

# Action Quality Assessment using Siamese Network-Based Deep Metric Learning

Hiteshi Jain, Gaurav Harit and Avinash Sharma

**Abstract**—Automated vision-based score estimation models can be used as an alternate opinion to avoid judgment bias. In the past works the score estimation models were learned by regressing the video representations to the ground truth score provided by the judges. However such regression-based solutions lack interpretability in terms of giving reasons for the awarded score. One solution to make the scores more explicable is to compare the given action video with a reference video. This would capture the temporal variations vis-à-vis the reference video and map those variations to the final score. In this work, we propose a new action scoring system as a two-phase system: (1) A Deep Metric Learning Module that learns similarity between any two action videos based on their ground truth scores given by the judges; (2) A Score Estimation Module that uses the first module to find the resemblance of a video to a reference video in order to give the assessment score. The proposed scoring model has been tested for Olympics Diving and Gymnastic vaults and the model outperforms the existing state-of-the-art scoring models.

**Index Terms**—automatic scoring, Siamese, LSTM, Deep Metric Learning

## I. INTRODUCTION

**H**UMANS strive to attain perfection and efficiency in their day-to-day activities and maximize the outcome of their actions. They improve their performances by comparing themselves with others or seeking feedback from experts. The task of Action Quality Assessment (AQA) involves quantification of the quality of action *i.e.* determining *how well* an action was performed. Unlike the traditional action recognition/classification task widely attempted in the computer vision community, where a subset of discriminative frames can help identify the type of action, the task of action assessment is more challenging as it requires quantitative and qualitative assessment of the entire action sequence. AQA has applications in many areas such as sports [16], [15], health-care [23], rehabilitation [14], exercise [6], [7] and actions of daily living [27].

The key challenge in AQA is modelling the human subjectivity as this task suffers from significant influence of human (or expert) bias. Traditionally, the subjectivity in action assessment has been addressed by considering feedbacks given by multiple experts, *e.g.* Olympics events allows averaging scores from multiple judges (while dropping the highest and lowest scores). Such biases may include nationalistic bias [2], [4], where judges give higher scores to performers from their own countries, a difficulty bias [12], where athletes attempting

more difficult routines receive higher execution scores, etc. Nevertheless, the human bias in scoring is inevitable and hence only partially addressable in any scenario. Thus, an automated AQA system should aim bringing more interpretability in predicted scores while being an objective source of alternate evaluation. However, such a solution needs to learn all possible nuances that a human expert/judge learns over the years.

Majority of the existing works on action quality assessment/scoring (*e.g.*, [16], [19], [15], [20]) perform holistic evaluation of action videos and learn to predict the assessment score from the video features. This is achieved by modeling the task of score estimation as a regression task where the action video representation is regressed to the ground truth score provided by the judges. Apart from holistic evaluation, it is also important to have a feedback on how different sub-actions contribute to the final score for better interpretability of the judgments. Parmar *et al.* [15] assume that all the sub-action clips of the performance contribute equally towards the final score. In addition to back propagating gradients for the final loss, they also back propagate gradients for the individual sub-actions. However, considering equal contributions of all clips during training is not realistic as different segments of an action have different contributions to the final score. In [8], [20], authors compute the contributions of different segments of a performance towards the final scoring. This is achieved by learning separate regression models for a segment towards the final score. These regression models then predict the sub-action scores of the test videos. The rank correlation between the predicted sub-action scores and the final scores define the level of contribution of the sub-actions towards the final score. For *e.g.* in diving, the splashes are seen to have the maximum contribution towards the scoring. Validation of the sub-actions scores of the individual videos has not been carried out.

A more meaningful solution to bring more objectivity and interpretability to an action assessment system is to consider a set of reference action videos as a pivot to evaluate a new action video. Such reference videos are expert's videos, which have the action performed with high precision. Thus, the problem of action quality assessment can be transformed into the problem of comparing a given action video with a reference video. However, one practical challenge is the lack of availability of a large number of reference videos to learn a regression model. Therefore, instead of learning how to compare a given action with an expert's action, we simply learn how to compare two actions. In our work we first train a metric learning module that can compare any two performances and predict a binary similarity label, and then learn a regression module that can use the learned metric to compare it with an expert's performance. We introduce a

Hiteshi Jain and Gaurav Harit are with Department of Computer Science and Engineering, Indian Institute of Technology Jodhpur, Karwar, Jodhpur, India, 342037 e-mail: jain.4@iiitj.ac.in.

Avinash Sharma is from International Institute of Information Technology, Gachibowli, Hyderabad, 500032, India

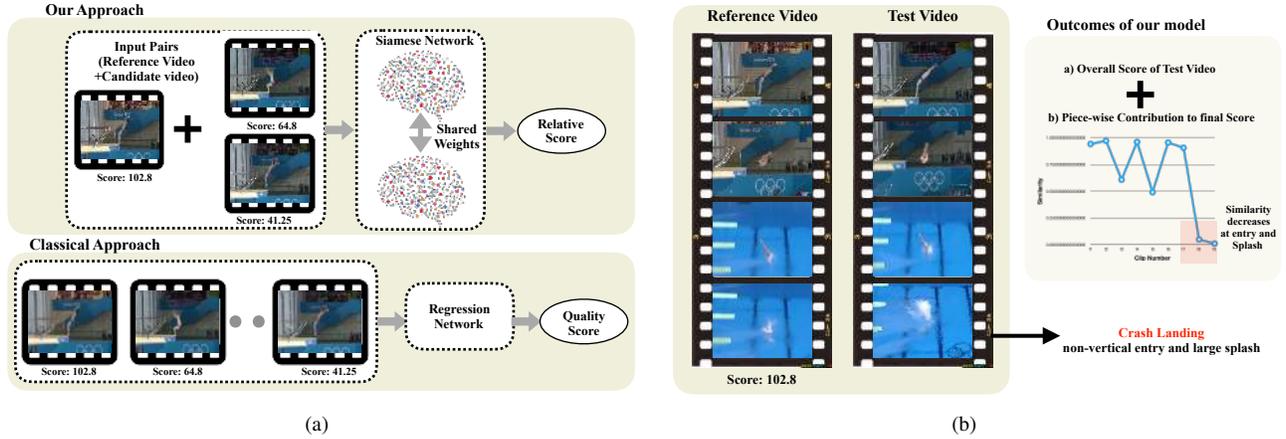


Fig. 1. Overview of our work; The task of scoring is transformed to the task of comparing a performance to an expert performance; Outputs of our model: Overall performance score and Clip-level feedback

novel deep learning-based approach for AQA in which the performances are compared with a reference performance to determine the score. The proposed scoring system works in two phases (Fig. 1):

- A Deep Metric Learning (DML) [24] module learns the similarity metric between two performances using the Siamese-LSTM network. The target label for this module is a binary label indicating similar/dis-similar actions. Training this module requires inputs as pairs of videos. This has the advantage that even smaller AQA datasets are able to generate a large number of pairs for training.
- In the second phase, a score prediction module is developed that uses the learned similarity metric and estimates the score of a video based on its resemblance to the reference video. The input for training this module is a pair – an action video and an expert’s video performing the same action.

Additionally, we attempt to find the segment (sub-activity) level similarity using the learned DML network and identify how different segments contribute to the final score of the performance. This helps to make the judgement more interpretable in terms of how a particular action in a video segment is performed compared to an expert performance. The sub-action (segment) level feedback can be verified by experts and can be used as a ground truth for future works.

The key contributions of this work are as follows:

- 1) We introduce a novel approach for automated AQA in which the predicted score depends on a relative comparison between the action performances in the test video and a reference expert video.
- 2) We train a DML network to learn the similarity metric between two action sequences based on the difference in the ground truth scores. Subsequently, a scoring network then regresses the concatenated representations (from the DML module) of the test video and a reference video to the final score.
- 3) *Feedback Proposal* - The learned DML network is also utilized to provide an unsupervised way of finding segment (sub-action) level contributions towards final scor-

ing thereby making the scoring more interpretable.

- 4) Through experiments we demonstrate the superiority of the proposed framework over the state-of-the-art solutions for AQA on publicly available Olympics diving and gymnastic vaults datasets.

The paper is organized as follows : Section II discusses the existing literature towards human action assessment and action scoring. Section III explains our scoring model architecture. Section IV describes a technique to generate sub-action level feedback for the performers. Section V includes the experimental results using the proposed approach and discusses the performance of the proposed model. Finally the paper is concluded in Section VI.

## II. RELATED WORK

The field of human action assessment is relatively new and there are only few works that have addressed the problem. Few early works [7], [6], [11], [10], [9] in the domain were hand crafted for specific actions and could not be generalised to different types of actions. Recently, there has been development of more generic assessment frameworks, that can broadly be divided into two categories: 1) Action Quality Assessment (AQA) and 2) Skill Assessment.

### A. Action Quality Assessment (AQA):

Quality scoring for human actions has been posed as a supervised regression task, where a model is learned to map the human action features to the ground-truth scores annotated by a domain expert.

Pirsiavash *et al.* [16] utilized the low frequency DCT/DFT components of the pose features as input to the Support Vector Regressor (SVR) to map to the final action quality score. The pose features were estimated using Flexible Parts Model [21] for each frame individually. Use of low frequencies filtered out high frequency noise due to pose estimation errors. Temporal resolution of features was improved with window-based extraction of frequency components.

In addition to the holistic scores prediction, the authors propose a feedback scheme in terms of giving directions

where the body parts should move to maximize the score. This is accomplished by differentiating the scoring function with respect to joint location. The maximum value of the gradient of the score with respect to the location of each joint indicates the direction that the performer must move to achieve the largest improvement in the score. The performance of the proposed assessment scheme was evaluated on the *MIT Olympics dataset* that consisted of two action categories: Olympics diving and figure-skating sports.

Venkataraman *et al.* [19] calculated the approximate entropy features using the estimated pose for each frame and concatenated them to get a high-dimensional feature. These features were seen to better encode the dynamical information compared to DCT and provide a better assessment over the MIT Olympics dataset.

Pose features provide interpretable feedback such as directions in which the limbs should be moved to maximize scores. However, errors in pose estimation can negatively impact the scoring models. Moreover, pose-based representations are not capable of modeling objects used during an action (such as sports ball or tools), and do not consider physical outcomes, e.g. splashes in diving, which may be important features for some activities.

Visual spatio-temporal features such as Convolution 3D features (C3D) [18] and Pseudo 3D convolution features (P3D) [30], learned from 3D convolution neural networks capture both appearance and subtle motion cues in videos naturally. These features have been recently utilized to represent human actions for the purpose of scoring [15][20][8]. Feature extraction has been done using pre-trained models [15][8] or learned in an end-to-end setup with a scoring objective [20].

Parmar and Morris [15] proposed three frameworks for human action quality assessment. These were: C3D-SVR, C3D-LSTM, and C3D-LSTM-SVR. The three frameworks differed in the way the features were aggregated and the regression was formulated. The authors proposed a new *UNLV sports dataset* which had twice the number of diving examples than the MIT dataset and also had gymnastic vault examples. The frameworks proved to be more efficient than the pose-based scoring works [16], [19].

The authors [15] introduced a new training protocol called the *incremental label training* to provide sub-action level feedback. It is expected that as an action advances in time, the score should build up (if the quality is good enough) or be penalized (if the quality is sub par). This intuition was formulated as a score getting accumulated as a non-decreasing function throughout an action. Thus a video is divided into equal sized non-overlapping clips such that each clip is assumed to have an equal contribution to the final score. Incremental-label training is used to guide the LSTM during the training phase to generate the final score along with intermediate score outputs (i.e. back-propagation occurs after each clip). The temporal score evolution as it changes through the LSTM structure is utilized to identify both “good” and “poor” segments of an action. However, training a network with an assumption that the clips have an equal contribution to the final score is not realistic.

Xiang *et al.* [20] proposed to divide a video into action

specific semantic segments and fused the segment-averaged P3D features to learn the final score. The segment-level features were learned by fine-tuning the P3D network to the score regression task. Li *et al.* [8] divided a video sample into 9 clips and used 9 different C3D networks dedicated to different stages of Diving. The clip-level features were concatenated and fed to a dense network to produce a final AQA score. The network was optimized using the ranking loss along with the L2 loss. The ranking loss ensured the right rank order of the predicted scores.

In both the aforementioned works [20], [8], authors contribute towards finding the contributions of different segments of a performance towards the final scoring. Different sub-action segments of all the training videos are separately regressed to the final scores. The rank correlation between the predicted sub-action scores and the final scores define the level of contribution of the sub-actions towards the final score. For e.g., in diving, the splashes are seen to have the maximum contribution towards the scoring.

More recently, Xu *et al.* [33] addressed the problem of scoring figure skating. Figure skating videos are long and last for few minutes unlike vaults and diving that last for only few seconds. The assessment of such videos requires evaluation of temporal segments as well as holistic performance. The authors proposed a deep architecture with 2 complementary components: Self-Attentive LSTM and Multi-scale Convolutional Skip LSTM, which learn the local and global sequential information respectively from a video. They introduced *FisV*, a new large scale figure skating video dataset for evaluating the proposed scheme. The attention weight matrix computed using the self attentive mechanism for a specific video signifies clip-level significance. A clip that has high weight in at least one row of the attention weight matrix shows an important technical movement and is otherwise considered insignificant. The multi-scale LSTM employs several parallel 1D convolution layers with different kernel sizes unlike C3D or P3D that have a fixed kernel size. A kernel with small size extracts the visual representation of action patterns lasting seconds in the videos. A kernel of large size tries to model the global information of the videos.

In this work, we develop an action assessment system which can provide segment-level feedback. For learning a scoring model we consider a set of reference videos, which we call as *expert videos*, i.e. the highest rated performers for the same type of action. We transform the problem of action assessment into the problem of (i) Generating a video representation which incorporates comparison with an expert video of the same type of action. (ii) Mapping the learned representation to the aggregate score using a regression model. The same network is also used to do a segment level comparison with a corresponding expert video segment and generate scores which depict assessment of the sub-action performed in that segment. The proposed system thus provides a more interpretable assessment.

Some works have addressed classifying performances into skill levels. Zia *et al.* [23] extract spatio-temporal interest points (STIPs) in the frequency domain to classify the performances into amateur, intermediate or expert skill levels.

Doughty *et al.* [27] learned convolutional features with ranking loss objective function to evaluate surgical, drawing, chopstick use and dough rolling skills. In their next work [28], they use temporal attention LSTMs to learn weighted contributions of different segments of the video to the final skill levels. The benefit of using Siamese networks for the task of learning the relative skills of performances has encouraged us to adopt a similar approach to the task of action quality assessment and provide interpretable feedback along with the scores.

In this work, we propose a new action quality assessment approach that transforms the problem of assessment into the problem of comparing a given action video with a reference video. Thus we require a model that can learn the similarity between two action sequences. This problem relates closely with the problem of estimating similarity of two signature samples. A Siamese network [31] is widely used as a deep metric learning-based (DML) approach [22] to learn the similarity between two sequences. The DML approach has been used for human action recognition in [32] and has shown promising results. While recognition is an inter-class discrimination problem, it remains to be explored if DMLs can learn differences between very similar sequences of the same class. The level of resemblance can be translated into scores pertaining to assessment of actions. We discuss our model in Section III.

### B. Action Quality Assessment Datasets

The past works have contributed Olympics dataset for 3 actions: Diving, Gymnastic Vaults and Figure Skating.

*Diving Dataset:* The MIT-Diving dataset [16] consists of 159 videos taken from 2012 Olympic mens 10-meter platform prelims round. The UNLV-diving dataset [15] is an extension to this dataset which includes semi-final and final round videos totaling to 370 videos, each having around 150 frames. A diving score is determined by the product of execution score, based on judge’s assessment of the quality of diving, multiplied by the diving difficulty score, which is a fixed value based on diving type. The scores vary between 0 (the worst) and 100 (the best).

More recently Parmar *et al.* [1] proposed a Multitask AQA Dataset with 1412 samples of diving videos including 10m Platform as well as 3m Springboard with both male and female athletes, individual or pairs of synchronized divers, and different views. The AQA score, dive type and the commentary for each sample has been included in the dataset.

*Vault Dataset:* The UNLV-Vault Dataset [15] dataset includes 176 videos. These videos are relatively short with an average length of about 75 frames. A vault score is determined by the sum of the execution score and the difficulty score. The score ranges from 0 (the worst) to 20 (the Best).

*Figure Skating:* Pirsiavash *et al.* [16] proposed the MIT-skate dataset has 150 videos with 24 frames per second. On an average, figure skating samples are 2.5 minutes long and with continuous view variation during a performance. The judges score ranges between 0 (worst) and 100 (best). Xu *et al.* [33] proposed the Fis-V dataset which consists of 3 times more videos than the MIT-skate dataset.

In this work we consider two action types - diving and gymnastic vaults. We do not consider figure skating because it is characterized by large variations of routines attempted by performers and does not have a standard template and therefore not suitable for assessment using comparison.

## III. PROPOSED SCORING MODEL

In this work, we address the problem of human action scoring as the problem of comparing a given action video with a reference video depicting a high rated performance. Our proposed assessment system is a two-phase system: 1) A *Deep Metric Learning Module* which utilizes a Siamese Network to learn a similarity metric between any two action sequences. 2) A *Score Estimation Module* which utilizes the learned Siamese Network and uses regression to estimate the score of a video based on its similarity to the expert execution of the same action.

We claim that learning a similarity metric before training a scoring regression model offers an advantage because even small training datasets can give a large number of pairs to train the DML network. A pair need not include an expert sequence. We explain these steps in detail.

### A. Deep Metric Learning Module

The Deep Metric Learning (DML) model is trained to learn the similarity between any two videos in a pair. The pair of sequences to be used as training input are sampled from the set of all videos. Fig. 2 shows the structure of the proposed metric learning module. This module takes two video sequences  $V_p$  and  $V_q$  as input and produces similar/dissimilar (1/0) label as output. Similar to [15], the clip-level C3D FC-6 layer features [18] are fed to an LSTM layer to generate a video-level description. The target labels to be learned for the model are binary valued 1/0, which classify the pairs to be either similar or dissimilar. The target label for each pair is decided based on the difference in their respective scores assigned by the judges. If the difference in the scores is less than a threshold  $th$ , the pairs are considered as similar (label 1), and otherwise dissimilar. The threshold is decided based on the type of activity being assessed. We would discuss the setting of this threshold in the experiments section.

The C3D model provides a more compact video representation than frame-level CNN. Since an entire clip is processed in one go, this results in fewer processing steps. For eg. a video with 145 frames can be represented as a sequence of 9 clips each containing 16 frames (input size required for the C3D network). Each of the two LSTM networks takes one sequence as input and produces individually two output vectors  $O_p \in R^M$  and  $O_q \in R^M$ . These vectors are then concatenated and fed into 2 fully connected dense layers,  $D_1$  and  $D_2$ . The output of the dense layer is passed to one fully connected Sigmoid layers to map the combined features of the two videos to a binary similar/non-similar classification output.

The effectiveness of the Deep Metric Learning module is in terms of how well it transforms the inputs to a feature space such that the videos with similar scores are closer in the feature

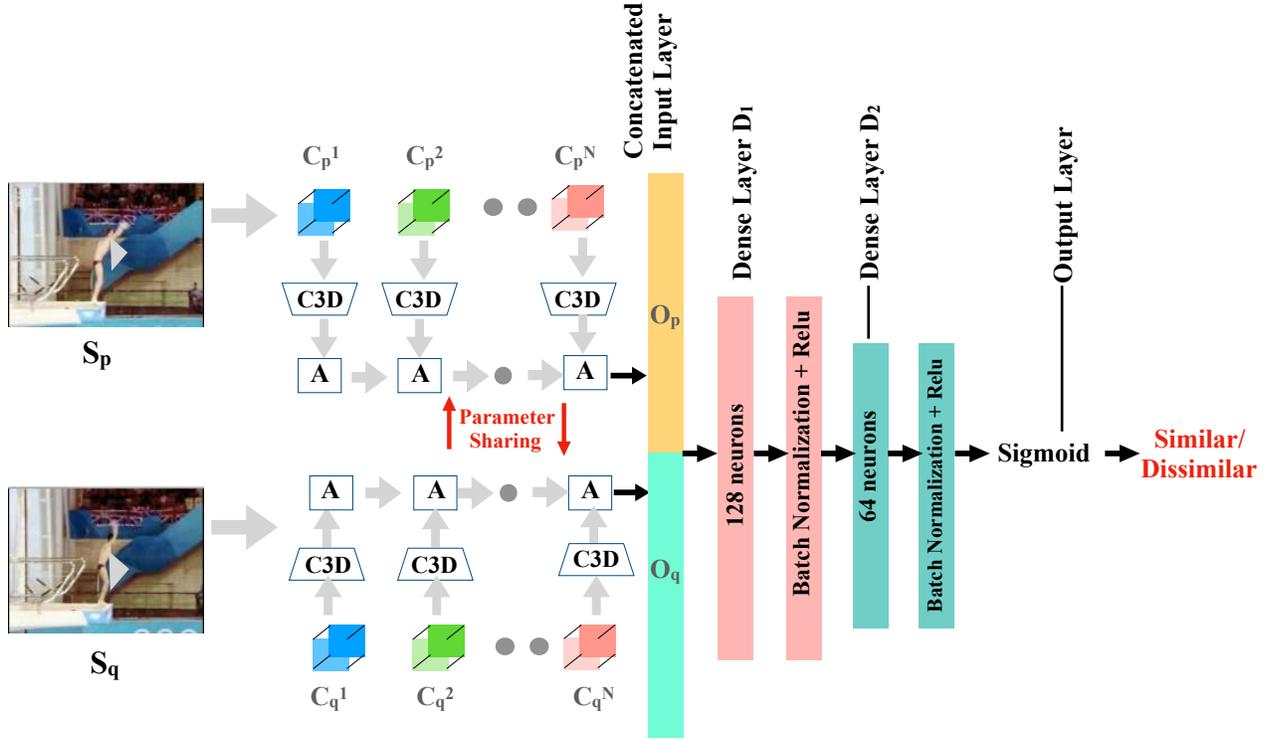


Fig. 2. Siamese Architecture for Deep Metric Learning

space and those with a large difference in scores are moved farther apart.

### B. Score Estimation Module using Expert bias

Our proposal for action quality assessment is to compare a video with an expert (highest rated) video and map the result of comparison to its final score. This is different from how the traditional scoring models have been working where a video representation is directly regressed to the scores.

An action can be performed in various ways. For eg. diving can be performed in forms like forward somersaults, inward somersaults, backward somersaults, etc. In order to do scoring in our model, each video is compared to an expert's performance of the same type of action. Further, we can have one or more experts for an action type while training the scoring model.

The scoring model utilizes the DML Siamese model, which is a subset of the architecture shown in Fig. 2. The model takes input as a pair with an expert and a non-expert video depicting the same type of action. The last sigmoid layer is replaced with a fully connected (FC) layer, so that the Siamese network model originally trained for classifying a pair of sequences as similar or dissimilar can now be used for the regression task and estimate the relative score of the video. The LSTM-Siamese weights are learned while training the DML module. These weights are fixed during the learning of the score estimation module. Instead, only the weights for the final connected layer which performs regression are learned.

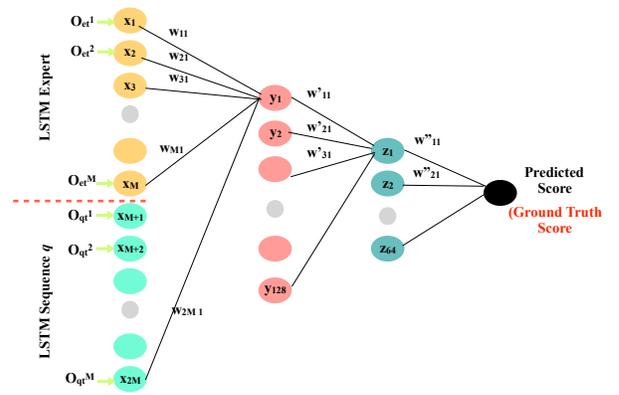


Fig. 3. Score Estimation Relative to Expert

To summarize, the DML module is trained with action pairs that include any action sequence  $V_p$  and  $V_q$  as input and the match or no-match labels as output, while the score estimation module is trained with the input as a pair of videos with  $E_t$  being an expert video of type  $t$  and  $V_{qt}$  is the  $q^{\text{th}}$  training video of the same action type  $t$  and the output is the ground truth score of the sequence  $V_{qt}$ .

*Expert performance as bias to the network* The Siamese embedding of the expert acts as a fixed but dive type dependent bias for the regression network, because the expert is fixed for multiple videos of the same action type. Here we show that the gradient of the loss used to update the weights  $w''$  (Fig. 3) of the last layer contains a contribution of expert's

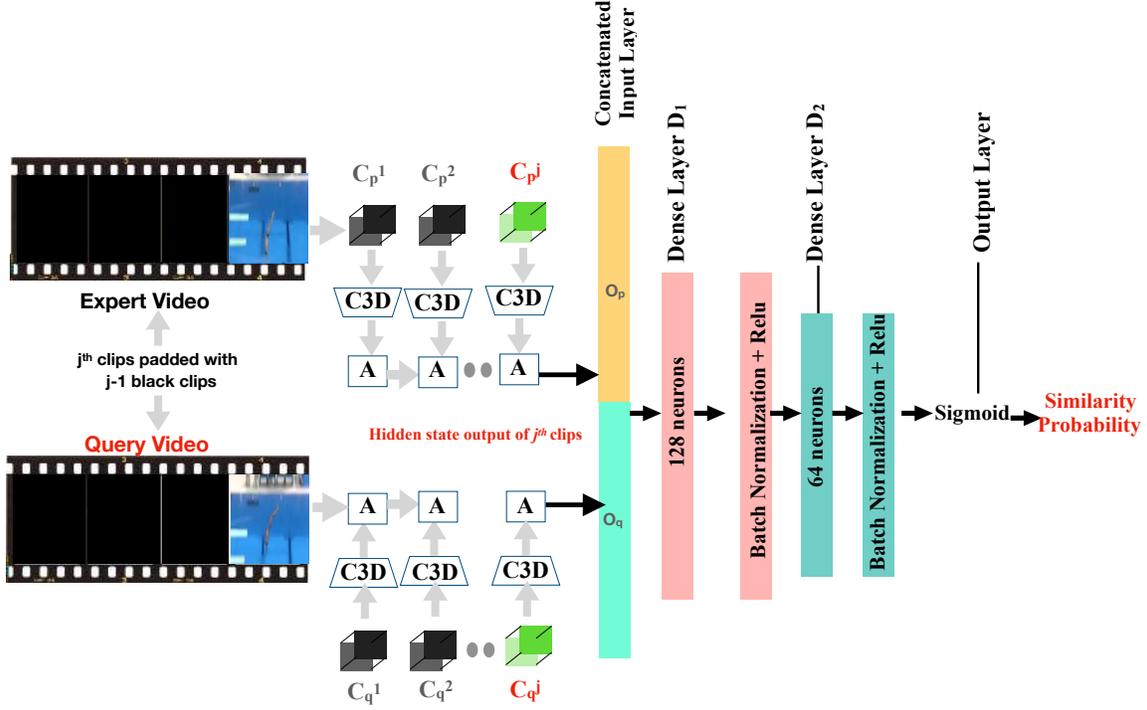


Fig. 4. Clip level similarity Estimation for evaluating individual clip contribution

bias which is different for different dive types. Let the LSTM embedding of the expert videos  $E_t$  for action type  $t$  in the pair be  $O_{et} \in R^M$  and that of the  $q^{th}$  training video be  $O_{qt} \in R^M$ . Let the concatenation of LSTM output of the two vectors be represented as vector  $X \in R^{2M}$  and likewise the outputs of the dense layers  $D_1$  and  $D_2$  be represented as  $Y \in R^{128}$  and  $Z \in R^{64}$ .

The weights of the dense layers ( $w$  from  $X$  to  $Y$  and  $w'$  from  $Y$  to  $Z$ ) are fixed and are previously learned during DML training. During score estimation only the weights from layer  $Z$  to the final regression layer are learned with the target of the ground truth score provided by judges. The scores are estimated following these steps:

- 1) The vector  $Y$  for action type  $t$  and  $q^{th}$  training video of the same type is computed as

$$y_{qi}^t = \sum_{j=1}^M w_{ji} x_j^t + \sum_{j=M+1}^{2M} w_{ji} x_{qj}^t$$

where the first term is the contribution from the expert and is common for all training videos of one type. We ignore the bias to keep the calculations easy.

- 2) Similarly the vector  $Z$  for action type  $t$  and  $q^{th}$  training video is computed as

$$\begin{aligned} z_{qk}^t &= \sum_{i=1}^{128} w'_{ik} \left( \sum_{j=1}^M w_{ji} x_j^t + \sum_{j=M+1}^{2M} w_{ji} x_{qj}^t \right) \\ &= \sum_{i=1}^{128} w'_{ik} \sum_{j=1}^M w_{ji} x_j^t + \sum_{i=1}^{128} w'_{ik} \sum_{j=M+1}^{2M} w_{ji} x_{qj}^t \end{aligned}$$

the first term is again the contribution from the expert, common for all training videos of one type and the second

term depends on the  $q^{th}$  training video of type  $t$ . Let the first term be represented by  $a_{ek}^t$  and the second term be  $z x_{qk}^t$ . Thus  $z_{qk}^t = a_{ek}^t + z x_{qk}^t$

- 3) The final score is thus computed as

$$S_q^{t'} = \sum_{k=1}^{64} w''_k (a_{ek}^t + z x_{qk}^t)$$

Let the ground truth score of the  $q^{th}$  video be  $S_q$ . The total loss over all training examples for all types is:

$$Loss = \sum_{t=1}^{\#action\_types} \sum_{q \in t} (S_q - \sum_{k=1}^{64} w''_k (a_{ek}^t + z x_{qk}^t))^2$$

Thus, while learning to estimate the scores of the videos, experts of different action types contribute in different ways during relative score prediction. To make the settings more general, we can consider the bias as a contribution by any kind of performance, e.g. it could be expert, intermediate or worst performer.

### C. Module Training

When training the DML module, there is a label imbalance problem because the number of non-matching pairs is much larger than the matching ones. To keep the training set balanced, we increase the match pairs by using augmentation techniques like video zoom, brightness alteration, temporal augmentation by dropping some random frames, etc. Training of the score estimation module follows the training of DML module. The score estimation module is trained using only the original training dataset, without using any augmentation.

#### IV. FEEDBACK PROPOSALS

Parmar *et al.* [15] proposed an incremental label training where it was expected that the temporal score evolution as it changes through the LSTM should be monotonically increasing. The “good” and “poor” components of an action are identified when the evolution violates monotonicity and the errors are expected to result in a loss of score. Following this intuition, the predicted sub-scores from the incremental training approach are analysed and the faulty clips are identified.

Events like Olympics Diving and Gymnastic vaults are annotated only by a final score provided by the judge. In order to make the assessments more interpretable and estimate how different parts of the videos contribute to the final score, we compute the similarity of each clip of a test video to the corresponding clip of the expert video. To compute the similarity of the  $j^{\text{th}}$  clips of the two videos, we first create the corresponding trimmed sequences which retain the  $j^{\text{th}}$  clips, but remove the succeeding frames and zero out all the preceding frames. Feeding the trimmed video as input to the LSTM provides the final hidden states at the end of the  $j^{\text{th}}$  clip. LSTM hidden states for the test and the expert video are concatenated and fed to the dense layers of the Siamese network. (Fig. 4) The final Sigmoid layer gives the probability of similarity of the two clips. A higher probability implies that the actions corresponding to the  $j^{\text{th}}$  clips of the two videos were performed in a similar way.

#### V. EXPERIMENTS

We compare our relative scoring approach to the baseline absolute scoring techniques in this section.

##### A. Dataset Splits

The experiments are based on the Olympics UNLV Diving Dataset and the Gymnastic Vaults Dataset[15]. For the diving dataset, 10 training splits are used to compare the performance of our model to the baseline models. For training the DML module we choose the top 3 most frequent dive types from the UNLV dataset while for training the scoring module we choose the top 6 most frequent dives. By examining the scoring performance on the dive types which were not seen during training, we can comment on the generalizability of DML training.

For training the DML module we use 10 splits of the training set, each with 130 videos while for training the score estimation module we use 10 splits, each with 174 training videos. We used 90 test videos which were not included during DML training.

We assign the match/mismatch targets as follows. We first compute the sum of the scores awarded by the three judges. The score awarded by a judge is in the range 0 to 10 and has a deviation of  $\pm 0.5$ . This implies that while judging the same performance a judge may not give exactly the same score every time, but may give a slightly different score with a tolerance of up to  $\pm 0.5$ . Thus the sum of the scores given by the three judges may have a total deviation  $\pm 1.5$ . Since the overall score is computed as the product of the judge’s score

with the difficulty level, which is  $3 \pm 0.3$ , the total deviation in the overall score is  $\pm 0.5 \times 0.33 \approx \pm 5$ . We consider two videos to be similar if their overall scores do not differ by more than  $\pm 5$ . The positive pairs to be used for training the Siamese Metric Learning network are identified as the ones with the two videos having a score difference of less than 5.

Based on this criteria, there are approximately 1300 positive pairs and 1900 negative pairs in the splits. We augment the training videos by applying augmentation techniques like variations in brightness, zooming effect, masking of the background and histogram equalization to give us about 26000 pairs.

The UNLV vaults dataset[15] contains 176 videos with 3 vault types out of which 138 vaults are performed in tuck position while the rest are free and twist vault types. We consider only the 138 tuck vaults for our experiment because the other vault types have inadequate samples for training and testing. Similar to the diving setup, we evaluate the scoring model for 5 train-test split, where the training set comprises 98 videos and test set comprises 40 videos. We train only the scoring module using the vault dataset and use the same DML module that was trained on the diving dataset. The vaults action is considered to show the generalizability of the DML model that has been trained for diving action.

##### B. Performance evaluation metrics

We evaluate the performance of the models using two metrics:

- 1) Rank correlation between the predicted and the ground truth scores : Spearman rank correlation  $\rho$  [16], [15] is used to measure performance. Higher  $\rho$  signifies better rank correlation between the true and predicted scores. This metric allows for non-linear relationship, however, it does not explicitly emphasizes the true score value, but the relative ranking.
- 2) Mean Square Error between the predicted and the ground truth scores: This metric emphasizes that the predicted score should be as close as possible to the score.

##### C. Baseline Works

We compare our works with three [16], [15], [20] baseline techniques for action score estimation. While [16] utilizes pose-based features, the other two techniques [15], [20] use spatio-temporal features. Also, [16], [15] are segmentation free approaches to score estimation i.e. the information about boundaries of the semantic segments is not available, while [20] requires the knowledge of the semantic boundaries. Both variants of the C3D+LSTM models proposed by Parmar *et al.* [15] i.e. final score labeling (F) and incremental training (I) are used in comparison. We use the code provided by the authors to evaluate their performance on our dataset splits.

##### D. Score Estimation Results

###### **Diving Results**

Tables I and II give the average and split-wise Rank correlation and Mean square error results for different baseline models.

TABLE I  
COMPARISON OF RANK CORRELATION; C:C3D, L:LSTM, F: FINAL, I: INCREMENTAL, S: SVR

Splits	Baseline Works						S3D	Dive Specific Regression models	Ours					
	Pirsiavash <i>et al.</i> [16]		Parmar <i>et al.</i> [15]						Xiang <i>et al.</i> [20]			Scoring with Bias		
	Pose + DCT	Pose + DFT	C + L (F)	C + L (I)	C + L + S(F)	C + L + S(I)			C + L (F)	Constant Expert Bias	Worst performer bias per dive	Best performer bias per dive		
1	0.46	0.42	0.65	0.6	0.62	0.46	0.52	0.69	0.67	0.69	0.71			
2	0.3	0.24	0.55	0.49	0.54	0.42	0.39	0.58	0.54	0.53	0.65			
3	0.53	0.42	0.63	0.62	0.65	0.5	0.39	0.71	0.69	0.7	0.71			
4	0.52	0.41	0.66	0.57	0.58	0.42	0.49	0.59	0.7	0.69	0.8			
5	0.43	0.4	0.68	0.58	0.61	0.67	0.42	0.68	0.62	0.66	0.73			
6	0.41	0.44	0.67	0.47	0.63	0.56	0.34	0.69	0.62	0.56	0.67			
7	0.49	0.56	0.65	0.49	0.58	0.49	0.58	0.68	0.64	0.7	0.72			
8	0.35	0.38	0.58	0.58	0.62	0.52	0.55	0.63	0.46	0.52	0.63			
9	0.49	0.52	0.66	0.63	0.66	0.59	0.47	0.65	0.58	0.72	0.65			
10	0.33	0.42	0.67	0.58	0.63	0.52	0.6	0.71	0.53	0.66	0.65			
<b>Average</b>	0.43	0.42	0.64	0.56	0.61	0.51	0.47	0.66	0.6	0.64	<b>0.69</b>			

TABLE II  
COMPARISON OF MEAN SQUARE ERROR; C:C3D, L:LSTM, F: FINAL, I: INCREMENTAL, S: SVR

Splits	Baseline Works						S3D	Dive Specific Regression models	Ours					
	Pirsiavash <i>et al.</i> [16]		Parmar <i>et al.</i> [15]						Xiang <i>et al.</i> [20]			Scoring with Bias		
	Pose + DCT	Pose + DFT	C+L(F)	C+L(I)	C + L + S (F)	C + L + S (I)			C + L (F)	Constant Expert Bias	Dive Specific Worst Performer Bias	Dive Specific Expert Performer Bias		
1	153.59	160.25	68.37	245.59	95.09	232.91	286.84	81.2	89.96	88.81	64.67			
2	212.12	228.48	107.81	252.03	108.04	226.36	192.4	108.26	114.5	98.54	101.31			
3	197.83	185.37	122.88	312.31	116.47	289.27	208.94	102.92	77.6	92.64	77.95			
4	214.02	228.15	104.11	344.2	124.25	383.6	262.73	133.27	100.46	92.96	72.43			
5	206.18	198.42	73.41	276.64	88.87	264.02	215.04	96.04	105.53	84.19	89.18			
6	184.56	169.37	84.11	249.26	108.19	234.83	208.62	72.43	97.18	97.36	85.08			
7	154.4	136.39	99.99	396.42	101.96	386.94	166.47	111.08	112.16	89.88	102.7			
8	153.96	158.76	106.34	223.33	102.89	225.67	160.34	110.68	146.15	137.58	110.17			
9	178.6	179.02	108.73	264.53	94.48	258.98	195.72	119.14	118.01	74.94	95.26			
10	231.25	213.58	74	327.81	113.33	303.89	231.18	80.81	110.17	90.32	87.2			
<b>Average</b>	188.65	185.78	94.97	289.21	105.35	280.65	212.83	101.58	107.18	94.72	<b>88.59</b>			

TABLE III  
EXPERT AND WORST PERFORMER DETAILS FOR DIVING

Dive Type	Expert Score	Count of Expert Performers	Worst Score	Count of Worst Performer
1	92.8	1	51.2	1
2	102.6	2	21.6	1
3	94.05	4	34.65	1
4	99.9	2	66.6	1
5	99.75	1	36.3	1
6	102.6	1	64.8	1

For comparing the scoring model with the other baselines, we evaluate three variants of our trained models where the biases (reference videos) were fixed as follows: 1) an expert bias of the first dive type is included in the pairs for all dive types; 2) dive specific expert performer bias is included in the pairs corresponding to the same dive type 3) dive specific worst performer bias is included in the pairs corresponding to the same dive type. Tables I and II illustrate that the bias from the dive specific expert videos outperforms the other two variants.

Table III gives the number of the best and the worst

performers for the Diving dataset and their respective scores for individual dive types. It is seen that the experts are more as compared to the worst performers. Thus we have around 380 pairs to train the scoring model in case of expert bias and 174 pairs in case of worst performer bias. This leads to an improved efficiency of scoring in case of expert bias as compared to the worst bias even when both lead to coherent Siamese embedding. Further, including a constant bias for the task of scoring gives a poor performance as compared to the case when we use a worst or expert performer as show in Table I and II where we choose expert of dive type 1 as the constant. This trend is followed even when we choose all 4 experts of dive type 3 which gives us around 700 input pairs. We achieve a still lower rank correlation of 0.62 and higher mean square error of 98.64 on average for the 10 split.

Further our proposed approach that constructs a video representation using the dive specific expert reference outperforms the traditional scoring models. Our model has the maximum rank correlation and the minimum mean square error. The improvement in the performance of our assessment model is because we make use of a representation that is specific to the

TABLE IV  
COMPARISON OF DIFFERENT TECHNIQUES : RANK CORRELATION AND MEAN SQUARE ERROR FOR INDIVIDUAL DIVE TYPES

Dive Type	Baseline Works														Ours	
	Pirsiavash <i>et al.</i> [16]				Parmar <i>et al.</i> [15]								Xiang <i>et al.</i> [20]			
	Pose + DCT		Pose + DFT		C + L (F)		C + L (I)		C+L +S (F)		C + L + S(I)		S3D		Best performer bias per dive	
RC	MSE	RC	MSE	RC	MSE	RC	MSE	RC	MSE	RC	MSE	RC	MSE	RC	MSE	
1	0.243	105.92	0.23	123.48	0.62	53.03	0.64	61.73	0.53	<b>41.17</b>	0.49	104.39	0.05	96.58	<b>0.67</b>	44.31
2	0.55	227.98	0.6	189.58	0.69	132.88	0.73	118.21	0.68	157.71	0.68	434.54	0.48	323.13	<b>0.76</b>	<b>105.62</b>
3	0.36	208.55	0.4	188.61	<b>0.69</b>	102.3	0.65	113.53	0.59	98.19	0.56	241.62	0.39	188.06	<b>0.69</b>	<b>95.43</b>
4	0.12	139.84	0.25	156.71	0.24	90.08	0.3	118.3	0.29	105.9	0.23	202.14	0.019	107.79	<b>0.34</b>	<b>81.99</b>
5	0.34	351.72	0.27	356.15	0.79	<b>90.34</b>	0.73	116.56	<b>0.8</b>	107.25	0.7	407.48	0.18	453.48	0.78	104.06
6	0.37	97.89	0.14	100.16	0.43	101.23	0.42	117.38	0.47	130.91	0.41	163.93	0.38	107.92	<b>0.48</b>	<b>97.16</b>

TABLE V  
STUDY OF IMPACT OF METRIC LEARNING ON FINAL SCORING

Model	# DML training pairs	RC	MSE
C+L (F) [15]	NA	0.64	96.89
Without DML With Expert Bias	NA	0.56	116.37
With DML + Expert Bias (without augmentation)	3000	0.57	138.29
With DML + Expert Bias (few augmented videos)	14000	0.626	107.8
With DML (large augmentation)	26000	<b>0.69</b>	<b>88.59</b>

dive type. It is the new embedding that considers an expert bias with the candidate video. We also learned different regression models (C3D + LSTM) for different dive types and it was seen that when compared to the use of bias-specific embeddings (our model), the rank correlation between the predicted scores from the 6 different regression models and the ground-truth scores is lower and the mean square error is higher. This result is indicated in Table II.

We also compare the performance of our model and the baseline models on evaluation of scores for the individual dive types. Table IV illustrates that our model outperforms the baseline models for 5 dive types and performs comparable for the 6<sup>th</sup> dive type too. Thus even when the DML module was not trained for 3 of the 6 dive types and the model has shown good generalization.

Next we check how the distance similarity metric learning impacts the performance of scoring. We skip the first phase of our framework *i.e.* the DML module and directly perform scoring relative to expert. Table V shows that without the DML phase the model gives a bad scoring performance. Following this the DML module is introduced and trained with and without data augmentation. With no augmentation there are too few (1300) pairs and the Metric Learning performance is poorer compared to when we include more pairs using augmentation(26000). Further, a bad Metric Learning adversely impacts the task of scoring. It is seen that with too few DML module training pairs, the performance of the scoring is lower than the original C3D-LSTM [15]. However, with data augmentation there is a positive impact on the performance of scoring and the DML module can identify intra-class variations.

Finally we compare the performance of our model with the

C3D-LSTM model [15] for gymnastic vaults. The other two baselines are not developed for gymnastic vaults.

### Vault Results

We fine-tuned our model and the C3D+LSTM model developed by Parmar *et al.* [15] for diving to that of gymnastic vaults. Table VI shows that our model performs better in terms of both Rank Correlation and Mean Square Error than most of the C3D+LSTM model variations introduced by Parmar *et al.*[15]. This shows that the distance metric learning is generic in terms of actions being compared. Thus our model not only performs better for the dive types not included in the DML training, but also gives better performance for gymnastic vaults after fine tuning.

### E. Feedback Proposals

Here we illustrate the efficacy of the Distance Metric Learning module in predicting the contributions of each clip of the action to the final score. Following the strategy of clip wise similarity computation proposed in the last section, we achieve a reasonably good feedback for the actions that correlate with the Olympics commentaries, which are provided by Parmar *et al.*[1]. Table VII shows the feedback proposals for some videos. We see that the similarity of the clips decreases in consistency with the video’s commentary.

We propose a scheme to compare the performance of our clip-level feedback to that provided by Parmar *et al.* [15]. For the test videos, the commentaries provided in the Multitask AQA diving dataset [1] are interpreted to find the faulty clips in the performance. This is considered as the ground truth for feedback proposals. The clip wise similarity computed from our DML module signify faulty clips when the Sigmoid output is less than 0.5 *i.e.* the clips are not similar to the corresponding clip of an expert execution.

We evaluate the precision and recall for the two systems in identifying the faulty clips. Table VIII lists the comparison results for the two systems. Our system outperforms the LSTM+C3D aggregations proposed by Parmar *et al.* [15]. The high recall of our scheme signifies fewer faulty clips that are missed. On the contrary the lower precision mostly occur due to misalignment of the clips.

### F. Analysis of various embedding

We analyse the LSTM embedding  $O_p$  and the Siamese embedding  $D_2$  (Figure 2, with weights learned during the DML

TABLE VI  
RESULTS OF FINE TUNING THE DIVING MODEL ON GYMNASTIC VAULTS; C: C3D, L: LSTM, F: FINAL LABEL SCORING, I: INCREMENTAL LABEL SCORING, S: SVR

Splits	Parmar <i>et al.</i> [15]								Ours	
	C + L (F)		C + L (I)		C + L + S (F)		C + L + S (I)		Best Performance Bias	
	RC	MSE	RC	MSE	RC	MSE	RC	MSE	RC	MSE
1	0.39	0.58	0.52	0.88	0.47	0.36	0.45	0.77	0.37	0.36
2	0.45	0.62	0.63	0.86	0.55	0.5	0.67	1.04	0.39	0.46
3	0.34	0.67	0.34	0.84	0.44	0.52	0.4	0.83	0.49	0.57
4	0.19	0.88	0.33	0.72	0.16	0.61	0.35	0.64	0.36	0.45
5	0.35	0.68	0.56	0.97	0.34	0.58	0.54	0.92	0.44	0.34
<b>Average</b>	<b>0.34</b>	<b>0.69</b>	<b>0.47</b>	<b>0.85</b>	<b>0.39</b>	<b>0.51</b>	<b>0.48</b>	<b>0.84</b>	<b>0.41</b>	<b>0.44</b>

TABLE VII  
CLIP LEVEL CONTRIBUTION TO THE FINAL SCORE

Video Number	Commentary	Clipwise similarity	Graph Interpretation
214	<eos> eight verse three and a half somersault oh he's a little short on the entry they're beautiful through the year as always is that top Chinese divers see he just didn't kick out early enough and <b>he's not vertical</b> at all although <b>somehow he manages to make such a small splash</b> that's great control underneath the water to pull that bubble down with him like many of the divers here they did take the chance to get a feel of the <eos>		<ul style="list-style-type: none"> <li>• Clip 8 - Entry Clip</li> <li>• Clip 9 - Splash</li> <li>• Similarity decreases on entry and increases in the next clip</li> </ul>
118	<eos> <b>beautiful is back to near his best</b> [Applause] [Music] three point two that's like it he loves to smile for the camera doesn't he that he loves this stage well he is a showman and to be fair even during the difficulties last night he came up smiling and waving at the camera I think he's one of those athletes Roger he just finds it a joy to be here wow I'm at the Olympics he's old he's done it all before but he just thrives on it and in 86 4 for Matthew he picked up another <eos>		The similarities of the clip to the expert clip is more than 0.5 in all cases which indicates that the performance is overall good.
44	<eos> <b>it's a bit of a crash-landing</b> I'm sad to say because Wolfram was having a good competition especially with his opening couple of dives and it's just gone missing for him confidence not there now short on position coming into the water and a well that's counting him out a sixty-four point eight is a disaster in <eos>		<ul style="list-style-type: none"> <li>• Clip 8 - Entry Clip</li> <li>• Clip 9 - Splash</li> <li>• The similarity decreases at the entry and splash clip</li> </ul>
292	<eos> and the crowd gasps and we'll just leave that one to the judges because you saw as well as we did he would have seen that a diver before him duomo in a similar position in terms of rankings make a mistake and he had a chance to capitalize just made him move up that one more spot but if anything <b>he's put in a worst dive</b> 39 6s Gilson so tough <eos>		Initial clips are padded ones and the similarity decreases in the last three clips indicating an overall performance.

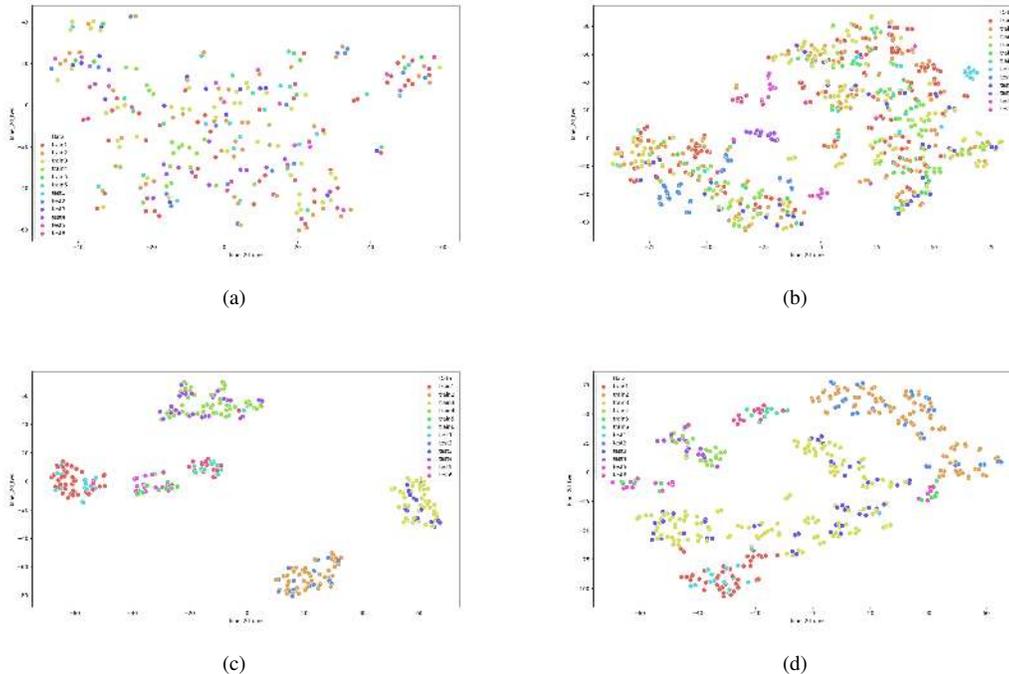


Fig. 5. t-SNE Plots of LSTM embedding and Siamese embedding of train and test data: (a) the LSTM embedding  $O_p$ ; (b) Siamese Embedding  $D_2$  with Constant Bias at  $O_p$ ; (c) Siamese Embedding  $D_2$  with Worst Performer Bias at  $O_p$ ; and, (d) Siamese Embedding  $D_2$  with Expert Performer Bias at  $O_p$

TABLE VIII  
PERFORMANCE COMPARISON OF FEEDBACK PROPOSAL SCHEMES

	Precision	Recall
<b>C3D + LSTM [15]</b>	0.18	0.52
<b>C3D + LSTM + SVR [15]</b>	0.11	0.36
<b>Ours</b>	0.58	0.95

module training. The 2D t-SNE plots of these embeddings for training and test data are plotted to evaluate their coherence across various dive types. Figure 5(a) shows the plots of the LSTM embedding. It is seen that the different dive types do not form clear clusters with the LSTM embedding. Similar non-distinct clusters are observed with the Siamese embedding  $D_2$  when a constant expert of dive type 3 is passed as an input in all the pairs (Figure 5(b)).

However, when the bias  $O_p$  is made dive-specific by either including the worst performer per dive or the best performer per dive, the Siamese embedding are seen to become coherent for individual dive types (Figure 5(c) and 5(d)). Thus an inclusion of a bias  $O_p$  specific to the dive types and using the  $D_2$  embedding for the task of scoring is a good choice, which is exhibited by improved performance as discussed in Table I and II.

## VI. CONCLUSIONS

Human action quality assessment or action scoring has been posed as a regression problem in the past works. The simple regression-based solutions lack interpretability. We introduce a novel approach for action scoring in which the performances are compared with the reference videos to estimate the final

score. A deep metric learning module learns the similarity metric between two videos using the difference in the scores of the videos. This module is then used to find the similarity of the video to the reference video. Such an approach can capture and assess the the temporal variations of videos with the reference video. Our experiments on Olympics Diving and Gymnastic vaults actions show that the proposed approach that includes a bias of the high-rated reference video performs better than the traditional scoring methods. Further, we introduce an unsupervised technique to provide the sub-action level feedback in order to make the scores more explicable. Use of deep learning metric to assess long term actions where alignment can be a major issue is the future direction of our work.

## REFERENCES

- [1] Paritosh Parmar and Brendan Tran Morris. What and how well you performed? a multitask learning approach to action quality assessment. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 304–313, 2019.
- [2] Charles J Ansorge and John K Scheer. International bias detected in judging gymnastic competition at the 1984 olympic games. *Research quarterly for exercise and sport*, 59(2):103–107, 1988.
- [3] Hazel Doughty, Dima Damen, and Walterio Mayol-Cuevas. Who’s better? who’s best? pairwise deep ranking for skill determination. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6057–6066, 2018.
- [4] John W Emerson, Miki Seltzer, and David Lin. Assessing judging bias: An example from the 2000 olympic games. *The American Statistician*, 63(2):124–131, 2009.
- [5] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.
- [6] Hiteshi Jain and Gaurav Harit. A framework to assess sun salutation videos. In *Proceedings of the Tenth Indian Conference on Computer Vision, Graphics and Image Processing*, page 29. ACM, 2016.

- [7] Hiteshi Jain and Gaurav Harit. Detecting missed and anomalous action segments using approximate string matching algorithm. In *Computer Vision, Pattern Recognition, Image Processing, and Graphics: 6th National Conference, NCVPRIPG 2017, Mandi, India, December 16-19, 2017, Revised Selected Papers 6*, pages 101–111. Springer, 2018.
- [8] Yongjun Li, Xiujuan Chai, and Xilin Chen. End-to-end learning for action quality assessment. In *Pacific Rim Conference on Multimedia*, pages 125–134. Springer, 2018.
- [9] Andrew S Gordon. Automated video assessment of human performance. In *Proceedings of AI-ED*, pages 16–19, 1995.
- [10] Marko Jug, Janez Perš, Branko Dežman, and Stanislav Kovačič. Trajectory based assessment of coordinated human activity. In *International Conference on Computer Vision Systems*, pages 534–543. Springer, 2003.
- [11] Matej Perše, Matej Kristan, Janez Perš, and Stanislav Kovačič. *Automatic evaluation of organized basketball activity using bayesian networks*. Citeseer, 2007.
- [12] Hillary N Morgan and Kurt W Rothhoff. The harder the task, the higher the score: Findings of a difficulty bias. *Economic Inquiry*, 52(3):1014–1026, 2014.
- [13] Alejandro Newell, Kaiyu Yang, and Jia Deng. Stacked hourglass networks for human pose estimation. In *European Conference on Computer Vision*, pages 483–499. Springer, 2016.
- [14] Paritosh Parmar and Brendan Tran Morris. Measuring the quality of exercises. *arXiv preprint arXiv:1608.09005*, 2016.
- [15] Paritosh Parmar and Brendan Tran Morris. Learning to score olympic events. In *Computer Vision and Pattern Recognition Workshops (CVPRW), 2017 IEEE Conference on*, pages 76–84. IEEE, 2017.
- [16] Hamed Pirsiavash, Carl Vondrick, and Antonio Torralba. Assessing the quality of actions. In *European Conference on Computer Vision*, pages 556–571. Springer, 2014.
- [17] Ilya Sutskever, Oriol Vinyals, and Quoc V Le. Sequence to sequence learning with neural networks. In *Advances in neural information processing systems*, pages 3104–3112, 2014.
- [18] Du Tran, Lubomir Bourdev, Rob Fergus, Lorenzo Torresani, and Manohar Paluri. Learning spatiotemporal features with 3d convolutional networks. In *Proceedings of the IEEE international conference on computer vision*, pages 4489–4497, 2015.
- [19] Vinay Venkataraman, Ioannis Vlachos, and Pavan K Turaga. Dynamical regularity for action analysis. In *BMVC*, pages 67–1, 2015.
- [20] Xiang Xiang, Ye Tian, Austin Reiter, Gregory D Hager, and Trac D Tran. S3d: Stacking segmental p3d for action quality assessment. In *2018 25th IEEE International Conference on Image Processing (ICIP)*, pages 928–932. IEEE, 2018.
- [21] Yi Yang and Deva Ramanan. Articulated pose estimation with flexible mixtures-of-parts. In *Computer Vision and Pattern Recognition (CVPR), 2011 IEEE Conference on*, pages 1385–1392. IEEE, 2011.
- [22] D Yi, Z Lei, and SZ Li. Deep metric learning for practical person re-identification (2014). *ArXiv e-prints*.
- [23] Aneeq Zia, Yachna Sharma, Vinay Bettadapura, Eric L Sarin, Mark A Clements, and Irfan Essa. Automated assessment of surgical skills using frequency analysis. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 430–438. Springer, 2015.
- [24] Dong Yi, Zhen Lei, Shengcai Liao, and Stan Z Li. Deep metric learning for person re-identification. In *2014 22nd International Conference on Pattern Recognition*, pages 34–39. IEEE, 2014.
- [25] Bilge Soran, Ali Farhadi, and Linda Shapiro. Generating notifications for missing actions: Don’t forget to turn the lights off! In *Proceedings of the IEEE International Conference on Computer Vision*, pages 4669–4677, 2015.
- [26] Chenxia Wu, Jiemi Zhang, Silvio Savarese, and Ashutosh Saxena. Watch-n-patch: Unsupervised understanding of actions and relations. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4362–4370, 2015.
- [27] Hazel Doughty, Dima Damen, and Walterio Mayol-Cuevas. Who’s better? who’s best? pairwise deep ranking for skill determination. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6057–6066, 2018.
- [28] Hazel Doughty, Walterio Mayol-Cuevas, and Dima Damen. The pros and cons: Rank-aware temporal attention for skill determination in long videos. *arXiv preprint arXiv:1812.05538*, 2018.
- [29] Hassan Ismail Fawaz, Germain Forestier, Jonathan Weber, Lhassane Idoumghar, and Pierre-Alain Muller. Evaluating surgical skills from kinematic data using convolutional neural networks. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 214–221. Springer, 2018.
- [30] Zhaofan Qiu, Ting Yao, and Tao Mei. Learning spatio-temporal representation with pseudo-3d residual networks. In *proceedings of the IEEE International Conference on Computer Vision*, pages 5533–5541, 2017.
- [31] Sumit Chopra, Raia Hadsell, Yann LeCun, et al. Learning a similarity metric discriminatively, with application to face verification. In *CVPR (1)*, pages 539–546, 2005.
- [32] Seyma Yucer and Yusuf Sinan Akgul. 3d human action recognition with siamese-1stm based deep metric learning. *arXiv preprint arXiv:1807.02131*, 2018.
- [33] Chengming Xu, Yanwei Fu, Bing Zhang, Zitian Chen, Yu-Gang Jiang, and Xiangyang Xue. Learning to score figure skating sport videos. *IEEE Transactions on Circuits and Systems for Video Technology*, 2019.