

Article

# VI-Net—View-Invariant Quality of Human Movement Assessment

Faegheh Sardari <sup>1,\*</sup> , Adeline Paiement <sup>2</sup> , Sion Hannuna <sup>1</sup> and Majid Mirmehdi <sup>1</sup> 

<sup>1</sup> Department of Computer Science, University of Bristol, Bristol BS8 1UB, UK; sh1670@bristol.ac.uk (S.H.); m.mirmehdi@bristol.ac.uk (M.M.)

<sup>2</sup> Université de Toulon, Aix Marseille Univ, CNRS, LIS, Marseille, France; adeline.paiement@univ-tln.fr

\* Correspondence: faegheh.sardari@bristol.ac.uk; Tel.: +44-(0)117-954 5139

Received: 11 August 2020; Accepted: 9 September 2020; Published: 15 September 2020

**Abstract:** We propose a view-invariant method towards the assessment of the quality of human movements which does not rely on skeleton data. Our end-to-end convolutional neural network consists of two stages, where at first a view-invariant trajectory descriptor for each body joint is generated from RGB images, and then the collection of trajectories for all joints are processed by an adapted, pre-trained 2D convolutional neural network (CNN) (e.g., VGG-19 or ResNeXt-50) to learn the relationship amongst the different body parts and deliver a score for the movement quality. We release the only publicly-available, multi-view, non-skeleton, non-mocap, rehabilitation movement dataset (QMAR), and provide results for both cross-subject and cross-view scenarios on this dataset. We show that VI-Net achieves average rank correlation of 0.66 on cross-subject and 0.65 on unseen views when trained on only two views. We also evaluate the proposed method on the single-view rehabilitation dataset KIMORE and obtain 0.66 rank correlation against a baseline of 0.62.

**Keywords:** movement analysis; view-invariant convolutional neural network (CNN); health monitoring

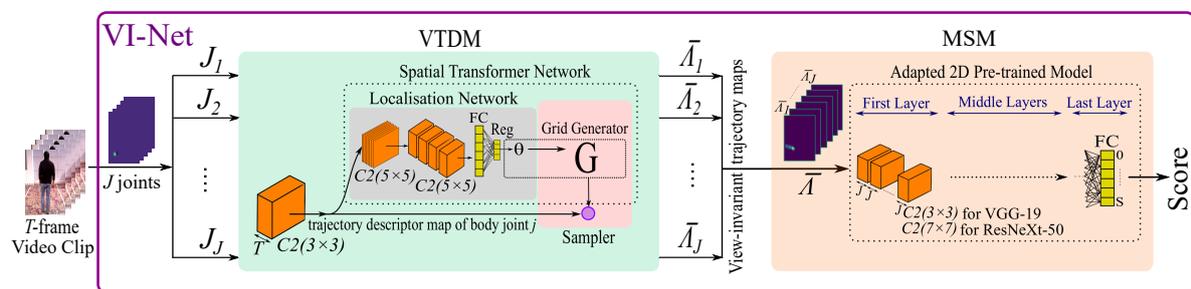
## 1. Introduction

Beyond the realms of action detection and recognition, action analysis includes the automatic assessment of the quality of human action or movement, for example, in sports action analysis [1–4], skill assessment [5,6], and patient rehabilitation movement analysis [7,8]. For example, in the latter application, clinicians observe patients performing specific actions in the clinic, such as walking or sitting-to-standing, to establish an objective marker for their level of functional mobility. By automating such mobility disorder assessment using computer vision, health service authorities can decrease costs, reduce hospital visits, and diminish the variability in clinicians' subjective assessment of patients.

Recent RGB (red, green, blue) based action analysis methods, such as References [2–4,6], are not able to deal with view-invariance when applied to viewpoints significantly different to their training data. To achieve some degree of invariance, some works such as References [7–13], have made use of 3D human pose obtained from (i) Kinect, (ii) motion capture, or (iii) 3D pose estimation methods. Although the Kinect can provide 3D pose efficiently in optimal conditions, it is dependent on several parameters, including distance and viewing direction between the subject and the sensor. Motion capture systems (mocaps) tend to be highly accurate and view-invariant, but obtaining 3D pose by such means is expensive and time consuming, since it requires specialist hardware, software, and setups. These make mocaps unsuitable for use in unconstrained home or clinical or sports settings. Recently, many deep learning methods, for example, References [14–18], have been proposed to extract 3D human pose from RGB images. Such methods (a) either do not deal with view-invariance and are trained from specific views on their respective datasets (for example, References [14,17] show that their methods fail when they apply them on poses and view angles which are different from their training

sets), (b) or if they handle view-invariance, such as References [19,20], then they need multiple views for training.

To the best of our knowledge, there is no existing *RGB-based, view-invariant* method that assesses the quality of human movement. We argue here that using temporal pose information from RGB, can be repurposed, instead of skeleton points, for view-invariant movement quality assessment. In the proposed end-to-end View-Invariant Network (VI-Net in Figure 1), we stack temporal heatmaps of each body joint (obtained from OpenPose [21]) and feed them into our view-invariant trajectory descriptor module (VTDM). This applies a 2D convolution layer that aggregates spatial poses over time to generate a trajectory descriptor map per body joint, which is then forged to be view-invariant by deploying the Spatial Transformer Network [22]. Next, in our movement score module (MSM), these descriptor maps for all body joints are put through an adapted pre-trained 2D convolution model, such as VGG-19 [23] or ResNeXt-50 [24], to learn the relationship amongst the joint trajectories and estimate a score for the movement. Note, OpenPose has been trained on 2D pose datasets which means that our proposed method implicitly benefits from joint labelling.



**Figure 1.** VI-Net has an view-invariant trajectory descriptor module (VTDM) and a movement score module (MSM) where the classifier output corresponds to a quality score.

Initially, we apply our method to a new dataset, called QMAR (dataset and code can be found at <https://github.com/fsardari/VI-Net>), that includes multiple camera views of subjects performing both normal movements and simulated Parkinsons and Stroke ailments for walking and sit-to-stand actions. We provide cross-subject and cross-view results on this new dataset. Recent works such as References [25–28], provide cross-view results only when their network is trained on multiple views. As recently noted by Varol et al. [29], a highly challenging scenario in view-invariant action recognition would be to obtain cross-view results by training from only one viewpoint. While we present results using a prudent set of two viewpoints only within a multi-view training scenario, we also rise to the challenge to provide cross-view results by training solely from a single viewpoint. We also present results on the single-view rehabilitation dataset KIMORE [30] which provides 5 different types of lower back exercises in real patients suffering from Parkinsons, Stroke, and back pain.

This work makes a number of contributions. We propose the first view-invariant method to assess quality of movement from RGB images and our approach does not require any knowledge about viewpoints or cameras during training or testing. Further, it is based on 2D convolutions only which is computationally cheaper than 3D temporal methods. We also present an RGB, multi-view, rehabilitation movement assessment dataset (QMAR) to both evaluate the performance of the proposed method and provide a benchmark dataset for future view-invariant methods.

The rest of the paper is organized as follows. We review the related works in Section 2 and our QMAR dataset in Section 3. Then, we present our proposed network in Section 4 and experimental results in Section 5. Finally, in Section 6, we conclude our work, discuss some of its limitations, and provide directions for future research.

## 2. Related Work

Action analysis has picked up relative pace only in recent years with the majority of works covering one of either physical rehabilitation, sport scoring, or skill assessment [13]. Here, we first consider example non-skeleton based methods (which are mainly on sport scoring), and then review physical rehabilitation methods as it is the main application focus of our work. Finally, given the lack of existing view-invariant movement analysis techniques, we briefly reflect on related view-invariant action recognition approaches.

**Non-Skeleton Movement Analysis**—A number of works have focused on scoring sports actions. Pirsiavash et al. [31] proposed a support vector machine (SVM) based method, trained on spatio-temporal features of body poses, to assess the quality of diving and figure-skating actions. Although their method estimated action scores better than human non-experts, it was less accurate than human expert judgments. More recently, deep learning methods have been deployed to assess the quality of sport actions in RGB-only data, such as References [1–4,32,33]. For example, Li et al. [1] divided a video into several clips to extract their spatio-temporal features by differently weighted C3D [34] networks and then concatenated the features for input to another C3D network to predict action scores. Parmar and Morris presented a new dataset and also used a C3D network to extract features for multi-task learning [3].

The authors of References [4,33] propose I3D [35] based methods to analyse human movement. Pan et al. [4] combine I3D features with pose information by building joint relation graphs to predict score movement. Tang et al. [33] proposed a novel loss function which addresses the intrinsic score distribution uncertainty of sport actions in the decisions by different judges. The use of 3D convolutions imposes a hefty memory and computational burden, even for a relatively shallow model, which we avoid in our proposed method. Furthermore, the performance of these methods are expected to drop significantly when they are applied on a different viewpoint since they are trained on appearance features which change drastically in varying viewpoints.

**Rehabilitation Movement Assessment**—Several works have focussed on such movement assessment, for example, References [7–10,36,37]. For example, Crabbe et al. [9] proposed a CNN network to map a depth image to high-level pose in a manifold space made from skeleton data. Then, the high level poses were employed by a statistical model to assess quality of movement for walking on stairs. In Reference [7], Sardari et al. extended the work in Reference [9] by proposing a ResNet-based model to estimate view-invariant high-level pose from RGB images where the high-level pose representation was derived from 3D mocap data using manifold learning. The accuracy of their proposed method was good when training was performed from all views, but dropped significantly on unseen views.

Liao et al. [8] proposed a long short term memory (LSTM) based method for rehabilitation movement assessment from 3D mocap skeleton data and proposed a performance metric based on Gaussian mixture models to estimate their score. Elkholy et al. [37] extracted spatio-temporal descriptors from 3D Kinect skeleton data to assess the quality of movement for walking on stairs, sit-down, stand-up, and walking actions. They first classified each sequence into normal and abnormal by making a probabilistic model from descriptors derived from normal subjects, and then scored an action by modeling a linear regression on spatio-temporal descriptors of movements with different scores. Khokhlova et al. [10] proposed an LSTM-based method to classify pathological gaits from Kinect skeleton data. They trained several bi-directional LSTMs on different training/validation sets of data. For classification, they computed the weighted mean of the LSTM outputs. All the methods that rely on skeleton data are either unworkable or difficult to apply to in-the-wild scenarios for rehabilitation (or sports or skills) movement analysis.

**View-Invariant Action Recognition**—As stated in References [26,29,38] amongst others, the performance of action recognition methods, such as References [34,35,39–41] to name a few, drops drastically when they test their models from unseen views, since appearance features change significantly in different viewpoints. To overcome this, some works have dealt with viewpoint

variations through skeleton data, for example, References [38,42–44]. For example, Rahmani et al. [38] train an SVM on view-invariant feature vectors from dense trajectories of multiple views in mocap data via a fully connected neural network. Zhang et al. [44] developed a two-stream method, one LSTM and one convolutional model, where both streams include a view adaptation and a classification network. In each case, the former network was trained to estimate the transformation parameters of 3D skeleton data to a canonical view, and the latter classified the action. Finally, the output of the two streams were fused by weighted averaging of the two classifiers' outputs.

As providing skeleton data is difficult for in-the-wild scenarios, others such as References [25–27,29,45] have focused on generating view-invariant features from RGB-D data. Li et al. [26] extract unsupervised view-invariant features by designing a recurrent encoder network which estimated 3D flows from RGB-D streams of two different views. In Reference [29], the authors generated synthetic multi-view video sequence from one view, and then trained a 3D ResNet-50 [40] on both synthetic and real data to classify actions. Among these methods, Varol et al. [29] is the only work that provides cross-view evaluation through single-view training, resulting in 49.4% accuracy on the UESTC dataset [46], which then was increased to 67.8% when they used additional synthetic multi-view data for training.

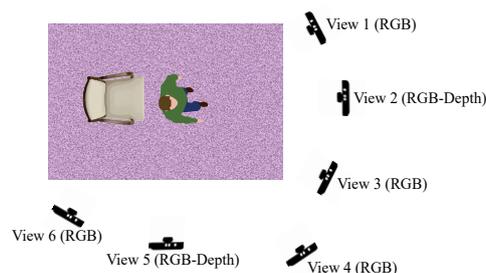
### 3. Datasets

There are many datasets for healthcare applications, such as References [8,30,37,47,48], which are single-view and only include depth and/or skeleton data. To the best of our knowledge, there is no existing dataset (bar one) that is suitable for view-invariant movement assessment from RGB images. The only known multi-view dataset is SMAD, used in Sardari et al. [7]. Although it provides RGB data recorded from 4 different views, it only includes annotated data for a walking action and the subjects' movements are only broadly classified into normal/abnormal, without any scores. Thus it is not a dataset we could use for comparative performance analysis.

Next, we first introduce our new RGB multi-view Quality of Movement Assessment for Rehabilitation dataset, QMAR. Then, we give the details of a recently released rehabilitation dataset KIMORE [30], a single-view dataset that includes RGB images and score annotations, making it suitable for single-view evaluation.

#### 3.1. QMAR

QMAR was recorded using 6 Primesense cameras with 38 healthy subjects, 8 female and 30 male. Figure 2 shows the position of the 6 cameras - 3 different frontal views and 3 different side views. The subjects were trained by a physiotherapist to perform two different types of movements while simulating two ailments, resulting in four overall possibilities: a return walk to approximately the original position while simulating Parkinsons (W-P), and Stroke (W-S), and standing up and sitting down with Parkinson (SS-P) and Stroke (SS-S). The dataset includes RGB and depth (and skeleton) data, although in this work we only use RGB. As capturing depth data from the 6 Primesense cameras was not possible due to infrared interference, the depth and skeleton data were retained from only view 2 at  $\approx 0^\circ$  and view 5 at  $\approx 90^\circ$ .



**Figure 2.** Typical camera views in the QMAR dataset with each one placed at a different height.

The movements in QMAR were scored by the severity of the abnormality. The score ranges were 0 to 4 for W-P, 0 to 5 for W-S and SS-S, and 0 to 12 for SS-P. A score of 0 in all cases indicates a normally executed action. Sample frames from QMAR are shown in Figure 3. Table 1 details the quality score or range and the number of frames and sequences for each action type. Table 2 details the number of sequences for each score.



**Figure 3.** Sample frames from QMAR dataset, showing all 6 views for (top row) walking with Parkinsons (W-P), (second row) walking with Stroke (W-S), (third row) sit-stand with Parkinsons (SS-P), and (bottom row) sit-stand with Stroke.

**Table 1.** Details of the movements in the QMAR dataset.

Action		Quality Score	# Sequences	#Frames/Video Min-Max	Total Frames
W	Normal	0	41	62–179	12,672
W-P	Abnormal	1–4	40	93–441	33,618
W-S	Abnormal	1–5	68	104–500	57,498
SS	Normal	0	42	28–132	9,250
SS-P	Abnormal	1–12	41	96–558	41,808
SS-S	Abnormal	1–5	74	51–580	47,954

**Table 2.** Details of abnormality score ranges in the QMAR dataset.

Action \ Score	#1	#2	#3	#4	#5	#6	#7	#8	#9	#10	#11	#12
W-P	4	8	16	12	-	-	-	-	-	-	-	-
W-S	10	14	19	15	10	-	-	-	-	-	-	-
SS-P	1	1	6	8	4	4	4	3	3	1	2	4
SS-S	3	19	19	13	20	-	-	-	-	-	-	-

### 3.2. KIMORE

This is the only RGB single-view rehabilitation movement dataset where the quality of movements have been annotated for quantitative scores. KIMORE [30] has 78 subjects (44 healthy, and 34 real patients suffering from Parkinson, Stroke, and back pain) performing five types of rehabilitation exercises (Ex #1 to Ex #5) for lower-back pain. All videos are frontal view - see sample frames in Figure 4.



Figure 4. Sample frames of KIMORE for five different exercises.

KIMORE [30] provides two types of scores,  $PO_S$  and  $CF_S$ , with values in the range 0 to 50 for each exercise as defined by clinicians.  $PO_S$  and  $CF_S$  represent the motion of upper limbs and physical constraints during the exercise respectively.

#### 4. Proposed Method

Although its appearance changes significantly when we observe an instance of human movement from different viewpoints, the 2D spatio-temporal trajectories generated by body joints in a sequence are affine transformations of each other. For example, see Figure 5, where the trajectory maps of just the feet joints appear different in orientation, spatial location and scale. Thus, our hypothesis is that by extracting body joint trajectory maps that are translation, rotation, and scale invariant, we should be able to assess the quality of movement from arbitrary viewpoints one may encounter in-the-wild.

The proposed VI-Net network has a view-invariant trajectory descriptor module (VTDM) that feeds into a subsequent movement score module (MSM) as shown in Figure 1. In VTDM, first a 2D convolution filter is applied on stacked heatmaps of each body joint over the video clip frames to generate a trajectory descriptor map. Then, the Spatial Transformer Network (STN) [22] is applied to the trajectory descriptor to make it view-invariant. The spatio-temporal descriptors from all body joints are then stacked as input into the MSM module, which can be implemented by an adapted, pre-trained CNN to learn the relationship amongst the joint trajectories and provide a score for the overall quality of movement. We illustrate the flexibility of MSM by implementing two different pre-trained networks, VGG-19 and ResNeXt-50, and compare their results. VI-Net is trained in an end-to-end manner. As the quality of movement scores in our QMAR dataset are discrete, we use classification to obtain our predicted score. Table 3 carries further details of our proposed VI-Net.



Figure 5. Walking example—all six views, and corresponding trajectory maps for feet.

**Table 3.** VI-Net’s modules:  $\{C2(d \times d, ch)\} \times n$ :  $n$  2D convolution filters with size  $d$  and  $ch$  channel size,  $MP(d \times d)$ : 2D max pooling with size  $d$ ,  $FC(N)$ : FC layer with  $N$  outputs.  $T$  is the # of clip frames,  $J$  is the # of joints and  $S$  is maximum score for a movement type.

	VTDM	MSM (Adapted VGG-19 or ResNeXt-50)
VI-Net	<b>1st layer:</b> $\{C2(3 \times 3, T)\} \times 1$ , BN, ReLU <b>Localisation Network:</b> $\{C2(5 \times 5, 1)\} \times 10, \{MP(2 \times 2)\}$ , ReLU, $\{C2(5 \times 5, 10)\} \times 10, \{MP(2 \times 2)\}$ , ReLU, $\{FC(32)\}$ , ReLU, $\{FC(4)\}$	<b>1st layer VGG-19:</b> $\{C2(3 \times 3, J)\} \times 64$ , BN, ReLU <b>1st layer ResNeXt-50:</b> $\{C2(7 \times 7, J)\} \times 64, \{MP(3 \times 3)\}$ , ReLU <b>Middle layers:</b> As in VGG-19/ResNeXt-50 <b>Last layer:</b> $\{FC(S + 1)\}$

**Generating a Joint Trajectory Descriptor**— First, we extract human body joint heatmaps by estimating the probability of each body joint at each image pixel, per frame for a video clip with  $T$  frames, by applying OpenPose [21]. Even though it may seem that our claim to be an RGB-only method may be undermined by the use of a method which was built by using joint labelling, the fact remains that OpenPose is used in this work as an existing tool, with no further joint labelling or recourse to non-RGB data. Other methods, e.g., Reference [49], which estimate body joint heatmaps from RGB images can equally be used.

To reduce computational complexity, we retain the first 15 joint heatmaps of the BODY-25 version of OpenPose. This is further motivated by the fact, highlighted in Reference [47], that the remaining joints only provide repetitive information. Then, for each body joint  $j \in \{1, 2, \dots, J=15\}$ , we stack its heatmaps over the  $T$ -frame video clip to get the 3D heatmap  $J_j$  of size  $W \times H \times T$  which then becomes the input to our VTDM module. To obtain a body joint's trajectory descriptor  $\Lambda_j$ , the processing in VTDM starts with the application of a convolution filter  $\Phi$  on  $J_j$  to aggregate its spatial poses over time, that is,

$$\Lambda_j = J_j * \Phi, \quad (1)$$

where  $\Lambda_j$  is of size  $W \times H \times 1$ . We experimented with both 2D and 3D convolutions, and found that a  $3 \times 3$  2D convolution filter yields the best results.

**Forging a View-Invariant Trajectory Descriptor**— In the next step of the VTDM module, we experimented with STN [22], DCN [50,51], and ETN [52] networks, and found STN [22] the best performing option to forge a view-invariant trajectory descriptor out of  $\Lambda_j$ .

STN can be applied to feature maps of a CNN's layers to render the output translation, rotation, scale, and shear invariant. It is composed of three stages. At first, a CNN-regression network, referred to as the localisation network, is applied to our joint trajectory descriptor  $\Lambda_j$  to estimate the parameters for a 2D affine transformation matrix,  $\theta = f_{loc}(\Lambda_j)$ . Instead of the original CNN in Reference [22], which applied 32 convolution filters followed by two fully connected (FC) layers, we formulate our own localisation network made up of 10 convolution filters followed by two FC layers. The rationale for this is that our trajectory descriptor maps are not as complex as RGB images, and hence fewer filters are sufficient to extract their features. The details of our localisation network's layers are provided in Table 3. Then, in the second stage, to estimate each pixel value of our view-invariant trajectory descriptor  $\bar{\Lambda}_j$ , a sampling kernel is applied on specific regions of  $\Lambda_j$ , where the centres of these regions are defined on a sampling grid. This sampling grid  $\Gamma_\theta(G)$  is generated from a general grid  $G = \{(x_i^g, y_i^g)\}, i \in \{1, \dots, W' \times H'\}$  and the predicted transformation parameters, such that

$$\begin{pmatrix} x_i^{\Lambda_j} \\ y_i^{\Lambda_j} \end{pmatrix} = \Gamma_\theta(G_i) = \begin{bmatrix} \theta_{11} & \theta_{12} \\ \theta_{21} & \theta_{22} \end{bmatrix} \times \begin{pmatrix} x_i^g \\ y_i^g \end{pmatrix}, \quad (2)$$

where  $\Gamma_\theta(G) = \{(x_i^{\Lambda_j}, y_i^{\Lambda_j}), i \in \{1, \dots, W' \times H'\}$  are the centers of the regions of  $\Lambda_j$  the sampling kernel is applied to, in order to generate the new pixel values of the output feature map  $\bar{\Lambda}_j$ . Jaderberg et al. [22] recommend the use of different types of transformations to generate the sampling grid  $\Gamma_\theta(G)$  based on the problem domain. In VTDM, we use the 2D affine transformations shown in Equation (2). Finally, the sampler takes both  $\Lambda_j$  and  $\Gamma_\theta(G)$  to generate a view-invariant trajectory descriptor  $\bar{\Lambda}_j$  from  $\Lambda_j$  at the grid points by bilinear interpolation.

**Assessing the Quality of Human Movement**— In the final part of VI-Net (see Figure 1), the collection of view-invariant trajectory descriptors  $\bar{\Lambda}_j$  for joints  $j \in \{1, 2, \dots, J=15\}$ , are stacked into a global descriptor  $\bar{\Lambda}$  and passed through a pre-trained network in the MSM module to assess the quality of movement of the joints. VGG-19 and ResNeXt-50 were chosen for their state-of-the-art performances, popularity, and availability. For VGG-19, its first layer was replaced with a new 2D convolutional layer, including  $3 \times 3$  convolution filters with channel size  $J$  (instead of 3 used for RGB input images), and for ResNeXt-50 its first layer was replaced with  $7 \times 7$  convolution filters with

channel size  $J$ . The last FC layer in each case was modified to allow movement quality scoring through classification where each score is considered as a class, that is, for a movement with maximum score  $S$ , where  $S=4$  for W-P,  $S=5$  for W-S and SS-S, and  $S=12$  for SS-P movements. The last FC layer of VI-Net has  $S + 1$  output units.

Although VGG-19/ResNeXt-50 were trained on RGB images, we still benefit from their pretrained weights, since our new first layers were initialised with their original first layer weights. The output of this modified layer has the same size as the output of the layer it replaces (Table 3), so the new layer is compatible with the rest of network. In addition, we normalize the pixel values of the trajectory heatmaps to be between 0 and 255, that is, the same as RGB images on which VGG and ResNeXt were trained on, and trajectory descriptor maps have shape and intensity variations - thus the features extracted from them would be as equally valid as for natural images on which VGG and ResNeXt operate.

## 5. Experiments and Results

We first report on two sets of experiments on QMAR to evaluate the performance of VI-Net to assess quality of movement, based around cross-subject and cross-view scenarios. Then, to show the efficiency of VI-Net on other datasets and movement types, we present its results also on the single-view KIMORE dataset. We used Pytorch on two GeForce GTX 750 GPUs. All networks were trained for 20 epochs using stochastic gradient descent optimization with initial learning rate of 0.001, and batch size 5. To evaluate the performance of the proposed method, we used Spearman's rank correlation as used in References [1,3,4].

**Dataset Imbalance**— It can be seen from Tables 1 and 2 that the number of sequences for score 0 (normal) is many more than the number of sequences for other individual scores, so we randomly selected 15 normal sequences for W-P, W-S, SS-S movements and 4 normal sequences for SS-P to mix with abnormal movements to perform all our experiments. To further address the imbalance, we applied offline temporal cropping to add new sequences.

**Network Training and Testing**— For each movement type, the proposed network is trained from scratch. In both the training and testing phases, video sequences were divided into 16-frame clips (without overlaps). In training, the clips were selected randomly from amongst all video sequences of the training set, and passed to VI-Net. Then, the weights were updated following a cross-entropy loss,

$$L_C(f, s) = -\log\left(\frac{\exp(f(s))}{\sum_{k=0}^S \exp(f(k))}\right), \quad (3)$$

where  $f(\cdot)$  is the  $S + 1$  dimensional output of the last fully connected layer and  $s$  is the video clip's ground truth label/score. In testing, every 16-frame clip of a video sequence was passed to VI-Net. After averaging the outputs of the last fully connected layer across each class for all the clips, we then set the score for the whole video sequence as the maximum of the clip scores (see Figure 6), that is,

$$s = \operatorname{argmax}_k(\bar{f}(k) = \frac{1}{M} \sum_{m=1}^M f_m(k)), \quad (4)$$

where  $k \in \{0, 1, \dots, S\}$  and  $M$  is the number of clips for a video.

**Comparative Evaluation**— As we are not aware of any other RGB-based view-invariant method to assess quality of movement, we are unable to compare VI-Net's performance to other methods under a cross-view scenario. However, for cross-subject and single-view scenarios, we evaluate against (a) a C3D baseline (fashioned after Parmar and Morris [3]) by combining the outputs of the C3D network to score a sequence in the test phase in the same fashion as in VI-Net, and (b) the pre-trained, fine-tuned I3D [35]. We also provide an ablation study for all scenarios by removing STN from VI-Net to analyse the impact of this part of the proposed method.

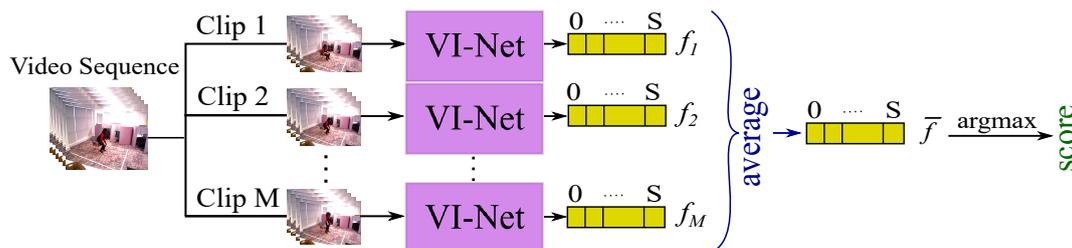


Figure 6. Scoring process for a full video sequence in testing phase.

### 5.1. Cross-Subject Quality of Movement Analysis

In this experiment, all available views were used in both training and testing, while the subjects performing the actions were distinct. We applied  $k$ -fold cross validation where  $k$  is the number of scores for each movement. Table 4 shows that VI-Net outperforms networks based on C3D (after Reference [3]) and I3D [35] for all types of movements, regardless of whether VGG-19 or ResNeXt-50 are used in the MSM module. While I3D results are mostly competitive, C3D performs less well due to its shallower nature, and larger number of parameters, exacerbated by QMAR's relatively small size. We show in Section 5.3 that C3D performs significantly better on a larger dataset.

As ablation analysis, to test the effectiveness of STN, we present VI-Net's results with and without STN in Table 4. It can be observed that the improvements with STN are not necessarily consistent across the actions since when all viewpoints are used in training, the MSM module gets trained on all trajectory orientations such that the effect of STN is often overridden.

Table 4. Comparative cross-subject results on QMAR. The bold numbers show the best result for each action type.

Method		Action				Avg	
		W-P	W-S	SS-P	SS-S		
Custom-trained C3D (after Reference [3])		0.50	0.37	0.25	0.54	0.41	
Pre-trained I3D [35]		0.79	0.47	0.54	0.55	0.58	
VI-Net	VTDM+MSM (VGG-19)	w/o STN	0.81	0.49	0.57	<b>0.74</b>	0.65
		w STN	0.82	0.52	0.55	0.73	0.65
	VTDM+MSM (ResNeXt-50)	w/o STN	<b>0.87</b>	<b>0.56</b>	0.48	0.72	0.65
		w STN	<b>0.87</b>	0.52	<b>0.58</b>	0.69	<b>0.66</b>

### 5.2. Cross-View Quality of Movement Analysis

We evaluate the generalization ability of VI-Net on unseen views by using cross-view scenarios, that is, distinct training and testing views of the scene, while data from all subjects is utilised. We also make sure that each test set contains a balanced variety of scores from low to high. Recent works such as References [25–28], provide cross-view results only when their network is trained on multiple views. As recently noted by Varol et al. [29], a highly challenging scenario in view-invariant action recognition would be to obtain cross-view results by training from only one viewpoint. Therefore, we performed the training and testing for each movement type such that (i) we trained from one view only and tested on all other views (as reasoned in Section 1), and in the next experiment, (ii) we trained on a combination of one frontal view (views 1 to 3) and one side view (views 4 to 6) and tested on all other available views. Since for the latter case there are many combinations, we show results for only selected views: view 2  $\approx 0^\circ$  with all side views, and view 5  $\approx 90^\circ$  with all frontal views.

Since in cross-view analysis all subjects are used in both training and testing, applying the C3D and I3D models would be redundant because they would simply learn the appearance and shape features of our participants in our study and their results would be unreliable.

In QMAR, when observing a movement from the frontal views, there is little or almost no occlusion of relevant body parts. However, when observing from side views, occlusions resulting from missing or noisy joint heatmaps from OpenPose, can occur for a few seconds or less (short-term), or for

almost the whole sequence (long-term). Short term occlusions are more likely in walking movements W-P and W-S, while long-term occlusions occur more often in sit-to-stand movements (SS-P and SS-S).

The results of our view-invariancy experiments, using single views only in training, are shown in Table 5. It can be seen that for walking movements W-P and W-S, VI-Net is able to assess the movements from unseen views well, with the best results reaching 0.73 and 0.66 rank correlation respectively (yellow highlights), and only relatively affected by short term occlusions. However, for sit-to-stand movements SS-P and SS-S, the long-term occlusions during these movements affect the integrity of the trajectory descriptors and the performance of VI-Net is not as strong, with the best results reaching 0.52 and 0.56 respectively (orange highlights). Note, for all action types, when VI-Net has STN with adapted ResNeXt, it performs best on average.

Table 6 shows the results for each movement type when one side view and one frontal view are combined for training. VI-Net's performance improves compared to the single-view experiment above with the best results reaching 0.92 and 0.83 for W-P and W-S movements (green highlights) and 0.61 and 0.64 for SS-P and SS-S movements (purple highlights), because the network is effectively trained with both short-term and long-term occluded trajectory descriptors. These results also show that on average VI-Net performs better with adapted ResNeXt-50 for walking movements (W-P and W-S) and with adapted VGG-19 for sit-to-stand movements (SS-P and SS-S). This is potentially because ResNext-50's variety of filter sizes are better suited to the variation in 3D spatial changes of joint trajectories inherent in walking movements compared to VGG-19's  $3 \times 3$  filters which can tune better to the more spatially restricted sit-to-stand movements. We also note that the fundamental purpose of STN in VI-Net is to ensure efficient cross-view performance is possible when the network is trained from a single-view only. It would therefore be expected and plausible that STN's effect would diminish as more views are used since the MSM module gets trained on more trajectory orientations (which we verified experimentally by training with multiple views).

**Table 5.** Cross-view results for all actions with single-view training. The bold numbers show the best result for each view of each action type; Yellow highlights: best results for W-P and W-S actions amongst all views, Orange highlights: best result for SS-P and SS-S actions amongst all views.

	View	VTDM+MSM (VGG-19)		VTDM+MSM (ResNeXt-50)		View	VTDM+MSM (VGG-19)		VTDM+MSM (ResNeXt-50)	
		w/o STN	w STN	w/o STN	w STN		w/o STN	w STN	w/o STN	w STN
W-P	1	0.51	<b>0.67</b>	0.64	<b>0.67</b>	1	0.51	0.43	0.60	<b>0.64</b>
	2	0.69	0.66	0.58	<b>0.72</b>	2	0.47	0.54	0.55	<b>0.62</b>
	3	0.62	0.66	0.63	<b>0.70</b>	3	<b>0.64</b>	0.56	0.61	0.59
	4	0.67	0.64	<b>0.72</b>	<b>0.72</b>	4	0.60	0.59	0.60	<b>0.66</b>
	5	0.67	0.67	0.68	<b>0.71</b>	5	0.62	0.60	0.62	<b>0.63</b>
	6	0.69	0.72	0.69	<b>0.73</b>	6	0.46	0.40	0.53	<b>0.60</b>
	<b>Avg</b>	0.64	0.67	0.65	<b>0.70</b>	<b>Avg</b>	0.55	0.52	0.58	<b>0.62</b>
SS-P	1	0.30	<b>0.32</b>	0.25	0.25	1	0.36	<b>0.49</b>	0.44	0.45
	2	0.27	0.31	0.31	<b>0.32</b>	2	0.47	0.40	<b>0.56</b>	<b>0.56</b>
	3	0.16	0.23	0.36	<b>0.43</b>	3	0.37	<b>0.52</b>	0.38	0.43
	4	0.10	0.34	0.44	<b>0.49</b>	4	0.38	0.34	0.41	<b>0.54</b>
	5	0.50	<b>0.52</b>	0.43	0.45	5	0.26	<b>0.50</b>	<b>0.50</b>	0.48
	6	0.41	0.24	<b>0.48</b>	0.44	6	0.21	<b>0.28</b>	0.13	0.16
	<b>Avg</b>	0.29	0.32	0.37	<b>0.39</b>	<b>Avg</b>	0.34	0.42	0.40	<b>0.43</b>

**Table 6.** Cross-view results for all actions with two-view training. The bold numbers show the best result for each combination of views of each action type; Green highlights: best results for W-P and W-S actions amongst all view combinations, Purple highlights: best results for SS-P and SS-S actions amongst all view combinations.

	View	VTDM+MSM (VGG-19)		VTDM+MSM (ResNeXt-50)			View	VTDM+MSM (VGG-19)		VTDM+MSM (ResNeXt-50)	
		w/o STN	w STN	w/o STN	w STN			w/o STN	w STN	w/o STN	w STN
W-P	2,4	0.77	0.81	0.87	<b>0.89</b>	W-S	2,4	0.58	0.72	<b>0.81</b>	0.73
	2,5	0.72	0.75	0.90	<b>0.92</b>		2,5	0.74	0.74	0.80	<b>0.81</b>
	2,6	0.75	0.76	0.73	<b>0.77</b>		2,6	0.64	0.67	<b>0.74</b>	0.68
	1,5	0.70	0.76	<b>0.80</b>	0.75		1,5	0.70	0.68	<b>0.83</b>	0.81
	3,5	0.73	0.79	<b>0.87</b>	0.84		3,5	0.66	0.66	<b>0.82</b>	0.79
	<b>Avg</b>	0.73	0.77	<b>0.83</b>	<b>0.83</b>		<b>Avg</b>	0.66	0.69	<b>0.80</b>	0.76
SS-P	2,4	<b>0.55</b>	0.52	0.41	0.46	SS-S	2,4	0.57	<b>0.64</b>	0.54	<b>0.64</b>
	2,5	<b>0.60</b>	0.53	0.49	0.46		2,5	0.62	0.56	<b>0.63</b>	0.61
	2,6	<b>0.48</b>	0.35	0.36	0.42		2,6	0.50	<b>0.62</b>	0.48	0.46
	1,5	0.46	<b>0.55</b>	0.39	0.52		1,5	<b>0.64</b>	0.53	0.48	0.58
	3,5	<b>0.61</b>	0.40	0.43	0.47		3,5	0.62	0.60	0.63	<b>0.67</b>
	<b>Avg</b>	<b>0.54</b>	0.47	0.41	0.46		<b>Avg</b>	<b>0.59</b>	<b>0.59</b>	0.55	0.58

### 5.3. Single-View Quality of Movement Analysis

Next, we provide the results of VI-Net on the single-view KIMORE dataset, to illustrate that it can be applied to such data too. KIMORE provides two types of scores,  $PO_S$  and  $CF_S$  (see Section 3.2) which have a strong correspondence to each other, such that if one is low for a subject, so is the other. Hence, we trained the network based on a single, summed measure to predict a final score ranging between 0 and 100 for each action type. We include 70% of the subjects for training and retain the remaining 30% for testing ensuring each set contains a balanced variety of scores from low to high.

Table 7 shows the results of C3D baseline (after Reference [3]), pre-trained, fine-tuned I3D [35] and VI-Net on KIMORE. It can be seen that VI-Net outperforms the other methods for all movement types except for Exercise #3. VI-Net with adapted VGG-19 performs better than with ResNeXt-50 for all movement types. This may be because, similar to sit-to-stand movements in QMAR, where VI-Net performs better with VGG-19, all movements types in KIMORE are also performed at the same location and distance from camera, and thus carry less variation in 3D trajectory space. This shows that our results are consistent in this sense across both datasets.

In addition, although all sequences in both training and testing sets have been captured from the same view, VI-Net's performance on average improves with STN. This can be attributed to STN improving the network generalization on different subjects. Also, unlike in QMAR's cross-subject results where C3D performed poorly, the results on KIMORE for C3D are promising because KIMORE has more data to help the network train more efficiently.

**Table 7.** Comparative results on the single-view KIMORE dataset. The bold numbers show the best result for each action type.

Method		Action					Average	
		Ex #1	Ex #2	Ex #3	Ex #4	Ex #5		
Custom-trained C3D (after Reference [3])		0.66	0.64	<b>0.63</b>	0.59	0.60	0.62	
Pre-trained I3D [35]		0.45	0.56	0.57	0.64	0.58	0.56	
VI-Net	VTDM+MSM (VGG-19)	w/o STN	0.63	0.50	0.55	<b>0.80</b>	<b>0.76</b>	0.64
		w STN	<b>0.79</b>	<b>0.69</b>	0.57	0.59	0.70	
	VTDM+MSM (ResNeXt-50)	w/o STN	0.55	0.42	0.33	0.62	0.57	0.49
		w STN	0.55	0.62	0.36	0.58	0.67	0.55

## 6. Conclusions

View-invariant human movement analysis from RGB is a significant challenge in action analysis applications, such as sports, skill assessment, and healthcare monitoring. In this paper, we proposed

a novel RGB based view-invariant method to assess the quality of human movement which can be trained from a relatively small dataset and without any knowledge about viewpoints used for data capture. We also introduced QMAR, the only multi-view, non-skeleton, non-mocap, rehabilitation movement dataset to evaluate the performance of the proposed method, which may also serve well for comparative analysis for the community. We demonstrated that the proposed method is applicable to cross-subject, cross-view, and single-view movement analysis by achieving average rank correlation 0.66 on cross-subject and 0.65 on unseen views when trained from only two views, and 0.66 on single-view setting.

VI-Net's performance drops in situations where long-term occlusions occur, since OpenPose fails in such cases to produce sufficiently consistent heatmaps - but in general many methods suffer from long-term occlusions, so such failure is expected. Another limitation of VI-Net is that it has to be trained separately for each movement type. For future work, we plan to apply 3D pose estimation methods to generate more robust joint heatmaps which would also be less troubled by occlusions. We also plan to develop multitask learning so that the network can recognize the movement type and its score simultaneously. Moreover, we aim to improve the performance of our method on unseen views by unsupervised training of view-invariant features from existing multi-view datasets for transfer to our domain.

**Author Contributions:** The authors contributed to the paper as follows: conceptualization, M.M. and F.S.; methodology, F.S. and M.M.; investigation, F.S. and M.M.; coding and validation, F.S.; formal analysis, F.S. and M.M.; data curation, F.S, M.M. and S.H.; writing—original draft preparation, F.S.; writing—review and editing, M.M.; supervision, M.M. and A.P.; project administration, M.M. All authors have read and agreed to the published version of the manuscript.

**Funding:** This work was performed under the SPHERE Next Steps Project funded by the UK Engineering and Physical Sciences Research Council (EPSRC), Grant EP/R005273/1.

**Acknowledgments:** The 1st author is grateful to the University of Bristol for her scholarship. The authors would also like to thank Alan Whone and Harry Rolinski of Southmead Hospital—clinical experts in Neurology—for fruitful discussions on the QMAR dataset.

**Conflicts of Interest:** The authors declare no conflict of interest. The funders had no role in the design of the study, in the collection, analyses, or interpretation of data, in the writing of the manuscript, or in the decision to publish the results.

## References

1. Li, Y.; Chai, X.; Chen, X. End-to-End Learning for Action Quality Assessment. In Proceedings of the Pacific Rim Conference on Multimedia, Hefei, China, 21–22 September 2018; pp. 125–134.
2. Parmar, P.; Tran Morris, B. Learning to Score Olympic Events. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshop, Honolulu, HI, USA, 21–26 July 2017; pp. 20–28.
3. Parmar, P.; Morris, B.T. 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, Long Beach, CA, USA, 16–20 June 2019; pp. 304–313.
4. Pan, J.H.; Gao, J.; Zheng, W.S. Action Assessment by Joint Relation Graphs. In Proceedings of the IEEE International Conference on Computer Vision, Seoul, Korea, 27 October–2 November 2019; pp. 6331–6340.
5. Fard, M.J.; Ameri, S.; Darin Ellis, R.; Chinnam, R.B.; Pandya, A.K.; Klein, M.D. Automated Robot-Assisted Surgical Skill Evaluation: Predictive Analytics Approach. *Int. J. Med. Robot. Comput. Assist. Surg.* **2018**, *14*, 1850. [[CrossRef](#)] [[PubMed](#)]
6. Doughty, H.; Mayol-Cuevas, W.; Damen, D. The Pros and Cons: Rank-Aware Temporal Attention for Skill Determination in Long Videos. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, Long Beach, CA, USA, 16–20 June 2019; pp. 7862–7871.
7. Sardari, F.; Paiement, A.; Mirmehdi, M. View-Invariant Pose Analysis for Human Movement Assessment from RGB Data. In Proceedings of the International Conference on Image Analysis and Processing, Trento, Italy, 9–13 September 2019; pp. 237–248.
8. Liao, Y.; Vakanski, A.; Xian, M. A Deep Learning Framework for Assessing Physical Rehabilitation Exercises. *IEEE Trans. Neural Syst. Rehabil. Eng.* **2019**, *28*, 468–477. [[CrossRef](#)] [[PubMed](#)]

9. Crabbe, B.; Paiement, A.; Hannuna, S.; Mirmehdi, M. Skeleton-free Body Pose Estimation from Depth Images for Movement Analysis. In Proceedings of the IEEE International Conference on Computer Vision Workshop, Santiago, Chile, 7–13 December 2015; pp. 70–78.
10. Khokhlova, M.; Migniot, C.; Morozov, A.; Sushkova, O.; Dipanda, A. Normal and Pathological Gait Classification LSTM Model. *Artif. Intell. Med.* **2019**, *94*, 54–66. [[CrossRef](#)] [[PubMed](#)]
11. Antunes, J.; Bernardino, A.; Smailagic, A.; Siewiorek, D.P. AHA-3D: A Labelled Dataset for Senior Fitness Exercise Recognition and Segmentation from 3D Skeletal Data. In Proceedings of the British Machine Vision Conference, Newcastle, UK, 3–6 September 2018; p. 332.
12. Blanchard, N.; Skinner, K.; Kemp, A.; Scheirer, W.; Flynn, P. “Keep Me In, Coach!”: A Computer Vision Perspective on Assessing ACL Injury Risk in Female Athletes. In Proceedings of the IEEE Winter Conference on Applications of Computer Vision, Waikoloa Village, HI, USA, 7–11 January 2019; pp. 1366–1374.
13. Lei, Q.; Du, J.X.; Zhang, H.B.; Ye, S.; Chen, D.S. A Survey of Vision-Based Human Action Evaluation Methods. *Sensors* **2019**, *19*, 4129. [[CrossRef](#)] [[PubMed](#)]
14. Wandt, B.; Rosenhahn, B. RepNet: Weakly Supervised Training of an Adversarial Reprojection Network for 3D Human Pose Estimation. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, Long Beach, CA, USA, 16–20 June 2019.
15. Zhao, L.; Peng, X.; Tian, Y.; Kapadia, M.; Metaxas, D.N. Semantic Graph Convolutional Networks for 3D Human Pose Regression. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, Long Beach, CA, USA, 16–20 June 2019.
16. Zhou, K.; Han, X.; Jiang, N.; Jia, K.; Lu, J. HEMlets Pose: Learning Part-Centric Heatmap Triplets for Accurate 3D Human Pose Estimation. In Proceedings of the IEEE International Conference on Computer Vision, Seoul, Korea, 27 October–2 November 2019.
17. Kolotouros, N.; Pavlakos, G.; Black, M.J.; Daniilidis, K. Learning to Reconstruct 3D Human Pose and Shape via Model-Fitting in the Loop. In Proceedings of the IEEE International Conference on Computer Vision, Long Beach, CA, USA, 16–20 June 2019; pp. 2252–2261.
18. Kocabas, M.; Athanasiou, N.; Black, M.J. VIBE: Video Inference for Human Body Pose and Shape Estimation. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, Seattle, WA, USA, 14–19 June 2020; pp. 5253–5263.
19. Qiu, H.; Wang, C.; Wang, J.; Wang, N.; Zeng, W. Cross view fusion for 3d human pose estimation. In Proceedings of the IEEE International Conference on Computer Vision, Long Beach, CA, USA, 16–20 June 2019; pp. 4342–4351.
20. Remelli, E.; Han, S.; Honari, S.; Fua, P.; Wang, R. Lightweight Multi-View 3D Pose Estimation through Camera-Disentangled Representation. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, Seattle, WA, USA, 14–19 June 2020; pp. 6040–6049.
21. Cao, Z.; Simon, T.; Wei, S.E.; Sheikh, Y. Realtime Multi-Person 2D Pose Estimation Using Part Affinity Fields. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, Honolulu, HI, USA, 21–26 July 2017; pp. 7291–7299.
22. Jaderberg, M.; Simonyan, K.; Zisserman, A.; Kavukcuoglu, K. Spatial Transformer Networks. In Proceedings of the Advances in Neural Information Processing Systems, Montreal, QC, Canada, 7–12 December 2015; pp. 2017–2025.
23. Simonyan, K.; Zisserman, A. Very Deep Convolutional Networks for Large-Scale Image Recognition. *arXiv* **2014**, arXiv:1409.1556.
24. Xie, S.; Girshick, R.; Dollár, P.; Tu, Z.; He, K. Aggregated Residual Transformations for Deep Neural Networks. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, Honolulu, HI, USA, 21–26 July 2017; pp. 1492–1500.
25. Wang, D.; Ouyang, W.; Li, W.; Xu, D. Dividing and Aggregating Network for Multi-View Action Recognition. In Proceedings of the European Conference on Computer Vision, Munich, Germany, 8–14 September 2018; pp. 451–467.
26. Li, J.; Wong, Y.; Zhao, Q.; Kankanhalli, M. Unsupervised Learning of View-Invariant Action Representations. In Proceedings of the Advances in Neural Information Processing Systems, Montreal, OC, Canada, 3–8 December 2018; pp. 1254–1264.

27. Lakhal, M.I.; Lanz, O.; Cavallaro, A. View-LSTM: Novel-View Video Synthesis Through View Decomposition. In Proceedings of the IEEE International Conference on Computer Vision, Seoul, Korea, 27 October–2 November 2019; pp. 7577–7587.
28. Li, W.; Xu, Z.; Xu, D.; Dai, D.; Van Gool, L. Domain Generalization and Adaptation Using Low Rank Exemplar SVMs. *IEEE Trans. Pattern Anal. Mach. Intell.* **2017**, *40*, 1114–1127. [[CrossRef](#)] [[PubMed](#)]
29. Varol, G.; Laptev, I.; Schmid, C.; Zisserman, A. Synthetic Humans for Action Recognition from Unseen Viewpoints. *arXiv* **2019**, arXiv:1912.04070.
30. Capecci, M.; Ceravolo, M.G.; Ferracuti, F.; Iarlori, S.; Monteriù, A.; Romeo, L.; Verdini, F. The KIMORE Dataset: Kinematic Assessment of Movement and Clinical Scores for Remote Monitoring of Physical Rehabilitation. *IEEE Trans. Neural Syst. Rehabil. Eng.* **2019**, *27*, 1436–1448. [[CrossRef](#)] [[PubMed](#)]
31. Pirsivash, H.; Vondrick, C.; Torralba, A. Assessing The Quality of Actions. In Proceedings of the European Conference on Computer Vision, Zurich, Switzerland, 6–12 September 2014; pp. 556–571.
32. Xiang, X.; Tian, Y.; Reiter, A.; Hager, G.D.; Tran, T.D. S3D: Stacking Segmental P3D for Action Quality Assessment. In Proceedings of the IEEE International Conference on Image Processing, Athens, Greece, 7–10 October 2018; pp. 928–932.
33. 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 Conference on Computer Vision and Pattern Recognition, Seattle, WA, USA, 13–19 June 2020; pp. 9839–9848.
34. Tran, D.; Bourdev, L.; Fergus, R.; Torresani, L.; Paluri, M. Learning Spatiotemporal Features With 3D Convolutional Networks. In Proceedings of the IEEE International Conference on Computer Vision, Santiago, Chile, 7–13 December 2015; pp. 4489–4497.
35. Carreira, J.; Zisserman, A. Quo Vadis, Action Recognition? a New Model and the Kinetics Dataset. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, Honolulu, HI, USA, 21–26 July 2017; pp. 6299–6308.
36. Tao, L.; Paiement, A.; Damen, D.; Mirmehdi, M.; Hannuna, S.; Camplani, M.; Burghardt, T.; Craddock, I. A Comparative Study of Pose Representation and Dynamics Modelling for Online Motion Quality Assessment. *Comput. Vis. Image Underst.* **2016**, *148*, 136–152. [[CrossRef](#)]
37. Elkholy, A.; Hussein, M.; Gomaa, W.; Damen, D.; Saba, E. Efficient and Robust Skeleton-Based Quality Assessment and Abnormality Detection in Human Action Performance. *IEEE J. Biomed. Health Inform.* **2019**, *24*, 208–291. [[CrossRef](#)]
38. Rahmani, H.; Mian, A.; Shah, M. Learning a Deep Model for Human Action Recognition from Novel Viewpoints. *IEEE Trans. Pattern Anal. Mach. Intell.* **2018**, *40*, 667–681. [[CrossRef](#)]
39. Feichtenhofer, C.; Fan, H.; Malik, J.; He, K. Slowfast Networks for Video Recognition. In Proceedings of the IEEE International Conference on Computer Vision, Seoul, Korea, 27 October–2 November 2019; pp. 6202–6211.
40. Hara, K.; Kataoka, H.; Satoh, Y. Can Spatiotemporal 3D CNNs Retrace the History of 2D CNNs and Imagenet? In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, Salt Lake City, UT, USA, 18–23 June 2018; pp. 6546–6555.
41. Lin, J.; Gan, C.; Han, S. Tsm: Temporal Shift Module for Efficient Video Understanding. In Proceedings of the IEEE International Conference on Computer Vision, Seoul, Korea, 27 October–2 November 2019; pp. 7083–7093.
42. Ke, Q.; Bennamoun, M.; An, S.; Sohel, F.; Boussaid, F. A New Representation of Skeleton Sequences for 3D Action Recognition. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, Honolulu, HI, USA, 21–26 July 2017; pp. 3288–3297.
43. Liu, M.; Liu, H.; Chen, C. Enhanced Skeleton Visualization for View Invariant Human Action Recognition. *Pattern Recog.* **2017**, *68*, 346–362. [[CrossRef](#)]
44. Zhang, P.; Lan, C.; Xing, J.; Zeng, W.; Xue, J.; Zheng, N. View Adaptive Neural Networks for High Performance Skeleton-Based Human Action Recognition. *IEEE Trans. Pattern Anal. Mach. Intell.* **2019**, *41*, 1963–1978. [[CrossRef](#)]
45. Liu, M.; Yuan, J. Recognizing Human Actions as the Evolution of Pose Estimation Maps. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, Salt Lake City, UT, USA, 18–23 June 2018; pp. 1159–1168.

46. Ji, Y.; Xu, F.; Yang, Y.; Shen, F.; Shen, H.T.; Zheng, W.S. A Large-Scale RGB-D Database for Arbitrary-View Human Action Recognition. In Proceedings of the ACM International Conference on Multimedia, Seoul, Korea, 22–26 October 2018; pp. 1510–1518.
47. Paiement, A.; Tao, L.; Hannuna, S.; Camplani, M.; Damen, D.; Mirmehdi, M. Online Quality Assessment of Human Movement from Skeleton Data. In Proceedings of the British Machine Vision Conference, Nottingham, UK, 1–5 September 2014; pp. 153–166.
48. Vakanski, A.; Jun, H.p.; Paul, D.; Baker, R. A Data Set of Human Body Movements for Physical Rehabilitation Exercises. *Data* **2018**, *3*, 2. [[CrossRef](#)]
49. Kocabas, M.; Karagoz, S.; Akbas, E. Self-supervised Learning of 3D Human Pose Using Multi-view Geometry. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, Long Beach, CA, USA, 16–20 June 2019; pp. 1077–1086.
50. Dai, J.; Qi, H.; Xiong, Y.; Li, Y.; Zhang, G.; Hu, H.; Wei, Y. Deformable Convolutional Networks. In Proceedings of the IEEE International Conference on Computer Vision, Venice, Italy, 22–29 October 2017; pp. 764–773.
51. Zhu, X.; Hu, H.; Lin, S.; Dai, J. Deformable Convnets v2: More Deformable, Better Results. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, Long Beach, CA, USA, 16–20 June 2019; pp. 9308–9316.
52. Tai, K.S.; Bailis, P.; Valiant, G. Equivariant Transformer Networks. In Proceedings of the International Conference on Machine Learning, Long Beach, CA, USA, 10–15 June 2019.



© 2020 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<http://creativecommons.org/licenses/by/4.0/>).