Semi-Supervised Teacher-Reference-Student Architecture for Action Quality Assessment

Wulian Yun[®], Mengshi Qi[®], Fei Peng[®], and Huadong Ma[®]

State Key Laboratory of Networking and Switching Technology, Beijing University of Posts and Telecommunications, China {yunwl,qms,pf0607,mhd}@bupt.edu.cn

Abstract. Existing action quality assessment (AQA) methods often require a large number of label annotations for fully supervised learning, which are laborious and expensive. In practice, the labeled data are difficult to obtain because the AQA annotation process requires domainspecific expertise. In this paper, we propose a novel semi-supervised method, which can be utilized for better assessment of the AQA task by exploiting a large amount of unlabeled data and a small portion of labeled data. Differing from the traditional teacher-student network, we propose a teacher-reference-student architecture to learn both unlabeled and labeled data, where the teacher network and the reference network are used to generate pseudo-labels for unlabeled data to supervise the student network. Specifically, the teacher predicts pseudo-labels by capturing high-level features of unlabeled data. The reference network provides adequate supervision of the student network by referring to additional action information. Moreover, we introduce confidence memory to improve the reliability of pseudo-labels by storing the most accurate ever output of the teacher network and reference network. To validate our method, we conduct extensive experiments on three AQA benchmark datasets. Experimental results show that our method achieves significant improvements and outperforms existing semi-supervised AQA methods. Our source code is available at https://github.com/wuli55555/TRS.

Keywords: Action Quality Assessment \cdot Semi-Supervised Learning \cdot Teacher-Reference-Student Architecture

1 Introduction

Action quality assessment (AQA) aims to assess how well the action is performed, which is an important task for video understanding and has been widely used for medical surgery skill assessment, sports activities scoring, etc. In recent years, many existing works [1, 6, 9, 25–27, 29, 47] have achieved significant progress on AQA task, which primarily relies on fully supervised learning and requires a large amount of annotated data for training. However, obtaining annotation labels for the AQA task is unique because the assessment process typically involves domain-specific knowledge, requiring annotators to have the relevant professional background and experience, resulting in a time-consuming and labor-intensive



Fig. 1: Illustrations of (a) fully supervised AQA task, (b) traditional teacher-student network and (c) our semi-supervised AQA task. The fully supervised methods utilize labeled data as inputs for learning to achieve the assessment, while our semi-supervised method leverages labeled data and unlabeled data as inputs to perform the assessment. In addition, differing from the traditional teacher-student network, we have a reference network designed for the AQA task to provide supervision information for the student.

process of label acquisition. Hence, it is important yet challenging to conduct better action assessments using very few labels.

Semi-supervised learning can effectively address the above issue by leveraging both labeled and unlabeled data for training, thereby enhancing the performance of deep learning models. Despite its significant advancements in various computer vision domains, *e.g.*, object detection [13, 23, 37] and action recognition [4, 46], there has been limited exploration of semi-supervised learning in the field of AQA. Therefore, this motivates us to propose a semi-supervised approach for AQA, as shown in Figure 1. Specifically, we leverage the commonly employed teacher-student network [35,39] in semi-supervised learning as the basis, in which the teacher generates pseudo-labels for unlabeled data to supervise the training of the student.

Ideally, the teacher network tends to predict more accurately results than the student network in semi-supervised learning, but there is no guarantee that the teacher's predictions invariably be better than the student's predictions. If the teacher generates low-fidelity pseudo-labels, it will affect the training of the student network and lead to poor assessment performance. Especially for tasks like AQA that require learning subtle variations in the whole action, it remains a great challenge to generate high-quality pseudo-labels.

Inspired by the fact that teachers cannot always guarantee the infallibility of their insights when instructing students, and often need to refer to other knowledge and compare similarities and differences to ensure an improvement in teaching quality, we propose a novel semi-supervised teacher-reference-student architecture for AQA. Specifically, our method mainly contains three networks, *i.e.*, teacher network, student network and reference network. The teacher network predicts quality scores by learning the high-level feature of each video, which first trains with labeled videos, and then generates pseudo-label for unlabeled videos to supervise the student network. The student network has the same structure as the teacher network. The reference network predicts pseudolabels by referring to other action information, thus providing additional supervision information for the student. In addition, considering that teachers have the memory capability, we construct a confidence memory to further improve the accuracy of pseudo-labels. The confidence memory stores the most accurate outputs ever predicted by the teacher network and reference network as output to supervise the student network, in which we introduce confidence to measure the reliability of pseudo-labels.

Our main contributions are summarized as follows:

- We propose a novel semi-supervised AQA method, which effectively leverages the knowledge from unlabeled and labeled data to improve the generalization of assessment.
- We introduce a reference network to provide additional supervision to the student network.
- We design a confidence memory to store the most accurate output predicted by the teacher network and reference network, which can improve the accuracy of pseudo-labels.
- Extensive experimental results on three widely used datasets, *i.e.*, MTL-AQA, Rhythmic Gymnastics and JIGSAWS, demonstrate the effectiveness and superiority of our method.

2 Related Work

Action quality assessment. Action quality assessment is an important task in video understanding [30–32,41] areas. Existing mainstream AQA methods can be categorized into regression-based [19, 25–27, 29, 34] and ranking-based methods [6,40]. Regression-based methods mainly formulate the AQA as a regression task, by employing scores as supervised signals. In early works, Pirisiavash et al. [29] utilize the Support Vector Machines [7] to regress spatio-temporal pose features to scores. Pan et al. [25] learn the detailed joint motion by constructing graphs from both spatial and temporal aspects. Tang et al. [34] adopt score distribution to depict the probability of the AQA score thereby addressing the problem of ambiguity of score labels. Besides, Li et al. [19] propose a pairwise contrastive learning network to learn relative scores by capturing subtle and critical differences between videos. Subsequently, some works [6,40] formulate AQA as a ranking problem to further improve the accuracy of AQA. However, the above methods rely on heavily samples with score labels, and obtaining these score labels is a time-consuming and labor-intensive process. In contrast, our method focuses on achieving effective AQA performance with only a limited number of labeled samples.

Semi-supervised learning. In recent years, semi-supervised learning has been widely used in various fields of computer vision, such as image classification [4,33, 46], semantic segmentation [22,38], object detection [13,21,23,37], etc. Existing

methods mainly divides into two branches, *i.e.* consistency-based [4,13,33,35,44]and pseudo-label based methods [21,23,39]. Consistency-based methods perturb the input data and leverage regularization loss to ensure that networks have consistent predictions for the same data under different augmentations. For example, Meanteacher [35] achieves semi-supervised learning by updating the parameters of the student to the teacher using the exponential moving average. While pseudo-label based methods employ the teacher network to generate pseudo-label predictions for unlabeled data, and then leverage these pseudo-labels to optimize the student network. A representative work is that NoisyStudent [39] utilizes the teacher to generate pseudo-labels for unlabeled data and introduces noise during the learning process to enhance the student's generalization ability. However, so far semi-supervised learning is rarely explored in AQA. Zhang et al. [45] learn the temporal dependence of actions in unlabeled videos by self-supervised segment feature recovery for semi-supervised learning. In contrast, we propose a novel semi-supervised teacher-reference-student architecture. This architecture utilizes the teacher network and reference network to generate pseudo-labels for unlabeled data, providing supervision for student training and facilitating comprehensive learning of unlabeled data.

3 Preliminary

3.1 Teacher-Student Network

Our semi-supervised architecture is based on a teacher-student network. The training process of the teacher-student network is mainly divided into two stages: *burn-in stage* and *teacher-student learning stage*.

1) Burn-in stage: First, the teacher network is pre-trained on the labeled data in burn-in stage, making the teacher network have enough regression ability to provide plausible pseudo-labels. The parameters of the teacher network are then used to initialize the student network.

2) Teacher-student learning stage: During this stage, strong augmentation and weak augmentation are first performed on the input unlabeled data of the student network and the teacher network, respectively. This manner prevents the student network from overfitting based on excessive pseudo-labels. And then the teacher network supervises the training of the student network by making pseudo-label predictions for unlabeled data, where the parameters θ_t of the teacher network are updated from the student's parameters θ_s via exponential moving average (EMA) [35], formulated as follows:

$$\theta_t = \alpha \theta_t + (1 - \alpha) \theta_s,\tag{1}$$

where $\alpha \in (0, 1)$ indicates the momentum. This strategy can prevent the teacher network from overfitting on limited labeled data. Throughout the semi-supervised training, the student network is optimized with labeled and unlabeled data. The loss function optimizes the student network can be formulated as the following:

$$L = L_{sup} + \beta L_{unsup},\tag{2}$$

where β is loss weight, L_{sup} denotes the supervised loss and L_{unsup} indicates the unsupervised loss. More details will be described in the following section.

3.2 Task Definition

In semi-supervised AQA, given a training set D comprising a labeled subset $D^{l} = \{(v_{i}^{l}, s_{i}^{l})\}_{i=1}^{N}$ and a unlabeled subset $D^{u} = \{v_{i}^{u}\}_{i=1}^{M}$. Specifically, v_{i}^{l} and v_{i}^{u} denote labeled video and unlabeled video, respectively. s_{i}^{l} represents the score label of v_{i}^{l} , M and N refer to the number of unlabeled and labeled videos, where $M \gg N$. Semi-supervised AQA aims to simultaneously leverage the limited labeled data D^{l} as well as unlabeled data D^{u} for training, and then assess the quality score for each test video.

4 Method

4.1 Pipeline

The overall framework of our method is illustrated in Figure 2. Specifically, our method consists of three main parts, *i.e.*, teacher network, student network and reference network. The teacher network learns the quality score of each video. The student network and the teacher network have the same architecture, differing mainly in how the weights are updated. The reference network learns the relative scores between different videos to provide additional supervision information for the student.

The training process of our method is mainly as follows: 1) the teacher network and reference network will use labeled videos v_p^l and v_q^l from labeled data D^l for a short period of training during the *burn-in stage*, enabling both the teacher network and the reference network with the ability to assess quality scores. To be specific, the teacher network takes v_p^l as input and directly estimates the score \hat{s}_p^l of v_p^l . The reference network takes a pair of videos $\langle v_p^l, v_q^l \rangle$ as input and estimates a relative score Δs^l between v_p^l and v_q^l . Then, the parameters θ_t of the teacher network are used to initialize the student network θ_s .

2) Followed by the teacher-student learning stage, which in our process we call the *teacher-reference-student learning stage*. The teacher network and the reference network are used to generate pseudo-labels for unlabeled data D^u , thus supervising the training of the student network. In this process, our method takes the unlabeled data v_i^u as well as the video v_i^l with known quality score s_i^l from labeled data D^l , where v_i^u is the input to the teacher network and v_i^u , v_i^l sever as the input to the reference network. We perform strong augmentation (*i.e.*, Gaussian, and RandomHorizontalFlip) on the unlabeled data to the student network and weak augmentation (*i.e.*, RandomHorizontalFlip) to the teacher network. Meanwhile, we introduce confidence memory to guarantee the reliability of pseudo-labels generated by the teacher network and the reference network, to ensure that the pseudo-label output by networks is high-fidelity. The scores of memory in the teacher network and the reference network are averaged as the



Fig. 2: The overall framework of our method. Our method starts with a burn-in stage of training, followed by the teacher-reference-student learning stage. It mainly contains a teacher network, a student network and a reference network. The teacher network and reference network predict pseudo-labels for unlabeled data and supervise the training of the student network. The student network has the same architecture as the teacher. Moreover, confidence memory is used to ensure the reliability of pseudo-labels generated by the teacher network and reference network.

final pseudo-label \bar{s} to supervise the student. The teacher's parameters θ_t are updated by EMA of the student's parameters θ_s as presented in Eq.(1). Next, we will present details about each component.

4.2 Teacher Network and Reference Network

We take the teacher-reference-student learning stage as an example to develop a detailed description of the teacher network and the reference network.

Teacher network aims to directly map the unlabeled video v^u to obtain the quality score \hat{s}^u based on the performance of human action and then supervise student's training. It can be defined as:

$$\hat{\mathbf{s}}^u = \mathbf{A}(v^u),\tag{3}$$

where $\mathbf{A}(\cdot)$ denotes the teacher network. Specifically, for the video v^u , we first follow [3, 40] to split the video into multiple overlapping snippets, where each snippet contains C consecutive frames. Then, I3D [5] is employed as the backbone to extract the feature of each snippet, thus obtaining video feature representation F^u . After that, we utilize the Mixer Layer [36] to extract higher-level stable features by encoding the temporal information of the video sequence. The Mixer Layer is comprised of token-mixing MLP and channel-mixing MLP, and each MLP contains two fully connected layers and a GELU [11]. Concretely, we leverage the feature F^u as input for Mixer Layer. F^u is first fed into token-mixing MLP for interaction between snippets at different temporal positions, and then the feature F^t output by token-mixing MLP are transposed and sent to channel-mixing MLP for independent interactions in each snippet dimension. The process can be formulated as:

$$F^{t} = F^{u} + f_{2}^{t} g(f_{1}^{t}(\text{LN}(F^{u}))),$$

$$F^{c} = F^{t} + f_{2}^{c} g(f_{1}^{c}(\text{LN}(F^{t}))),$$
(4)

where $LN(\cdot)$ represents the LayerNorm [2] operation, \boldsymbol{g} denotes GELU, f_1^t , f_2^t and f_1^c , f_2^c denote the learnable weight matrices of fully-connected layers in *token-mixing* MLPs and *channel-mixing* MLPs, respectively. Finally, the feature F^c generated by Mixer Layers is fed into the regression layer to obtain the quality score \hat{s}^u of the teacher network.

Reference network aims to refer to other action information to improve the regression learning capability of the network and the accuracy of pseudo-labels. The reference network employs the attention mechanism to capture differences and similarities between videos, generating pseudo-labels and providing additional supervision information for the student network. To be specific, given an video v^l with known quality score s^l and an unlabeled video v^u , the reference network aims to estimate the relative score Δs^u of $\langle v^u, v^l \rangle$:

$$\Delta \mathbf{s}^{u} = \Upsilon \left(v^{u}, v^{l} \right), \tag{5}$$

where $\Upsilon(\cdot)$ denotes the reference network, video v^l comes from the labeled set D^l , v^u comes from unlabeled set D^u .

In detail, the reference network first performs feature extraction on video v^l by backbone I3D to obtain feature F^l . Subsequently, we utilize cross-attention learning to learn the relative relationship by distinguishing subtle differences in the temporal level of video pair $\langle v^u, v^l \rangle$. The feature F^u of v^u is treated as a query, and the feature F^l of v^l is served as a key and value, respectively. The learning process can be formulated as follows:

$$Attention(\cdot) = Softmax \left(QK^T / \sqrt{d_k} \right) V,$$

$$F' = Attention(LN(F^u, F^l)) + F^u,$$

$$F'' = MLP (LN(F')) + F',$$
(6)

where Q, K, V are the query, key and value matrices, $LN(\cdot)$ denotes the LayerNorm operation, and d is the dimension of query and key features. Finally, we predict the score Δs^u through the regression layer based on feature F'' of interaction between the video pair $\langle v^u, v^l \rangle$.

4.3 Confidence Memory

To further enhance the accuracy of pseudo-labels, we take into account that teachers have memory capability and thus want to utilize memory to store the best output of the teacher network and reference network. However, how to measure whether the network prediction results are reliable is a key issue. Consider that in object detection [10, 16] and pose estimation [17, 18] tasks, the reliability of outputs usually relies on the confidence measure. Therefore, we propose a confidence memory that stores the best-reliable score output of the teacher network and the reference network by confidence evaluation during the teacher-reference-student learning stage.

Specifically, we set an empty confidence memory M^t , which is used to store the quality score and confidence of each unlabeled data predicted by the teacher network. In each training iteration, the teacher network retrieves videos' confidence $\boldsymbol{\sigma}$ and quality score from memory using a read operation, if the confidence $\boldsymbol{\sigma}$ of the current video prediction is higher than the confidence stored in memory M^t , the score of the current prediction is archived into memory M^t by a write operation. Simultaneously, the confidence $\boldsymbol{\sigma}$ in memory M^t is also updated. Regarding the confidence estimation, we leverage a single variate Gaussian following [10, 17] to generate the probability distribution $P_{\boldsymbol{\Theta}}(\mathbf{x})$ of the predicted score:

$$P_{\Theta}(\mathbf{x}) = \frac{1}{\sqrt{2\pi}\boldsymbol{\sigma}} \exp\left(-\frac{(\mathbf{x}-\boldsymbol{\mu})^2}{2\boldsymbol{\sigma}^2}\right),\tag{7}$$

where Θ represents learnable parameters, μ denotes the predicted score of the teacher network. σ denotes the standard deviation, which is used as a measure of prediction uncertainty. When $\sigma \to 0$, our network is confident in the predicted result. In this way, we construct another confidence memory bank M^r to store high-fidelity pseudo-labels predicted by the reference network. It is worth noting that the predicted scores stored in memory M^t are \hat{s}^u , and the predicted scores in memory M^r are $\hat{s}^u = s^l + \Delta s^u$.

4.4 Optimization and Inference

The final optimization objective of our method contains both supervised loss \mathcal{L}_{sup} and unsupervised loss \mathcal{L}_{unsup} . The supervised loss \mathcal{L}_{sup} is mainly composed of \mathcal{L}_{reg}^s and \mathcal{L}_{reg}^r . The main goal of \mathcal{L}_{reg}^s is to minimize the error between the score s^s predicted by the student network and the ground truth s, and \mathcal{L}_{reg}^r aims to minimize the error between the predicted relative score Δ s of reference network and the corresponding ground truth s, s^l. The loss function of \mathcal{L}_{sup} can be defined as:

$$\mathcal{L}_{sup} = \mathcal{L}_{reg}^{s} + \mathcal{L}_{reg}^{r},$$

$$\mathcal{L}_{reg}^{s} = -\log P_{\Theta}(\mathbf{x}) \Big|_{\mathbf{x}=\mathbf{s}} \propto \log \hat{\boldsymbol{\sigma}} + \frac{(\mathbf{s} - \mathbf{s}^{s})^{2}}{2\hat{\boldsymbol{\sigma}}^{2}},$$

$$\mathcal{L}_{reg}^{r} = -\log P_{\Theta}(\mathbf{x}) \Big|_{\mathbf{x}=\Delta \mathbf{s}} \propto \log \hat{\boldsymbol{\sigma}} + \frac{(\Delta \mathbf{s} - |\mathbf{s} - \mathbf{s}^{l}|)^{2}}{2\hat{\boldsymbol{\sigma}}^{2}}.$$
(8)

The unsupervised loss \mathcal{L}_{unsup} can be formulated as:

$$\mathcal{L}_{unsup} = -\log P_{\Theta}(\mathbf{x}) \Big|_{\mathbf{x}=\bar{\mathbf{s}}} \propto \log \hat{\boldsymbol{\sigma}} + \frac{(\bar{\mathbf{s}} - \mathbf{s}^s)^2}{2\hat{\boldsymbol{\sigma}}^2}, \tag{9}$$

where s^s denote the quality score predicted by the student network, \bar{s} is the final pseudo-label provided by the teacher network and reference network. Finally, the optimization object for our method can be summarized as:

$$L_{total} = L_{sup} + \beta L_{unsup},\tag{10}$$

where β denotes loss weight. During the inference stage, we leverage only the student network to predict quality scores for test videos. Overall, the complete training process of our method is described in Algorithm 1.

Algorithm 1: The training process of our method.

	Input: Labeled data $(v^l, s^l) \in D^l$, unlabeled data $v^u \in D^u$, burn-in step b,							
	maximum iteration E .							
	Output: Teacher model θ_t , Student model θ_s .							
1	1 for $e < E$ do							
2	if $e < b$ then							
3	Predict scores \hat{s}^l and Δs^l of v^l by Teacher θ_t and Reference θ_f ;							
4	Calculate L_{sup} and Update θ_t , θ_f by \mathcal{L}_{sup} in Eq.(8);							
5	end							
6	if $e = b$ then							
7	Initialize θ_s with θ_t							
8	end							
9	if $e > b$ then							
10	Predict pseudo-label \bar{s} of v^u by Teacher θ_t and Reference θ_f ;							
11	Predict score s^s of v^u by Student θ_s ;							
12	Calculate L_{unsup} with (\bar{s}, s^s) ;							
13	Calculate L_{sup} with D^l ;							
14	Update θ_s and θ_f by Eq.(10) and Eq.(8);							
15	Update θ_t by θ_s with EMA in Eq.(1)							
16	end							
17	end							
18	s Return θ_t and θ_s							

5 Experiments

5.1 Datasets

We conduct our experiments on the three commonly-used benchmark datasets, including MTL-AQA [27], Rhythmic Gymnastics [42] and JIGSAWS [8].

MTL-AQA contains 1,412 diving videos, where videos are collected from 16 different international competitions. Following the previous work [19,34,40], we use 1,059 training videos to train our model and 353 test videos for evaluation. Rhythmic Gymnastics comprises 1,000 videos from international competitions, consisting of four gymnastics routines, *i.e.*, ball, ribbon, hoop, and clubs. Each routine has 250 videos, each of which is approximately 1.6 minutes in length

9

and has a frame rate of 25 frames. Following [42,45], we train separate models for each kind, with 200 videos used for training and 50 videos used for evaluation. **JIGSAWS** is a surgical activity dataset, which consists of 103 videos from three different types of surgical activities: Suturing, Needle Passing and Knot Tying. To assess different aspects of surgical actions, each video sample is annotated with multiple scores and the final score is the sum of all annotations.

5.2 **Evaluation Metric**

Following [19, 34, 40], we exploit the widely-adopted Spearman's rank correlation (Sp.Corr) ρ as evaluation metric to evaluate the performance of our method. Spearman's rank correlation mainly calculates the difference between groundtruth scores and predicted scores, which ranges from -1 to 1 and the higher value of ρ indicates better performance. ρ can be defined as:

$$\rho = \frac{\sum_{i} (x_{i} - \bar{x}) (y_{i} - \bar{y})}{\sqrt{\sum_{i} (x_{i} - \bar{x})^{2} \sum_{i} (y_{i} - \bar{y})^{2}}}$$
(11)

where x and y denote the ranking of two series, respectively.

5.3**Implementation Details**

We implement our method with PyTorch [28] and all models are trained by Adam optimizer [15]. We utilize the I3D [5] backbone pre-trained on Kinetics [14] as the feature extractor. The total training epoch is set as 150. For MTL-AQA [27] dataset, we sample 103 frames for each video following [40] and segment them into 10 overlapping snippets each containing 16 continuous frames with a stride of 10 frames. The learning rate is 2e-4. For Rhythmic Gymnastics [42] dataset, following previous work [45], we uniformly sample 2 frames per second from each original video, resulting in a sampled video of about 180 frames. Then, we divide the sampled video into 18 snippets. The learning rate is set to 4e-4. For JIGSAWS [8] dataset, 160 frames are extracted from each video, which are divided into 10 overlapping snippets, each containing 16 consecutive frames. The learning rate is 1e-4. Four-fold cross-validation is performed during the experiment by following [34, 45]. All the video frames are resized to 224×224 . The weight β of unsupervised loss is updated with training epoch t following an exponential warming-up function [20]: $\lambda(t) = 0.2 \times e^{-5(1-t/200)^2}$.

5.4Comparison with State-of-the-Art Methods

We compare our method with existing state-of-the-art AQA methods on three public benchmarks, the results are shown in Table 1, Table 2 and Table 3. **Results on MTL-AQA dataset.** We first compare our method with SOTA fully supervised methods (*i.e.*, SVR [29], C3D-AVG-STL [27], C3D-AVG-MTL [27], USDL [34], CoRe [40] and TPT [3]) and semi-supervised methods (*i.e.*, COREG [48],

Mothods	Semi-supervision		
Methods	10%	40%	
SVR* [29]	0.427	0.565	
C3D-AVG-STL* [27]	0.561	0.632	
C3D-AVG-MTL* [27]	0.584	0.656	
USDL^* [34]	0.530	0.646	
$CoRe^*$ [40]	0.773	0.894	
TPT^* [3]	0.723	0.781	
COREG [48]	0.487	0.526	
Pseudo-labels [12]	0.622	0.716	
VAT [24]	0.635	0.724	
$S^{4}L$ [43]	0.621	0.721	
S^4 AQA [45]	0.676	0.746	
Ours	0.825	0.901	

 Table 1: Comparison with existing AQA methods on MTL-AQA dataset. * denotes the fully supervised method. The best results are highlighted in bold.

Pseudo-labels [12] VAT [24], S^4L [43] and S^4AQA [45]) on MTL-AQA dataset. Our method utilizes the student network to directly predict scores without additional data. In contrast, CoRe and TPT employ a complex multi-exemplar voting strategy that requires multiple additional training data to be introduced for testing. To ensure fairness, we adapt their testing strategy to randomly introduce a video from the training set classified as unlabeled for testing, and randomly select video for training. As shown in Table 1, our method can obtain obvious gains over existing semi-supervised methods while also surpassing fully supervised methods. On the MTL-AQA dataset with 40% labels, our method outperforms the results of semi-supervised method S^4AQA [45] by 20.8%. Also, compared with S^4AQA on only 10% label, the performance gains by our method are significant, *i.e.*, + 22.0%. In summary, the results demonstrate that our method can sufficiently utilize both unlabeled and labeled data to improve the performance of AQA.

Results on Rhythmic Gymnastic dataset. Our method achieves state-ofthe-art results in four categories of the Rhythmic Gymnastic dataset with 40% of labeled data, shown in Table 2. We can find that our method outperforms semisupervised method S^4 AQA [45] by 0.187 and fully supervised method ACTION-NET [42] by 0.211, demonstrating the effectiveness of our method.

Results on JIGSAWS dataset. Table 3 shows the evaluation results on JIG-SAWS dataset with 50% of labeled data. From the table, our model achieves competitive performance compared to other methods. This shows that we can achieve better results not only in the sports domain but also in the surgical domain, proving the generalizability of our method.

5.5 Ablation Study

The impact of different components. We evaluate the effect of different components of our method, the results are shown in Table 4. We set up the

Mathada	Rhythmic Gymnastic(40%)				
Methods	Ball	Clubs	Hoop	Ribbon	AVG
SVR^* [29]	0.175	0.243	0.261	0.309	0.248
ACTION-NET [*] [42]	0.196	0.403	0.319	0.305	0.308
COREG [48]	0.230	0.338	0.331	0.268	0.292
Pseudo-labels [12]	0.183	0.330	0.346	0.305	0.292
VAT [24]	0.208	0.355	0.345	0.292	0.301
$S^{4}L$ [43]	0.209	0.325	0.324	0.290	0.288
S^4 AQA [45]	0.248	0.388	0.307	0.357	0.342
Ours	0.558	0.518	0.567	0.475	0.529

Table 2: Comparison with existing AQA methods on Rhythmic Gymnastic dataset. * denotes the fully supervised method. The best results are highlighted in bold.

Table 3: Comparison with existing AQA methods on JIGSAWS dataset. * denotes the fully supervised method. The best results are highlighted in bold.

Mothoda	JIGSAWS(50%)			
Wiethous	Suturing	Needle Passing	Knot Tying	AVG
USDL* [34]	0.439	0.351	0.680	0.505
$CoRe^*$ [40]	0.622	0.776	0.753	0.717
Pseudo-labels [12]	0.445	0.501	0.714	0.566
VAT [24]	0.524	0.526	0.749	0.612
$S^{4}L$ [43]	0.455	0.529	0.730	0.585
S^4 AQA [45]	0.533	0.552	0.813	0.655
Ours	0.696	0.810	0.753	0.753

traditional teacher-student network in Sec 3.1 as the baseline, and then gradually add confidence memory, and reference network to the baseline. Where 1) *Base* denotes the traditional teacher-student network; 2) Base+TM denotes add confidence memory into the teacher network of 1); 3) Base+RN denotes add a reference network into 1); 4) Base+RN+TM denotes add confidence memory in teacher network of 3); 5) *Ours* denotes add confidence memory into the reference network of 4), *i.e.*, our full model. From the results, it can be observed that with the increase of different components, the assessment results gradually increase with the use of different proportions of labeled data, which is enough to prove the importance of different components.

The impact of memory updating strategy. To validate the effectiveness of the updating memory by confidence measure in our method, we explore another memory updating strategy, *i.e.*, direct updating based on the output of the current epoch, and the experimental results are shown in Base+RN of Table 4. We can find the direct update strategy underperforms the model with the confidence-based updating (*i.e.*, Base+RN+TM). The results can prove that updating the memory according to the confidence measure is effective.

Sensitivity to different proportions of labeled data. To verify the sensitivity of our method when using different proportions of labeled data, we conducted experiments with 10%, 20%, 30% and 40% of labeled data, and the results are shown in Table 5. As seen from the table, the prediction performance of our method gradually increases with the increase of labeled data. Meanwhile, we

Table 4: Ablation study on each component of our method on MTL-AQA dataset. 'Base' indicates baseline, 'RN' denotes the reference network, and 'TM' and 'RM' represent confidence memory of the teacher network and the reference network, respectively.



Fig. 3: Scatter plot comparison results of score prediction between baseline and our method on MTL-AQA dataset. The black line indicates the ground truth, while the pink points represent the predicted scores.

have found that our proposed method with only 10% of labeled data outperforms the S^4 AQA [45] with 40% labeled data in Table 1, and these findings are sufficient to show that our proposed method can effectively learn the rich representations from the unlabeled video.

5.6 Qualitative Analysis

To further intuitively prove the effectiveness of our proposed method, we visualize the representation of the feature distribution of our method and baseline on 10% labeled data by using a scatter plot, as shown in Figure 3. We observe that the prediction results of our model can achieve better alignment with the ground truth compared to the baseline, and these results indicate that our method by introducing the reference network and confidence memory can provide more accurate pseudo-labels to supervise the student network.

In addition, we show the results of four test examples in Figure 4, and we show the predicted scores of the baseline, our proposed method, and ground truth. The baseline and our method's models are trained on 10% of the labeled data. As can be seen from the figure, the error between the predicted score of our method and the ground truth is smaller than the error between the baseline and the ground truth, which can prove the superiority of the assessment performance of our method.



 Table 5: The experiment results of different proportions of labeled data.

Fig. 4: Case study with qualitative results on four video samples from the MTL-AQA dataset, which presents the predicted scores comparisons of baseline, our method and ground truth (GT). 'Baseline' means the traditional teacher-student network.

6 Conclusion

In this paper, we propose a novel semi-supervised teacher-reference-student architecture for the AQA task, which can leverage both unlabeled and labeled data to improve the generalization of assessment. Specifically, the teacher network predicts pseudo-labels for unlabeled data thereby supervising the student network. The reference network provides additional supervision information for the student by referring to other actions. Besides, the confidence memory is introduced to further enhance the accuracy of pseudo-labels generated by the teacher network and reference network. Extensive experiments validate the effectiveness and superiority of our method over the other existing methods.

Acknowledgements

This work is partly supported by the Funds for the Innovation Research Group Project of the NSFC under Grant 61921003, the NSFC Project under Grant 62202063, Beijing Natural Science Foundation (No.L243027) and 111 Project under Grant B18008.

References

- An, Q., Qi, M., Ma, H.: Multi-stage contrastive regression for action quality assessment. In: Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing. pp. 4110–4114 (2024)
- 2. Ba, J.L., Kiros, J.R., Hinton, G.E.: Layer normalization. arXiv preprint arXiv:1607.06450 (2016)
- Bai, Y., Zhou, D., Zhang, S., Wang, J., Ding, E., Guan, Y., Long, Y., Wang, J.: Action quality assessment with temporal parsing transformer. In: Proceedings of the European Conference on Computer Vision. pp. 422–438 (2022)
- Berthelot, D., Carlini, N., Goodfellow, I., Papernot, N., Oliver, A., Raffel, C.A.: Mixmatch: A holistic approach to semi-supervised learning. Advances in Neural Information Processing Systems 32, 5050–5060 (2019)
- Carreira, J., Zisserman, A.: Quo vadis, action recognition? a new model and the kinetics dataset. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 4724–4733 (2017)
- Doughty, H., Damen, D., Mayol-Cuevas, W.: Who's better? who's best? pairwise deep ranking for skill determination. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 6057–6066 (2018)
- Drucker, H., Burges, C.J., Kaufman, L., Smola, A., Vapnik, V.: Support vector regression machines. Advances in Neural Information Processing Systems 9, 155– 161 (1996)
- Gao, Y., Vedula, S.S., Reiley, C.E., Ahmidi, N., Varadarajan, B., Lin, H.C., Tao, L., Zappella, L., Béjar, B., Yuh, D.D., Chen, C.C.G., Vidal, R., Khudanpur, S., Hager, G.: Jhu-isi gesture and skill assessment working set (jigsaws): A surgical activity dataset for human motion modeling. In: Proceedings of the Modeling and Monitoring of Computer Assisted Interventions. vol. 3, p. 3 (2014)
- Gedamu, K., Ji, Y., Yang, Y., Shao, J., Shen, H.T.: Fine-grained spatio-temporal parsing network for action quality assessment. IEEE Transactions on Image Processing 32, 6386–6400 (2023)
- He, Y., Zhu, C., Wang, J., Savvides, M., Zhang, X.: Bounding box regression with uncertainty for accurate object detection. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 2888–2897 (2019)
- 11. Hendrycks, D., Gimpel, K.: Gaussian error linear units (gelus). arXiv preprint arXiv:1606.08415 (2016)
- Hou, P., Geng, X., Huo, Z.W., Lv, J.Q.: Semi-supervised adaptive label distribution learning for facial age estimation. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 31, pp. 2015–2021 (2017)
- 13. Jeong, J., Lee, S., Kim, J., Kwak, N.: Consistency-based semi-supervised learning for object detection. Advances in Neural Information Processing Systems **32** (2019)
- Kay, W., Carreira, J., Simonyan, K., Zhang, B., Hillier, C., Vijayanarasimhan, S., Viola, F., Green, T., Back, T., Natsev, P., et al.: arXiv preprint arXiv:1705.06950 (2017)
- Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980 (2014)
- Lee, Y., Hwang, J.W., Kim, H.I., Yun, K., Kwon, Y., Bae, Y., Hwang, S.J.: Localization uncertainty estimation for anchor-free object detection. In: Proceedings of the European Conference on Computer Vision Workshops. pp. 27–42 (2023)
- Li, C., Lee, G.H.: Generating multiple hypotheses for 3d human pose estimation with mixture density network. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 9879–9887 (2019)

- 16 W. Yun et al.
- Li, J., Bian, S., Zeng, A., Wang, C., Pang, B., Liu, W., Lu, C.: Human pose regression with residual log-likelihood estimation. In: Proceedings of IEEE/CVF International Conference on Computer Vision. pp. 11005–11014 (2021)
- Li, M., Zhang, H.B., Lei, Q., Fan, Z., Liu, J., Du, J.X.: Pairwise contrastive learning network for action quality assessment. In: Proceedings of the European Conference on Computer Vision. pp. 457–473 (2022)
- 20. Liu, Y., Zhu, L., Pei, S., Fu, H., Qin, J., Zhang, Q., Wan, L., Feng, W.: From synthetic to real: Image dehazing collaborating with unlabeled real data. In: Proceedings of the 29th ACM international conference on multimedia. pp. 50–58 (2021)
- Liu, Y.C., Ma, C.Y., Kira, Z.: Unbiased teacher v2: Semi-supervised object detection for anchor-free and anchor-based detectors. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 9809–9818 (2022)
- Liu, Y., Tian, Y., Chen, Y., Liu, F., Belagiannis, V., Carneiro, G.: Perturbed and strict mean teachers for semi-supervised semantic segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 4248–4257 (2022)
- Mi, P., Lin, J., Zhou, Y., Shen, Y., Luo, G., Sun, X., Cao, L., Fu, R., Xu, Q., Ji, R.: Active teacher for semi-supervised object detection. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 14462– 14471 (2022)
- Miyato, T., Maeda, S.i., Koyama, M., Ishii, S.: Virtual adversarial training: a regularization method for supervised and semi-supervised learning. IEEE Transactions on Pattern Analysis and Machine Intelligence 41(8), 1979–1993 (2018)
- Pan, J.H., Gao, J., Zheng, W.S.: Action assessment by joint relation graphs. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 6331–6340 (2019)
- Parmar, P., Morris, B.T.: Learning to score olympic events. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops. pp. 20–28 (2017)
- Parmar, P., Morris, B.T.: What and how well you performed? a multitask learning approach to action quality assessment. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 304–313 (2019)
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., et al.: Pytorch: An imperative style, highperformance deep learning library. Advances in Neural Information Processing Systems 32, 8026–8037 (2019)
- 29. Pirsiavash, H., Vondrick, C., Torralba, A.: Assessing the quality of actions. In: Proceedings of the European Conference on Computer Vision. pp. 556–571 (2014)
- Qi, M., Qin, J., Li, A., Wang, Y., Luo, J., Van Gool, L.: stagnet: An attentive semantic rnn for group activity recognition. In: Proceedings of the European Conference on Computer Vision. pp. 101–117 (2018)
- Qi, M., Qin, J., Yang, Y., Wang, Y., Luo, J.: Semantics-aware spatial-temporal binaries for cross-modal video retrieval. IEEE Transactions on Image Processing 30, 2989–3004 (2021)
- Qi, M., Wang, Y., Li, A., Luo, J.: Stc-gan: Spatio-temporally coupled generative adversarial networks for predictive scene parsing. IEEE Transactions on Image Processing 29, 5420–5430 (2020)
- Sohn, K., Berthelot, D., Carlini, N., Zhang, Z., Zhang, H., Raffel, C.A., Cubuk, E.D., Kurakin, A., Li, C.L.: Fixmatch: Simplifying semi-supervised learning with consistency and confidence. Advances in Neural Information Processing Systems 33, 596–608 (2020)

- Tang, Y., Ni, Z., Zhou, J., Zhang, D., Lu, J., Wu, Y., Zhou, J.: Uncertainty-aware score distribution learning for action quality assessment. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 9836– 9845 (2020)
- Tarvainen, A., Valpola, H.: Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. Advances in Neural Information Processing Systems 30, 1195–1204 (2017)
- Tolstikhin, I.O., Houlsby, N., Kolesnikov, A., Beyer, L., Zhai, X., Unterthiner, T., Yung, J., Steiner, A., Keysers, D., Uszkoreit, J., et al.: Mlp-mixer: An all-mlp architecture for vision. Advances in Neural Information Processing Systems 34, 24261–24272 (2021)
- 37. Wang, X., Yang, X., Zhang, S., Li, Y., Feng, L., Fang, S., Lyu, C., Chen, K., Zhang, W.: Consistent-teacher: Towards reducing inconsistent pseudo-targets in semi-supervised object detection. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 3240–3249 (2023)
- Wang, Y., Wang, H., Shen, Y., Fei, J., Li, W., Jin, G., Wu, L., Zhao, R., Le, X.: Semi-supervised semantic segmentation using unreliable pseudo-labels. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 4238–4247 (2022)
- Xie, Q., Luong, M.T., Hovy, E., Le, Q.V.: Self-training with noisy student improves imagenet classification. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 10684–10695 (2020)
- Yu, X., Rao, Y., Zhao, W., Lu, J., Zhou, J.: Group-aware contrastive regression for action quality assessment. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 7899–7908 (2021)
- Yun, W., Qi, M., Wang, C., Ma, H.: Weakly-supervised temporal action localization by inferring salient snippet-feature. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 38, pp. 6908–6916 (2024)
- 42. Zeng, L.A., Hong, F.T., Zheng, W.S., Yu, Q.Z., Zeng, W., Wang, Y.W., Lai, J.H.: Hybrid dynamic-static context-aware attention network for action assessment in long videos. In: Proceedings of the 28th ACM international conference on multimedia. pp. 2526–2534 (2020)
- Zhai, X., Oliver, A., Kolesnikov, A., Beyer, L.: S4l: Self-supervised semi-supervised learning. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 1476–1485 (2019)
- 44. Zhang, B., Wang, Y., Hou, W., Wu, H., Wang, J., Okumura, M., Shinozaki, T.: Flexmatch: Boosting semi-supervised learning with curriculum pseudo labeling. Advances in Neural Information Processing Systems 34, 18408–18419 (2021)
- Zhang, S.J., Pan, J.H., Gao, J., Zheng, W.S.: Semi-supervised action quality assessment with self-supervised segment feature recovery. IEEE Transactions on Circuits and Systems for Video Technology **32**(9), 6017–6028 (2022)
- 46. Zhao, Z., Zhou, L., Duan, Y., Wang, L., Qi, L., Shi, Y.: Dc-ssl: Addressing mismatched class distribution in semi-supervised learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 9747– 9755 (2022)
- Zhou, K., Ma, Y., Shum, H.P.H., Liang, X.: Hierarchical graph convolutional networks for action quality assessment. IEEE Transactions on Circuits and Systems for Video Technology 33(12), 7749–7763 (2023)
- Zhou, Z.H., Li, M.: Semi-supervised regression with co-training. In: Proceedings of the 19th International Joint Conference on Artificial Intelligence. pp. 908–913 (2005)