-based GRU model to classify hand-crafted face and body features from videos and

detect the level of engagement in the EmotiW dataset [6]. Whitehill et al. in [5] proposed different feature-extraction (box filters and Gabor features) and classification (SVM and GentleBoost) combinations to detect the level of engagement of students from single images in their dataset.

Huang et al. [13] proposed Deep Engagement Recognition Network (DERN) which combines bidirectional LSTM and attention mechanism to classify extracted features from faces and detect the level of engagement. They achieved 60% engagement level detection accuracy on the DAiSEE dataset. Wang et al. [14] proposed a CNN architecture to classify facial landmarks and features extracted from faces to detect the level of engagement with an accuracy of 57% on the DAiSEE dataset.

In the **end-to-end** approaches, the raw frames of videos or images are fed to a deep CNN classifier/regressor to detect engagement. Gupta et al. [7] introduced the DAiSEE dataset and established benchmark results using different end-to-end convolutional video-classification techniques, including the InceptionNet [18], C3D [22], and Long-term Recurrent Convolutional Networks (LRCN) [23], achieving 46.4%, 56.1%, and 57.9% accuracy, respectively.

Geng et al. [15] utilized C3D classifier along with the focal loss to classify the level of engagement in the DAiSEE dataset, and achieved 56.2% accuracy. Zhang et al. [16] proposed a modified version of the Inflated 3D (I3D) along with the weighted cross entropy loss to classify the level of engagement in the DAiSEE dataset, and achieved 52.35% accuracy. In addition to the original four-level engagement classification problem in the DAiSEE, Zhang et al. [16] considered the engagement detection in the DAiSEE as a two-class classification problem. They changed the labels from low and very-low levels of engagement to not-engaged and from high and very-high levels of engagement to engaged. They trained the I3D model to solve the binary, not-engaged/engaged, classification problem, and obtained 98.82% accuracy. The above two loss functions, the focal loss in [15] and the weighted cross entropy loss in [16], are used to tackle the problem of imbalanced data distribution in the DAiSEE dataset.

Liao et al. [24] proposed Deep Facial Spatio-Temporal Network (DFSTN) for students’ engagement detection in online learning. Their model for engagement detection contains two modules, a pretrained SE-ResNet-50 is used for extracting spatial features from faces, and a LSTM with global attention for generating an attentional hidden state. They evaluated their method on the DAiSEE dataset and achieved an accuracy of 58.84%.

Dewan et al. [25] and [26], modified the original four-level video engagement annotations in the DAiSEE dataset, and defined two and three-level engagement detection problems based on the labels of other emotional states in the DAiSEE dataset. They also changed the video engagement detection problem in the DAiSEE dataset to image engagement de-

tection problem and performed their experiments on 1800 images extracted from videos in the DAiSEE dataset. For the two-level engagement detection (not-engaged and engaged), the face images originally labeled as bored, confused, and frustrated are assigned the label not-engaged. The labels for the engaged face images has not been changed. For the three-level engagement detection (not-engaged, normally-engaged, and very-engaged), the engaged face images with intensity 1 and 2 are given the label normally-engaged. The engaged face images with the intensity values higher than 2 are given the label very-engaged. The face images with labels bored, confused, and frustrated are given the label not-engaged [25]. In [25], they used Local Directional Pattern (LDP) to extract person-independent edge features for different facial expressions, Kernel Principal Component Analysis (KPCA) to capture the nonlinear correlations among the extracted features, and Deep Belief Networks (DBN) to classify the extracted features and detect the level of engagement. They achieved 90.89%, and 87.25% accuracy for two and three-level engagement detection, respectively. In [26], the extracted face regions are fed to different 2D CNN architectures to detect the level of engagement, and achieved 92.33% accuracy for two-level engagement detection. The authors in [25] and [26] have altered the original *video* engagement detection problem in the DAiSEE dataset [7] to *image* engagement detection problem and have not evaluated their methods on the video test set in the DAiSEE dataset (described in Section IV). Therefore, their reported accuracy would be hard to verify for generalization of results and their methods working with *images* cannot be compared to other works on the DAiSEE dataset that use *videos*.

We formulate the student engagement problem from videos as a spatio-temporal classification problem. We first extract spatial features from raw frames using a ResNet architecture. Since the engagement behavior varies across different video frames, we build a TCN network on the spatial features to model the temporal variation. Next, we describe the ResNet+TCN architecture.

III. RESNET+TCN ARCHITECTURE FOR DETECTING STUDENT ENGAGEMENT

Figure. 1 shows the structure of the ResNet+TCN hybrid neural network architecture for determining student engagement levels in videos. This method is end-to-end, without the need to extract manual or handcrafted features from videos or frames, and the features are learned on the fly while jointly training the network. The input to the ResNet+TCN is a sequence of raw frames of video, and the output is detected ordinal class corresponding to the level of engagement of student in the video.

As discussed above, engagement is a spatio-temporal affective state that takes place in consecutive frames of video over time, its detection requires analyzing video spatially and temporally. The ResNet and TCN are powerful deep neural

network architectures for spatial and temporal analysis of large volumes of data. In previous works, spatial and temporal neural network architectures have been combined in different video-analysis applications. For instance, Donahue et al. [23] combined 2D CNN and LSTM and proposed LRCN for activity recognition and description in video. Petridis et al. [27] combined 3D CNN and ResNet (for spatial feature extraction from video) with bidirectional GRU (for temporal analysis) for lip reading in video. Ma et al. [28] modified the method proposed by Petridis et al. [27] by substituting the GRU with the TCN [19] for lip reading in video.

We adapt an altered version of the architecture proposed in [28] and propose a new architecture using ResNet and TCN. In [28], spatio-temporal features are extracted from videos using 3D CNN. The extracted features by 3D CNN is given to ResNet, and the output of ResNet is given to TCN for temporal analysis. In our architecture, spatial features are extracted from frames using ResNet, and the outputs of ResNet are given to TCN for temporal analysis.

In our proposed architecture, the ResNet [18] and TCN [19] are combined to model and jointly train on the spatio-temporal data in sequences of frames in video. We choose TCN because of its superiority in modeling sequences of larger length and retaining memory of history in comparison to generic recurrent architectures such as LSTMs and GRUs [19]. The ResNet extracts spatial features from single frames, and the TCN models the temporal changes in sequence of frames and outputs the detected level of engagement. Given a video sequence, the input to the ResNet+TCN is an $L \times C \times H \times W$ tensor, where L , C , H , and W correspond to the number of frames, number of channels, frame height, and frame width, respectively. After removing the final fully-connected layer of the standard Resnet18 [29], it is used as the (trainable) feature extractor from single frames of the input video. The extracted feature vectors from the consecutive frames are considered as the multi-dimensional input to the consecutive time steps of a dilated TCN [19] to model temporal information in the video. The output of the final time step of the TCN is fed to a fully-connected layer and a softmax function to detect the level of student engagement from the input video.

One common problem in the existing student engagement detection datasets is the highly imbalanced distribution of engagement levels [7], [6], [5]. The number of samples with low levels of engagement is much less than the number of samples with high levels of engagement. In such sample distribution, it is highly likely that most of the minority-level samples are classified to the majority levels of engagement. To tackle this problem, we use weighted cross entropy loss function [29]. The weights of the cross entropy loss function for each class are determined according to the proportion of the samples in the corresponding class to all the samples in the training set.

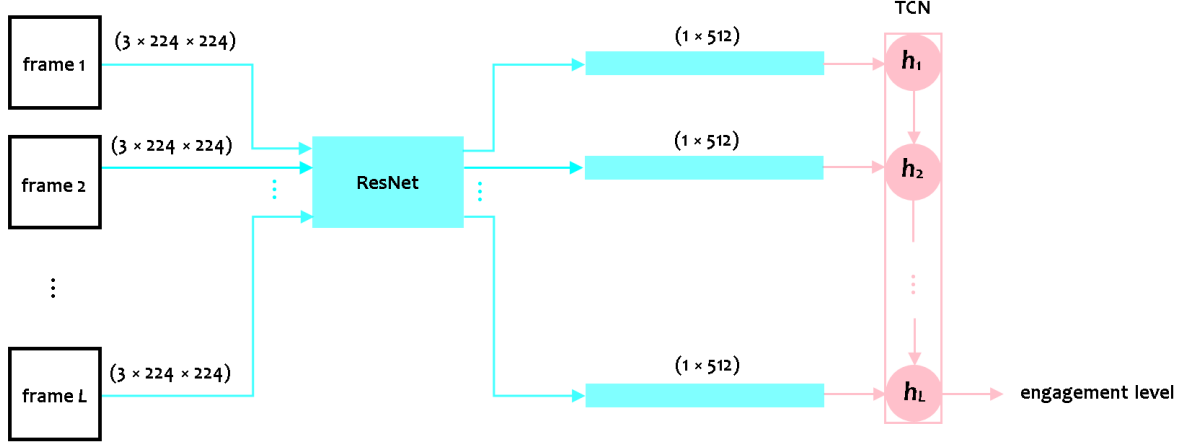


Figure 1: The proposed end-to-end architecture, ResNet+TCN, for engagement level detection from video. The input to the ResNet+TCN is a sequence of L raw frames of video, and the output is detected ordinal class corresponding to the level of engagement of student in the video. The 2D ResNet extracts spatial features from consecutive video frames, and the TCN analyzes the temporal changes in video frames to detect the level of engagement. The extracted feature vectors (by ResNet) from the consecutive frames are considered as the input to the consecutive time steps of the TCN. One fully connected layer after the last time step of the TCN outputs the predicted class.

Another problem encountered during training the network is that due to random sampling in each training iteration, the batches may not contain samples from all the classes because the minority class samples may be left out. Since, in most of training iterations, there are no samples of the minority classes in training batches, the weighted loss function will not be able to affect the network to be trained by the minority class samples. To circumvent this problem, we adopted a customized sampling strategy in which the samples of all classes are included in each batch of the Stochastic Gradient Descent (SGD) during training [29].

level	train	validation	test
0	34	23	4
1	213	143	84
2	2617	813	882
3	2494	450	814
total	5358	1429	1784

Table I: The number of samples in train, validation, and test sets in different levels of engagements in the DAiSEE dataset [7], see Section. IV-A.

IV. EXPERIMENTAL RESULTS

A. Dataset

The DAiSEE dataset, introduced by Gupta et al. [7], is used to evaluate the performance of the ResNet+TCN in comparison to the previous related methods. The dataset contains 9,068 videos captured from 112 students in online courses for recognizing their affective states of boredom, confusion, engagement, and frustration in the wild. We

only focus on the student engagement, with four levels of engagement as level 0 (very low), 1 (low), 2 (high), and 3 (very high). These scores were given based on Whitehill et al. [5]. The length, frame rate, and resolution of the videos are 10 seconds, 30 fps, and 640×480 pixels. Table. I shows the distribution of samples in train, validation, and test sets according to the authors of the DAiSEE dataset [7]. We use these sets in our experiments to fairly compare the ResNet+TCN method with the previous methods on the DAiSEE dataset. We combined 5358 and 1429 videos in the train and validation sets to train ResNet+TCN model and report results on 1784 test videos. As can be seen in Table. I, the dataset is highly imbalanced, only 0.63%, 1.61%, and 0.22% percent of train, validation, and test sets are in the level 0 (corresponding to very low engagement).

B. Experimental Setting

The videos were down-sampled, temporally and spatially, to get $50 \times 3 \times 224 \times 224$ ($L \times C \times H \times W$) tensors as inputs to the ResNet+TCN architecture (as discussed in Section. III). The $3 \times 224 \times 224$ dimensions are the standard dimension for input to the ResNet18 [29]. The SGD is used for parameter optimization with learning rate and batch size of 0.001 and 5, respectively. The ResNet extracts feature vectors of dimension 512 from the consecutive frames and feeds them to the TCN. The parameters of the TCN, giving the best results, are as follows, 8, 128, 7, and 0.25 for the number of levels, number of hidden units, kernel size, and dropout [19]. We implemented the experiments in PyTorch [29] on a server with 64 GB of RAM and NVIDIA Tesla P100 PCIe 12 GB GPU. The code of Resnet+TCN is available at <https://github.com/abedICODES/ResNet-TCN>.

C. Results

For comparing the ResNet+TCN network with other works, we took reported results from the following methods: video-level and frame-level InceptionNet [7], C3D [7], I3D [16], DERN [13], and DFSTN [24]. In addition, we implemented the combination of the ResNet with LSTM, and C3D (up to the layer pool-5) [22] with LSTM (one-layer unidirectional with 128 hidden neurons) to investigate their performance compared to the ResNet+TCN method. Figure. 2 shows the results of applying different end-to-end methods to the four-class engagement level detection problem in the DAiSEE dataset. In this figure, the accuracy is shown for the previous works and the methods we implemented to evaluate on the DAiSEE dataset. Figure. 2 shows that the ResNet+TCN method achieves the highest accuracy of 63.9%, which is 2.75% higher than the best performing method of Resnet with LSTM (61.15%) and 3.9% higher than the state-of-the-art DERN method (60%), and higher than other previous methods. These results show the superiority of the ResNet+TCN in modeling students' engagement level as a spatio-temporal classification problem.

None of the previous works on the DAiSEE dataset, working with the original four-class annotations, reported their confusion matrices for test set [7], [16], [14], [13], [15], [24], and only reported the accuracy results. Therefore, it is hard to determine the individual performance of their methods on each of the engagement levels. We implemented some of the previous methods, including C3D (feature extraction) [7], C3D (fine tuning) [7], C3D + LSTM [30], and C3D averaging + LSTM [30] to investigate their confusion matrices compared to the proposed architectures. Figure. 3 shows the confusion matrices of the ResNet+TCN compared to the other implemented methods. It can be observed that due to the highly-imbalanced sample distribution, none of the methods are able to classify any samples to the first two classes correctly.

Liao et al. [24] reported their confusion matrix for one fold of cross validation during training on the DAiSEE dataset. They didn't report the confusion matrix on the test set of the DAiSEE dataset. Based on their report, they also observed that all the samples were classified to the two higher levels of engagement, and their method cannot correctly classify any samples to the two lower levels of engagement.

As described in Section. III, after using the customized sampling strategy and weighted loss function in our proposed ResNet+TCN architecture, some samples are correctly classified to the first two levels of engagement (Figure. 3 (h)), albeit at the cost of reducing the overall accuracy.

V. CONCLUSION AND FUTURE WORK

In this paper, we presented a new end-to-end spatio-temporal hybrid architecture, ResNet+TCN, for determining the level of engagement among students in an online




Figure 2: Engagement-level classification accuracy of different methods on the DAiSEE dataset [7], (a) video-level InceptionNet [7], (b) frame-level InceptionNet [7], (c) C3D feature extraction [7], (d) C3D averaging + LSTM [30], (e) I3D [16], (f) ResNet + TCN with sampling and weighted loss (proposed), (g) C3D + LSTM [30], (h) LRCN [23], (i) C3D fine tuning [22], (j) DFSTN [24], (k) C3D + TCN (proposed), (l) DERN [13], (m) ResNet + LSTM (proposed), (n) ResNet + TCN (proposed).

classroom setting. We evaluated the performance of the ResNet+TCN method and compared to several previous end-to-end methods. The Resnet+TCN architecture showed improved results in comparison to the state-of-the-art engagement classification accuracy on the DAiSEE dataset. Our results showed that it is very challenging to detect the minority engagement level with very few samples in a supervised classification setting. When class weight term is added to the loss function, some minority level samples were detected at the cost of more false alarms and reducing the overall accuracy of the classifier. In future, we aim to evaluate the performance of the proposed method on other publicly available user engagement datasets, including [6]. We aim to work on utilizing end-to-end models on extracted body/face regions of videos. We also aim to investigate the effectiveness of different features [10], [6], [31] in engagement detection and develop neural network models to detect engagement using those features. One possible direction of research to detect the minority levels of engagement is to apply anomaly detection approaches based on autoencoders and generative adversarial networks.

		predicted labels				
true labels		0	1	2	3	
	0	0	0	3	1	
	1	0	0	35	49	
	2	0	0	395	487	
	3	0	0	351	463	
(a)						
		predicted labels				
true labels		0	1	2	3	
	0	0	0	4	0	
	1	0	0	54	30	
	2	0	0	422	460	
	3	0	0	205	609	
(b)						
		predicted labels				
true labels		0	1	2	3	
	0	0	0	4	0	
	1	0	0	78	6	
	2	0	0	570	312	
	3	0	0	374	440	
(c)						
		predicted labels				
true labels		0	1	2	3	
	0	0	0	3	1	
	1	0	0	32	52	
	2	0	0	495	387	
	3	0	0	379	435	
(d)						
		predicted labels				
true labels		0	1	2	3	
	0	0	0	3	1	
	1	0	0	62	22	
	2	0	0	623	259	
	3	0	0	339	475	
(e)						
		predicted labels				
true labels		0	1	2	3	
	0	0	0	4	0	
	1	0	0	77	7	
	2	0	0	577	305	
	3	0	0	322	492	
(f)						
		predicted labels				
true labels		0	1	2	3	
	0	0	0	4	0	
	1	0	0	64	20	
	2	0	0	616	266	
	3	0	0	290	524	
(g)						
		predicted labels				
true labels		0	1	2	3	
	0	1	0	2	1	
	1	2	10	39	33	
	2	16	38	505	323	
	3	6	4	362	442	
(h)						

Figure 3: Engagement-level confusion matrices of different methods on the DAiSEE dataset [7], (a) C3D feature extraction [22], (b) C3D fine tuning [22], (c) C3D + LSTM [30], (d) C3D averaging + LSTM [30], (e) ResNet + LSTM (proposed), (f) C3D + TCN (proposed), (g) ResNet + TCN (proposed), (h) ResNet + TCN with weighted sampling and weighted loss (proposed).

REFERENCES

- [1] K. Mukhtar, K. Javed, M. Arooj, and A. Sethi, "Advantages, limitations and recommendations for online learning during covid-19 pandemic era," *Pakistan journal of medical sciences*, vol. 36, no. COVID19-S4, p. S27, 2020.
- [2] A. Nuara, M. Fabbri-Destro, E. Scalona, S. E. Lenzi, G. Rizolatti, and P. Avanzini, "Telerehabilitation in response to constrained physical distance: an opportunity to rethink neurorehabilitative routines," *Journal of neurology*, pp. 1–12, 2021.
- [3] B. J. Venton and R. R. Pompano, "Strategies for enhancing remote student engagement through active learning," 2021.
- [4] Ö. Sümer, P. Goldberg, S. D'Mello, P. Gerjets, U. Trautwein, and E. Kasneci, "Multimodal engagement analysis from facial videos in the classroom," *arXiv preprint arXiv:2101.04215*, 2021.
- [5] J. Whitehill, Z. Serpell, Y.-C. Lin, A. Foster, and J. R. Movellan, "The faces of engagement: Automatic recognition of student engagement from facial expressions," *IEEE Transactions on Affective Computing*, vol. 5, no. 1, pp. 86–98, 2014.
- [6] A. Dhall, G. Sharma, R. Goecke, and T. Gedeon, "Emotiiv 2020: Driver gaze, group emotion, student engagement and physiological signal based challenges," in *Proceedings of the 2020 International Conference on Multimodal Interaction*, 2020, pp. 784–789.
- [7] A. Gupta, A. D'Cunha, K. Awasthi, and V. Balasubramanian, "Daisee: Towards user engagement recognition in the wild," *arXiv preprint arXiv:1609.01885*, 2016.
- [8] P. Guhan, M. Agarwal, N. Awasthi, G. Reeves, D. Manocha, and A. Bera, "Abc-net: Semi-supervised multimodal gan-based engagement detection using an affective, behavioral and cognitive model," *arXiv preprint arXiv:2011.08690*, 2020.
- [9] K. Doherty and G. Doherty, "Engagement in hci: conception, theory and measurement," *ACM Computing Surveys (CSUR)*, vol. 51, no. 5, pp. 1–39, 2018.
- [10] M. A. A. Dewan, M. Murshed, and F. Lin, "Engagement detection in online learning: a review," *Smart Learning Environments*, vol. 6, no. 1, pp. 1–20, 2019.
- [11] O. M. Nezami, M. Dras, L. Hamey, D. Richards, S. Wan, and C. Paris, "Automatic recognition of student engagement using deep learning and facial expression," in *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*. Springer, 2019, pp. 273–289.
- [12] S. D'Mello, E. Dieterle, and A. Duckworth, "Advanced, analytic, automated (aaa) measurement of engagement during learning," *Educational psychologist*, vol. 52, no. 2, pp. 104–123, 2017.
- [13] T. Huang, Y. Mei, H. Zhang, S. Liu, and H. Yang, "Fine-grained engagement recognition in online learning environment," in *2019 IEEE 9th international conference on electronics information and emergency communication (ICEIEC)*. IEEE, 2019, pp. 338–341.
- [14] Y. Wang, A. Kotha, P.-h. Hong, and M. Qiu, "Automated student engagement monitoring and evaluation during learning in the wild," in *2020 7th IEEE International Conference on Cyber Security and Cloud Computing (CSCloud)/2020 6th IEEE International Conference on Edge Computing and Scalable Cloud (EdgeCom)*. IEEE, 2020, pp. 270–275.
- [15] L. Geng, M. Xu, Z. Wei, and X. Zhou, "Learning deep spatiotemporal feature for engagement recognition of online courses," in *2019 IEEE Symposium Series on Computational Intelligence (SSCI)*. IEEE, 2019, pp. 442–447.

- [16] H. Zhang, X. Xiao, T. Huang, S. Liu, Y. Xia, and J. Li, "An novel end-to-end network for automatic student engagement recognition," in *2019 IEEE 9th International Conference on Electronics Information and Emergency Communication (ICEIEC)*. IEEE, 2019, pp. 342–345.
- [17] G. M. Sinatra, B. C. Heddy, and D. Lombardi, "The challenges of defining and measuring student engagement in science," 2015.
- [18] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich, "Going deeper with convolutions," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2015, pp. 1–9.
- [19] S. Bai, J. Z. Kolter, and V. Koltun, "An empirical evaluation of generic convolutional and recurrent networks for sequence modeling," *arXiv preprint arXiv:1803.01271*, 2018.
- [20] J. Wu, B. Yang, Y. Wang, and G. Hattori, "Advanced multi-instance learning method with multi-features engineering and conservative optimization for engagement intensity prediction," in *Proceedings of the 2020 International Conference on Multimodal Interaction*, 2020, pp. 777–783.
- [21] B. Zhu, X. Lan, X. Guo, K. E. Barner, and C. Boncelet, "Multi-rate attention based gru model for engagement prediction," in *Proceedings of the 2020 International Conference on Multimodal Interaction*, 2020, pp. 841–848.
- [22] D. Tran, L. Bourdev, R. Fergus, L. Torresani, and M. Paluri, "Learning spatiotemporal features with 3d convolutional networks," in *Proceedings of the IEEE international conference on computer vision*, 2015, pp. 4489–4497.
- [23] J. Donahue, L. Anne Hendricks, S. Guadarrama, M. Rohrbach, S. Venugopalan, K. Saenko, and T. Darrell, "Long-term recurrent convolutional networks for visual recognition and description," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2015, pp. 2625–2634.
- [24] J. Liao, Y. Liang, and J. Pan, "Deep facial spatiotemporal network for engagement prediction in online learning," *Applied Intelligence*, pp. 1–13, 2021.
- [25] M. A. A. Dewan, F. Lin, D. Wen, M. Murshed, and Z. Uddin, "A deep learning approach to detecting engagement of online learners," in *2018 IEEE SmartWorld, Ubiquitous Intelligence & Computing, Advanced & Trusted Computing, Scalable Computing & Communications, Cloud & Big Data Computing, Internet of People and Smart City Innovation (SmartWorld/SCALCOM/UIC/ATC/CBDCom/IOP/SCI)*. IEEE, 2018, pp. 1895–1902.
- [26] M. Murshed, M. A. A. Dewan, F. Lin, and D. Wen, "Engagement detection in e-learning environments using convolutional neural networks," in *2019 IEEE Intl Conf on Dependable, Autonomic and Secure Computing, Intl Conf on Pervasive Intelligence and Computing, Intl Conf on Cloud and Big Data Computing, Intl Conf on Cyber Science and Technology Congress (DASC/PiCom/CBDCom/CyberSciTech)*. IEEE, 2019, pp. 80–86.
- [27] S. Petridis, T. Stafylakis, P. Ma, F. Cai, G. Tzimiropoulos, and M. Pantic, "End-to-end audiovisual speech recognition," in *2018 IEEE international conference on acoustics, speech and signal processing (ICASSP)*. IEEE, 2018, pp. 6548–6552.
- [28] P. Ma, Y. Wang, J. Shen, S. Petridis, and M. Pantic, "Lip-reading with densely connected temporal convolutional networks," in *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 2021, pp. 2857–2866.
- [29] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga et al., "Pytorch: An imperative style, high-performance deep learning library," *arXiv preprint arXiv:1912.01703*, 2019.
- [30] P. Parmar and B. Tran Morris, "Learning to score olympic events," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, 2017, pp. 20–28.
- [31] J. Yang, K. Wang, X. Peng, and Y. Qiao, "Deep recurrent multi-instance learning with spatio-temporal features for engagement intensity prediction," in *Proceedings of the 20th ACM international conference on multimodal interaction*, 2018, pp. 594–598.