

# ExeChecker: Where Did I Go Wrong?

Yiwen Gu<sup>1</sup>, Mahir Patel<sup>1</sup>, and Magrit Betke<sup>1</sup>

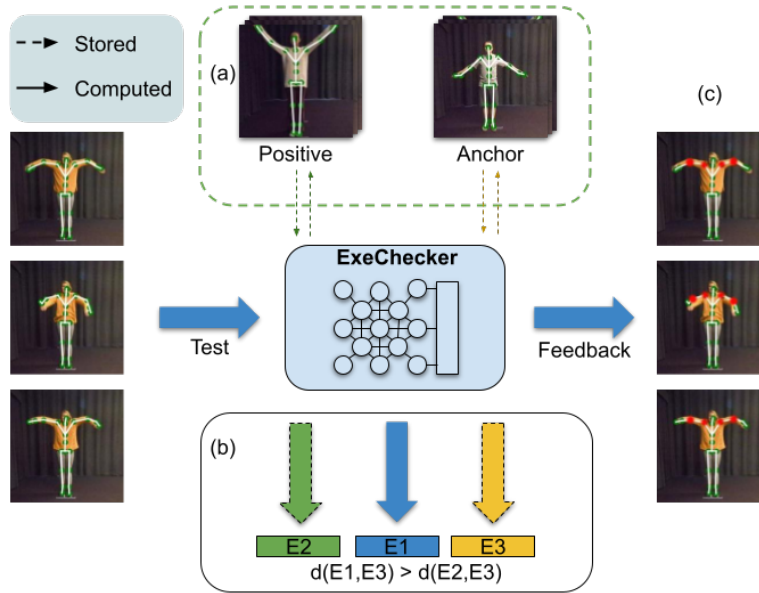
Boston University, Boston MA 02125, USA  
{yiweng, mahirp, betke}@bu.edu

**Abstract.** In this paper, we present a contrastive learning based framework, **ExeChecker**, for the interpretation of rehabilitation exercises. Our work builds upon state-of-the-art advances in the area of human pose estimation, graph-attention neural networks, and transformer interpretability. The downstream task is to assist rehabilitation by providing informative feedback to users while they are performing prescribed exercises. We utilize a contrastive learning strategy during training. Given a tuple of correctly and incorrectly executed exercises, our model is able to identify and highlight those joints that are involved in an incorrect movement and thus require the user’s attention. We collected an in-house dataset, **ExeCheck**, with paired recordings of both correct and incorrect execution of exercises. In our experiments, we tested our method on this dataset as well as the UI-PRMD dataset and found ExeCheck outperformed the baseline method using pairwise sequence alignment in identifying joints of physical relevance in rehabilitation exercises.

**Keywords:** Movement assessment · Contrastive learning

## 1 Introduction

Rehabilitation exercises help individuals recover from injuries, surgeries, or chronic conditions. In the clinical setting, physical therapists consider the physical conditions of the patients and prescribe them different sets of exercises or the same exercise but with different standards. In most cases, patients learn the prescribed exercises during their clinic visits but then perform them regularly in their homes. In the absence of a therapist, patients do not receive adequate guidance and instructive feedback. This may lead to a lack of motivation and incorrect exercise performance, negatively impacting the patient’s health. To help patients better adhere to the prescribed exercises, researchers build home-based systems that take the patient’s movements as input, analyze their performance, and automatically provide feedback [15, 21]. However, most feedback generated by those systems is too general ("excellent/ good/ needs more work") or too obscure (some score). Such feedback does not provide information regarding which movements the patients did incorrectly and hence need attention. In this paper, we propose a framework, **ExeChecker**, that tells incorrect from the correct movement for a set of rehabilitation exercises and points out the joints that cause the incorrect performance (Fig. 1).



**Fig. 1:** ExeChecker interprets the movement of skeleton joints extracted from RGB video and provides feedback to the user about which joints to pay attention to. (a) We first generate a database by pre-computing and storing the embeddings for the positive and anchor sequences. (b) During inference, we compute embedding E1 for the input sequence  $M$  and test it against the positive (E2) and anchor (E3) pair. (c) By using the attention score from our graph attention network, we highlight the joints (red) that require adjustments.

**ExeChecker** takes skeleton-based human poses, leverages the existing spatial-temporal graph attention transformer, and adopts a triplet learning strategy. During the training, the model is forced to minimize the distances between the positive pairs (consisting of the sequences of correctly conducted exercises) and maximize the distances of negative pairs (consisting of correctly and incorrectly conducted exercises) in the embedding space. As a result, the learned embedding focuses on the dissimilarities between the incorrect and correct exercises.

We conducted extensive experiments on two datasets, UI-PRMD [25] and our in-house collected dataset **ExeCheck**. Both datasets feature common rehabilitation exercises and paired recordings of both correct and incorrect performance.

Our contributions are summarized as follows:

- We collaborated with a physical therapist who specializes in exercise rehabilitation and collected an in-house dataset that includes ten frequently prescribed exercises for patients with Parkinson’s disease. For each exercise, we paired the correct performance with incorrect ones that represent the most common mistakes.
- We propose a novel strategy that leverages similarity learning to identify joints that result in incorrect movement and require more attention from

the patients. The identification of those joints is self-supervised. We compose triplets of correctly and incorrectly performed exercises and force the model to learn embeddings that maximize the differences between the correctly and incorrectly performed exercises while minimizing the distances among the correctly performed ones.

- We invented a scoring system to evaluate the performance of models in terms of correctly identifying the joints-of-attention (JoA), which are the joints with physical relevance leading to incorrectness when doing an exercise. We provide informative visual feedback that highlights those joints for users to see "Where did I go wrong?".

The collected dataset, the 3D joint positions, and the attention scores per joint will be publicly available at <http://www.cs.bu.edu/faculty/betke/ExeChecker>.

## 2 Related Work

Movement assessment, also called action quality assessment (AQA), refers to the evaluation and analysis of physical movements or actions to determine their quality, efficiency, and effectiveness. This assessment is useful in various fields, including sports, physical therapy, and robotics. The assessment could be classifying the form of the movement [5, 13, 16, 27], or providing a final score representing the quality of the movements [3, 6, 10, 14, 17, 18, 22, 24].

In the field of sports, the quality of an action may be assessed by a judge or a panel of judges. Parmar *et al.* [18] constructed a dataset on Olympic events from YouTube videos and proposed the use of spatio-temporal features from 3D convolutional neural networks (C3D) for predicting quality assessment scores [18]. Their work was extended by considering assessment as a multitask [17]; besides computing an action quality score, the extended model also provides a summary of action events (e.g., number somersaults) and generates captions. Pan *et al.* [14] combined features extracted from local image patches around joints with pose information by building temporal and spatial joint-relation graphs to assess action performance. Tang *et al.* [24] proposed an uncertainty-aware score-distribution learning approach to address the intrinsic ambiguity in the score labels caused by different judges.

When action assessment scores are not readily available, movements can also be assessed by classifying the form, *i.e.*, whether it is correct or what error category it belongs to. Ogata *et al.* [13] proposed a ResNet-based model analyzing temporal distance matrices for the classification of squat exercises. Parmar *et al.* [16] introduced a Fitness-AQA dataset on weight lifting videos in gyms and proposed two self-supervised approaches, pose contrastive learning and motion disentangling, to assess errors in the execution of the exercises.

Contrastive learning has also been used to improve classification accuracy for the diagnosis of motor dysfunction in infants. Gao *et al.* [5] presented a motor assessment model to identify "fidgety movements" (FM). In their model, they included a branch trained with a triplet loss to differentiate positive (FM) and

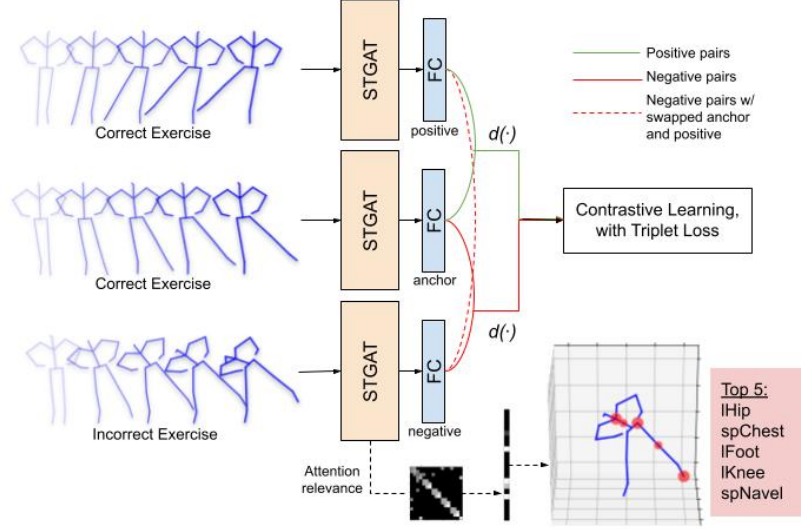
negative (non-FM) representations. In our work, we also designed a triplet loss, here to distinguish positive and negative examples of correct exercising.

In the field of rehabilitation, exercise performance is typically assessed by physical therapists. Data collection can be conducted in controlled environments, *e.g.*, a laboratory or clinic, using cameras and/or marker-based or markerless motion capture systems. The two datasets most widely used for research in exercise assessment are *KInematic Assessment of MOvement and Clinical Scores for Remote Monitoring of Physical REhabilitation* (KIMORE) [2] and the *University of Idaho-Physical Rehabilitation Movement Data* (UI-PRMD) [25]. Both datasets provide skeletal joint data. The KIMORE dataset includes RGB-D exercise recordings of a group of 34 participants who suffer from chronic motor disabilities and a group of 44 healthy participants. KIMORE provides score annotations where the quality of the exercise performance is evaluated by physical therapists. The UI-PRMD involves joint data of 10 healthy participants conducting physical therapy movements in correct and incorrect forms. UI-PRMD does not include score annotations. Instead, in their subsequent work [10], the researcher group proposed a scoring function that maps the values of the performance metric into a movement quality score in the range between 0 and 1. They adopted the log-likelihood of a Gaussian Mixture Model as the performance metric for evaluating the movements [10]. Other works on rehabilitation exercise assessment exploited the inherent spatial and temporal relations in the skeletal joints and developed graph-based neural networks for quality assessment. More generally, in the field of human action recognition, various strategies that incorporate spatial and temporal information into graph convolutional networks or graph attention transformers have been investigated (survey [28]). Correspondingly, rehabilitation-orientated graph-based models with attention mechanisms have been proposed [3, 12, 20] and evaluated using UI-PRMD and/or KIMORE as benchmark datasets. In our study, we conducted experiments with UI-PRMD, in addition to our ExeCheck dataset, for their paired collection of correctly and incorrectly performed exercises.

### 3 Method

Given an exercise movement  $M = [\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_T]$  where  $\mathbf{x}_t \in \mathbb{R}^{N \times 3}$  is the human pose at time  $t$  and is represented by  $N$  skeletal joints in 3D positions. ExeChecker 1) classifies  $M$  as either correctly or incorrectly performed exercise; and 2) if  $M$  is of the incorrect form, ExeChecker outputs the body joints that contribute to the incorrectness.

The architecture of ExeChecker is a triplet network, shown in Fig. 2. It leverages the design of a state-of-the-art spatial-temporal graph transformer for feature extraction, *i.e.*, embedding generation (Section 3.1), and then uses the triplet training strategy (Section 3.2) to maximize the distance in the embedding space between correct and incorrect movements. After being trained, the output of the attention map from the feature extraction module is used to identify the Joint of Attention if the input sequence is an incorrect movement.



**Fig. 2:** The ExeChecker framework: The top five body joints are identified that distinguish correct and incorrect exercise movements using contrastive learning.

### 3.1 Feature Extraction with STGAT

Spatial-temporal graph transformers are designed for the task of skeleton-based action recognition. Since the core of human action recognition is to obtain discriminative feature representation [19], it aligns with our goal of finding discriminating feature embeddings that separate the correct and incorrect exercise movements. A variety of graph transformers has been proposed over the past years, and we recommend a recent survey [28] for readers interested in this topic. In our implementation, we adopted the same STGAT architecture as described in the work of Hu *et al.* [9] as our feature extractor. Compared to prior works, they added a local spatial-temporal feature aggregator, which enables the model to better capture local cross-spacetime joint relationships and makes it more sensitive to similar action pairs. Briefly, at a given timestamp  $t$ , each joint node in pose  $\mathbf{x}_t$  is not only connected to all the other joints (plus itself) in the current pose but also every joint in  $\tau$  neighboring poses. As a result, the cross-spacetime adjacency matrix at time  $t$  is  $A_\tau^t$  is  $\tau \times N \times N$  in dimension. By incorporating  $A_\tau^t$  into graph operations and applying the multi-head self-attention, the output of STGAT can be obtained as

$$Y^t = \sigma\left(\frac{1}{H} \sum_{h=1}^H W_h X_\tau^t A_{\tau,h}^t\right) \quad (1)$$

where  $X_\tau$  is a sequence of  $\tau$  poses centered at timestamp  $t$  and  $H$  is the number of attention heads. For our experiments, we set  $\tau = 3$  and  $H = 8$ .

### 3.2 Triplets Training

Triplet networks have proven to be suitable for applications where fine-grained distinction between classes is important, such as face recognition [7, 23] and fine-grained image retrieval [4, 29]. Triplet loss [26] is a triple-based ranking loss function defined as

$$\mathcal{L}(a, p, n) = \max(0, d(a, p) - d(a, n) + \mu), \quad (2)$$

where  $a, p$ , and  $n$  denote the anchor, positive and negative samples respectively,  $d(\cdot)$  is a distance function and  $\mu$  is a margin parameter. This ranking loss measures the violation of the ranking order of the embedded features inside the triplet. Pairwise distance metrics such as Euclidean distance or cosine similarity are valid distance functions.

The margin  $\mu$  in Equation (2) introduces an extra hyperparameter during the triplet training. With a margin too small, the network can quickly learn to satisfy the triplet constraint, leading to a triplet loss that drops to zero early in training. As a result, the learned embeddings may not be well separated, and the network fails to learn discriminative features for different classes. On the other hand, if the margin is too large, the network may find it hard to satisfy the triplet constraint and struggle to converge. In a separate study [8], a softmax function is applied to the distance pairs, creating a ratio measure. The proposed triplet ratio loss can be formulated as

$$\mathcal{L}(a, p, n) = \left( \frac{e^{d(a, p)}}{e^{d(a, p)} + e^{d(a, n)}} \right)^2. \quad (3)$$

In our experiments, we followed the in-triplet hard negative mining with anchor swap strategy [1]. Briefly, if  $d(p, n) < d(a, n)$ , we swap anchor and positive samples. This ensures that the hardest negative inside the triplet is used for backpropagation. We initially used the margin loss (Equation (2)) for training the triplets. The training was unstable, and the loss did not converge despite trying various margin values. Consequently, we use the ratio loss (Equation (3)) for our experiments.

### 3.3 Visual Feedback on Incorrect Joint Movement

In order to evaluate if the model correctly identifies the joints of attention, we design a scoring system described as follows.

Given an exercise  $e$  and the corresponding set of Joints of Attention,  $JoA_e$ , the number of elements inside denoted as  $N_e$  ( $N_e = |JoA|$ ) varies depending on the  $e$ . Denote the complete joints set that represents a skeletal pose as  $J = \{J_i\}, (i = 1, 2, \dots, N)$ , we have  $JoA \subset J$ . For each exercise, the  $i$ th joint is assigned a score  $s_i$  such that

$$s_i = \begin{cases} 1 & \text{if } J_i \in JoA \\ 0 & \text{otherwise.} \end{cases} \quad (4)$$

Our method computes a score for each joint. We first min-max normalize the scores to  $[0, 1]$ . Let  $\hat{s}_i$  denote the computed score for joint  $J_i$ , we calculate score of JoA as

$$S_{JoA} = \frac{1}{N_e} \sum_{i=1}^N s_i \times \hat{s}_i, \quad (5)$$

where  $S_{JoA} \in [0, 1]$ .

## 4 ExeCheck Dataset

The ExeCheck consists of RGB-D videos of 10 rehabilitation exercises performed by 7 healthy subjects. Each exercise has a paired performance in both correct and incorrect forms by the same subject with 5 movement repetitions.

**Data Collection.** We collaborated with a physical therapist from a rehabilitation science program from an accredited university, who selected ten routinely prescribed rehabilitation exercises for patients with Parkinson’s disease and demonstrated to us both their correct and incorrect performance (Table 1). Specifically, the incorrect form highlights the most common mistakes observed by therapists. A total of 7 healthy subjects were recruited, and their consent for participating in the study was obtained according to a protocol approved by the university’s Institutional Review Board.

Each subject performed 10 exercises wearing relatively tight-fitting clothes. They performed each exercise five times, first in a correct way and then in an incorrect way, by purposely making at least one of the common mistakes listed in Table 1. The data was collected in a lab setting using the MS Azure Kinect sensor. We recorded RGB-D videos for each exercise at a framerate of 30 fps.

**Data Processing and Annotation.** We processed the recordings offline via the Azure Kinect Body Tracking SDK [11] to extract skeletal joint information. For each video frame, we obtained the positions and orientations of 32 joints. We used the 3D coordinates of 21 of them for our experiments. The chosen joints include 17 joints compatible with the Human3.6M dataset format, plus the left and right hands and feet. We annotated each exercise with the joints of attention, where the correct exercise (CE) and incorrect exercise (IE) share the same set. We further segmented the repetitions of each exercise manually.

We will release our dataset, which includes both the original videos and processed joint files, as well as the repetition annotations.

## 5 Experiments and Results

### 5.1 Datasets

We conducted experiments on two datasets, our **ExeCheck** dataset (see Section 4) and the **UI-PRMD** dataset [25].

The original UI-PRMD dataset includes joint positions of 10 persons performing 10 exercises, each with 2 sequences for correct and incorrect movements.

**Table 1:** Joints of Attention in the ExeCheck dataset identified by our model. The letter in the brackets indicates the participant’s side (L: left, R: right) that is actively engaged in the exercise. If not indicated, both sides are engaged. For incorrectly performed exercises, we asked the participants to err on a specific side and the corresponding joints of attention are denoted by \*.

Exercise Name	Included Common Mistakes	Joints of Attention
Arm Circle (0-AC)	Hunching shoulders Not keeping arms extended straight	lShoulder, rShoulder, lElbow, rElbow
Forward Lunge (1-FL) [L]	Leaning forward Letting the knee go past the toe Slamming right knee on the ground	spChest, lFoot, rKnee
High Knee Raise (2-HKR) [L]	Unstable upper body and wobbling Not raising the knee high enough Slamming foot (dropping the leg too quickly)	spNavel, spChest, lKnee, lFoot
Hip Abduction (3-HA) [L]	Tilting the upper body to the opposite side for compensation	lHip, spNavel, spChest
Leg Extension (4-LE) [L]	Leaning the upper body forward Not fully extending the leg	spNavel, spChest, lHip, lKnee
Shoulder Abduction (5-SA)	Lifting the arms unequally and tilting the head for compensation Hunching shoulders	rShoulder*, rElbow*, neck lShoulder
Shoulder External Rotation (6-SER) [L]	Not keeping the elbow close to the body	lShoulder, lElbow
Shoulder Flexion (7-SF)	Not lifting arms equally high enough	lShoulder*, lElbow*
Side Step Squat (8-SSS) [L]	Rotating the upper body Leaning the upper body forward	lHip, rHip, spNavel, spChest
Squat (9-Sq)	Bending knees inward (valgus knees) Not keeping the back upright and straight	lKnee, rKnee, spChest, spNavel

In some sequences in the UI-PRMD dataset, the published joint positions, especially those in the lower body, model unrealistic movements and suffer from unstableness. To circumvent this issue, we re-processed the original RGB videos using the state-of-the-art off-the-shelf model for human pose estimation Pose-formerV2 [30]. We also manually segmented the repetitions of each exercise based on the video time stamps. The extracted human poses have 17 joints following the Human3.6M dataset format. We further annotated the corresponding JoAs for each exercise based on the description provided in the UI-PRMD paper. Their descriptions of the non-optimal movements, *i.e.*, the incorrect exercises, point out joints that need to be attended to. Moreover, their exercise set overlaps with ours, so some knowledge for JoA could be transferred. The difference mainly



lies in the different joint sets. Our ExeCheck dataset has two spine joints (spine navel and spine chest) while the Human3.6M format has only one (torso). As an alternative, if, for an exercise, the incorrect version is described as "Subject unable to maintain upright trunk posture [..]" we include 'torso' and 'neck' in the JoA set of this exercise.

We will release our re-processed UI-PRMD skeleton data as well as the annotations together with ExeCheck dataset.

## 5.2 Experimental Methodology

To train and evaluate ExeChecker, we split the ExeCheck dataset and used the extracted skeleton joints of subjects 1–6 for training and those of subject 7 for testing. We used the segmented sequences and augmented them by flipping the left and right sides, resulting in 120 sequences for correct and incorrect movements in the training set and 20 in the testing set per exercise. We then exhaustively composed all possible combinations of triplets. Anchors and positives are sampled from correct exercises regardless of orders due to the in-triplet swap strategy, and negatives are sampled from incorrect exercises. As a result, for each exercise in the ExeCheck,  $\binom{60}{2} \times \binom{60}{1} = 106,200$  triplets are used for training and  $\binom{10}{2} \times \binom{10}{1} = 450$  triplets are used in validation for each exercise.

For the UI-PRMD, we used the skeleton joints of subjects 1–9 of UI-PRMD for training and those of subject 10 for testing, which corresponds to approximately (some movement sequences do not have 10 repetitions) 180 sequences per correct and incorrect exercises in the training set and 20 per correct and incorrect exercises in the testing set after repetition segmentation and mirror augmentation. We then composed triplets in a similar way as we did for the ExeCheck dataset, resulting in approximately 2.8 million ( $\binom{180}{2} \times \binom{180}{1}$ ) triplets for training and 3800 ( $\binom{20}{2} \times \binom{20}{1}$ ) triplets for validation for each exercise. We randomly/uniformly sample 160 frames and then randomly/centrally crop 128 frames for training/test splits for the ExeCheck dataset and 104 then 88 for the UI-PRMD dataset. All the models are trained using the AdamW optimizer with a start learning rate of 0.001. We step the learning rate scheduler exponentially at a ratio of 0.9 twice every epoch. All models are trained on a 48G GPU node for a maximum of 10 epochs. We use a batch size of 20 for both datasets.

We use the same configuration as used by Hu *et al.* [9] for extracting features with STGAT. We set the dimension of the embedding space for triplets training to 128 for all the models. For calculating the JoA scores, we pull the cross-spacetime attention map of the input sequence, take the average of the 8 heads, use the center frame as suggested by the authors, and sum the attention values each joint received as their raw joint scores.

## 5.3 Average JoA Results

We report the results of a baseline that uses the hop-adjusted Canonical Time Warping (CTW) [31]. CTW is an extension of Dynamic Time Warping (DTW)

that incorporates canonical correlation analysis as a measure of spatial alignment. Zhou and De La Torre showed that CTW outperforms DTW in human motion alignment [31]. Following their method, we first aligned the segmented correct (CE) and incorrect exercises (IE). For a fair comparison, we used the normalized sequences as inputs for ExeChecker for the CTW as well. The aligned path is then used for indexing the corresponding time points in CE and IE. We aggregated the Euclidean distance per joint between the aligned CE and IE and considered those distances as the joint’s raw scores. A higher score, *i.e.*, a larger distance, suggests that the joint in IE is further away from CE, indicating an incorrect movement. For JoA evaluation, our method min-max normalizes the raw scores to zero and one.

In our preliminary studies, we found that terminal joints like hands and feet tend to always have higher values and yield low JoA scores for both ExeCheck and UI-PRMD datasets. We note that distal joints with higher degrees of freedom compared to proximal joints are biased to higher scores due to the nature of the alignment. To correct for this bias, we adjusted the raw joint score  $s^r$  obtained from CTW by its number of hops to the root (pelvis) such that the adjusted score  $s_i$  of joint  $i$  is equal to  $s_i^r/(h_i+1)$ , with  $h_i$  being its hop number. This adjustment increased the JoA scores for most of the exercises. We use the adjusted CTW joint scores to calculate JoA scores as a baseline and the visualizations.

We report average JoA scores of the 10 exercises in ExeCheck dataset in Table 2 and those of UI-PRMD in Table 3.

**Table 2:** Average JoA scores per exercise on ExeCheck Dataset. ExeChecker outperforms CTW for 9 of 10 exercises.

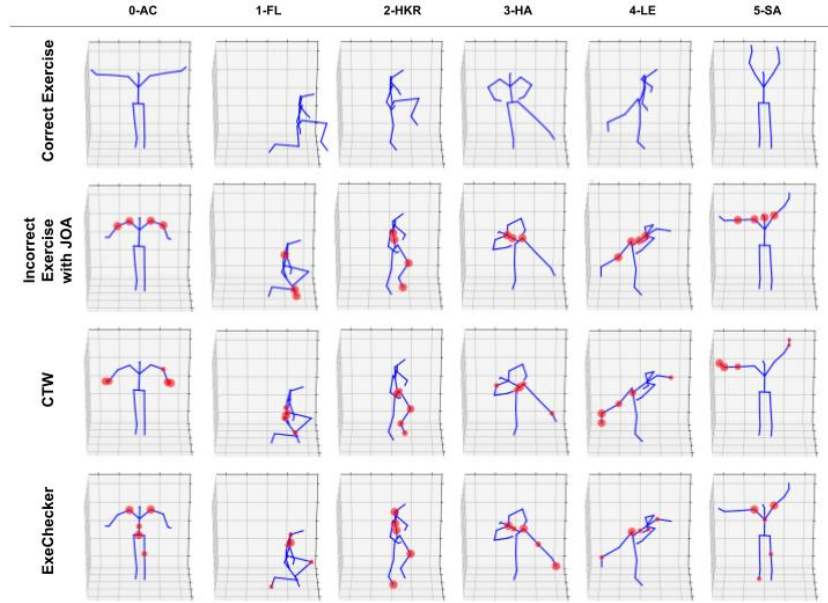
Exercises	0-AC	1-FL	2-HKR	3-HA	4-LE	5-SA	6-SER	7-SF	8-SSS	9-Sq
CTW	0.281	0.208	0.356	0.254	0.299	0.189	0.588	0.259	0.346	0.367
ExeChecker	<b>0.429</b>	<b>0.403</b>	<b>0.402</b>	<b>0.633</b>	<b>0.388</b>	<b>0.223</b>	0.207	<b>0.373</b>	<b>0.387</b>	<b>0.488</b>

**Table 3:** Average JoA scores per exercise on UI-PRMD [25]. We used the same order of exercises as in Table 1. ExeChecker outperforms CTW for all but one exercise.

Exercises	m01	m02	m03	m04	m05	m06	m07	m08	m09	m10
CTW	0.414	0.479	0.473	0.165	0.305	0.382	0.183	0.550	0.278	0.295
ExeChecker	<b>0.543</b>	0.454	<b>0.517</b>	<b>0.517</b>	<b>0.447</b>	<b>0.733</b>	<b>0.539</b>	<b>0.614</b>	<b>0.543</b>	<b>0.664</b>

#### 5.4 Qualitative Results

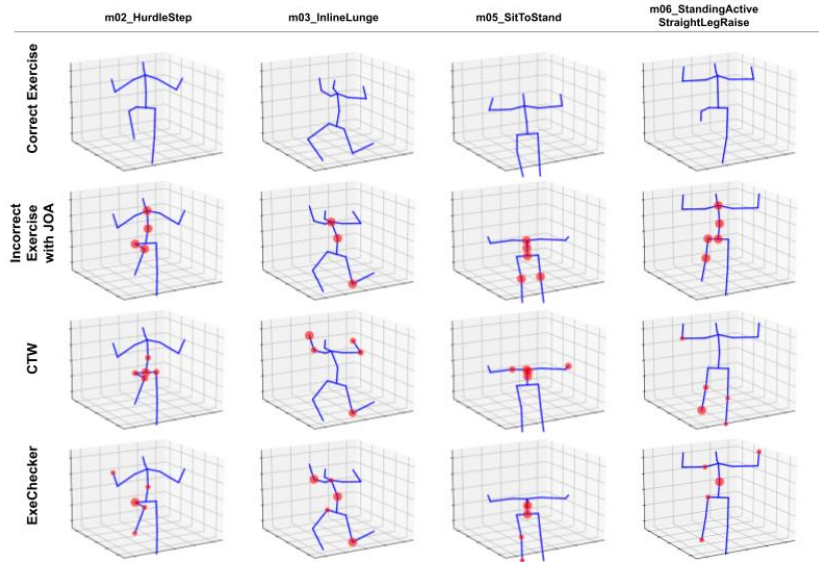
We present our qualitative results on the ExeCheck dataset in Figure 3 and results on the UI-PRMD dataset in Figure 4. Both figures show that ExeChecker



**Fig. 3: Qualitative results on the ExeCheck dataset.** The first row shows the correct form of each exercise, and the following three rows show an example of the incorrect form per exercise. The skeletal joints in the second row are overlaid with the corresponding JoAs (Table 1) shown as red circles of equal sizes. For CTW and ExeChecker, we highlight the top 5 joints with the highest attention scores evaluated on the incorrect exercise. The size of the circle is proportional to the computed score. Note that the number of JoAs varies among exercises; hence, the number of red circles in row 2 differs between exercises. The JoAs and their scores identified by ExeChecker (row 4) are better than those identified by CTW (row 3) in matching the expert-defined JoAs shown in row 2.

in general focuses on the right joints that need attention and such feedback is more informative than the CTW. Specifically, as shown in the first column in the Figure 3, 0-AC (Arm Circle), although missing the elbows, ExeChecker correctly focuses on the shoulders with great attention, which is more appropriate than CTW’s output; CTW assigns the highest joint scores to the wrists and hands despite the hop-adjustment. The second column of 1-FL (Forward Lunge) presents an interesting case. We list the left foot as one of the joints of attention for the forward lunge because a too-small step is the reason causing the knee to pass the toe, a common error that can put too much stress on the knee joint and could eventually lead to knee injury. Although the ExeChecker did not catch the left foot, it highlighted the left knee as shown in the example.

The third column, 2-HKR (High Knee Raise) shows that ExeChecker not only highlights the knee, which in this incorrectly performed exercise was not raised to enough height, but also focuses on the joints along the spine (spNavel, spChest,



**Fig. 4: Qualitative results on the UI-PRMD** Similar to Figure 3, we show and compare our results with CTW and JoA. Compared to the CTW, ExeChecker is good at detecting the torso joint, which is of key importance because it indicates that the subject did not maintain the body trunk upright, a very common mistake for a variety of exercises.

and neck being extra), which is proper feedback for a wobbling upper body. CTW is only able to catch the knee problem but fails to detect the wobbling. The column of 3-HA further demonstrated the strength of ExeChecker in detecting the tilting of the upper body with additional attention to the leg in motion. The column of 4-LE shows that ExeChecker pays the heaviest attention to the correct hip that is responsible for the leg extension. It still catches the upper body leaning forward problem in the top 5 but with less attention.

The following column of 5-SA shows that ExeChecker correctly pays the highest attention to the shoulders; it then assigns attention to the chest and lower body instead of the neck and the elbows though. Overall, ExeChecker is able to highlight the subtle but err-leading movements like hunching shoulders and unstable upper body, which are usually missed by CTW, which is based on pairwise alignment. For the plots of exercises 6–9 of ExeCheck, please refer to Supplementary Figure 2.

Qualitative results by models trained on UI-PRMD are shown in Figure 4. In the first column, the incorrect hurdle step features an non-neutral femur position caused by the right hip and a slightly bent body trunk. ExeChecker highlights the hip with the greatest attention and then the leg (knee, ankle) and torso. In comparison, CTW assigns the greatest joint score to the pelvis. In the second column, the inline lunge, both ExeChecker and CTW highlight the rear knee,

which should not touch the ground. However, CTW does not catch the torso, which should be highlighted because the subject did not maintain an upright trunk posture. The third column is about sit-to-stand. The subject incorrectly raised the pelvis while standing up. CTW catches the torso and neck but misses the pelvis. It instead assigns a score to the head. ExeChecker on the other hand, misses the neck but it correctly focuses on the pelvis, which is the more important one. In addition, the subject here had a slight knee valgus collapse on the right side and ExeChecker pays moderate attention to that knee. In the last column, straight leg raise, the incorrect movement is about the body trunk not being upright, deviated pelvis, and insufficient right hip flexion, which leads to not raising the leg high enough. CTW misses the torso again and focuses mostly on the right ankle and then the knee. On the other hand, ExeChecker notices the hip after assigning most of the attention to the torso.

### 5.5 Ablation Studies

We investigated the effects of different distance functions used in the triplet loss, which revealed that no distance function is a clear winner for all exercises (Table 4).

We also investigated the effects of bone orientations and report our results in the same table. We decouple the skeleton and each joint essentially is replaced by the bone vector from its parent to itself. That is, joint  $J_i$  instead of  $(x_i, y_i, z_i)$  is now represented by  $(\delta x_i, \delta y_i, \delta z_i)$ , where  $(\delta x_i, \delta y_i, \delta z_i) = (x_i, y_i, z_i) - (x_{f(i)}, y_{f(i)}, z_{f(i)})$ ,  $f(\cdot)$  denotes a functional relationship and  $f(i)$  is the parent joint of  $J_i$  following the kinematic chain. The root joint, pelvis ( $J_0$ ) has  $(\delta x_0, \delta y_0, \delta z_0) = (0, 0, 0)$ . In general, the decoupled sequence shows degraded JoA scores possibly due to the information loss when subtracting the parent joints. Arguably, the cosine similarity works better with the decoupled sequences.

**Table 4:** Average JoA scores per exercise on ExeCheck Dataset. ExeChecker uses Euclidean distances for computing triplet loss. EC-cos use cosine similarity for computing the pairwise distances between (anchor, positive) and (anchor, negative) pairs. Row 3 and 4 use bone vectors as the model inputs with EC-Bone using Euclidean distances and EC-Bone-cos using cosine similarities as their distance metric.

Exercises	0-AC	1-FL	2-HKR	3-HA	4-LE	5-SA	6-SER	7-SF	8-SSS	9-Sq
ExeChecker	0.429	0.403	0.402	0.633	0.388	0.223	0.207	0.373	0.387	0.488
EC-cos	0.295	0.397	0.638	0.274	0.444	0.674	0.591	0.377	0.588	0.608
EC-Bone	0.031	0.493	0.038	0.145	0	0.536	0	0.343	0.114	0.004
EC-Bone-cos	0.123	0.376	0.034	0.439	0.250	0.814	0.717	0.540	0.513	0.154

## 6 Conclusions

In this paper, we presented ExeChecker, a framework for identifying and visualizing the joints of attention in incorrectly conducted exercises. We further proposed a JoA scoring method to quantitatively evaluate the identification. Both quantitative and qualitative results show that ExeChecker in general focuses on the right joints. Benefiting from contrastive learning using triplets, ExeChecker can detect small yet crucial joint movements that make the exercise incorrect, like shoulder joints for struggling shoulders, and spine joints for wobbling, leaning, or tilting of the upper body. Those small movements are usually missed by the baseline method CTW. Importantly, the composition of the triplets is based on the paired correct and incorrect sequences, provided by our ExeCheck dataset. As to the best of our knowledge, ExeCheck and UI-PRMD are the only two datasets that have paired movement sequences. Moreover, in addition to skeletal sequences, ExeCheck also provides RGB-D videos. A limitation of the ExeCheck dataset is that it only includes a set of common mistakes and the JoA labels are specific to those mistakes. Future work could expand the dataset by adding incorrect variants and annotating the corresponding JoAs. A user study is planned that analyzes how helpful the visual feedback showing "Where did I go wrong?" is for exercising participants, particularly Parkinson's patients. We plan to incorporate the ExeChecker visualization into the ExerciseCheck platform [15] for this user study.

## Acknowledgements

We thank the participants of our data collection study, and Mona Jalal and Shreya Pandit for helping with the data acquisition. We are grateful to Timothy Nordahl, Doctor of Physical Therapy, for his instructions on the correct performance of rehabilitation exercises in our study, and for his advice on how our participants should move to illustrate typical incorrect performance of these exercises. We thank Professor Aleksandar Vakanski, University of Idaho, for kindly providing us access to the original videos of the UI-PRMD dataset.

## References

1. Baltas, V., Riba, E., Ponsa, D., Mikolajczyk, K.: Learning local feature descriptors with triplets and shallow convolutional neural networks. In: British Machine Vision Conference (BMVC) (2016). <https://doi.org/10.5244/C.30.119>, 11 pages
2. 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 Transactions on Neural Systems and Rehabilitation Engineering* **27**(7), 1436–1448 (July 2019). <https://doi.org/10.1109/TNSRE.2019.2923060>
3. Deb, S., Islam, M.F., Rahman, S., Rahman, S.: Graph convolutional networks for assessment of physical rehabilitation exercises. *IEEE Transactions on Neural Systems and Rehabilitation Engineering* **30**, 410–419 (2022). <https://doi.org/10.1109/TNSRE.2022.3150392>

4. D’Innocente, A., Garg, N., Zhang, Y., Bazzani, L., Donoser, M.: Localized triplet loss for fine-grained fashion image retrieval. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops. pp. 3910–3915 (June 2021)
5. Gao, Q., Yao, S., Tian, Y., Zhang, C., Zhao, T., Wu, D., Yu, G., Lu, H.: Automating General Movements Assessment with quantitative deep learning to facilitate early screening of cerebral palsy. *Nature Communications* **14**(8294), 1–11 (2023). <https://doi.org/10.1038/s41467-023-44141-x>
6. Gu, Y., Pandit, S., Saraee, E., Nordahl, T., Ellis, T., Betke, M.: Home-based physical therapy with an interactive computer vision system. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) Workshops (Oct 2019). <https://doi.org/10.1109/ICCVW.2019.00320>
7. Hermans, A., Beyer, L., Leibe, B.: In defense of the triplet loss for person re-identification. *arXiv preprint arXiv:1703.07737* (2017). <https://doi.org/10.48550/arXiv.1703.07737>
8. Hoffer, E., Ailon, N.: Deep metric learning using triplet network. In: Similarity-based pattern recognition: Third International Workshop, SIMBAD 2015, Copenhagen, Denmark, October 12–14, 2015. Proceedings 3. pp. 84–92. Springer (2015). [https://doi.org/10.1007/978-3-319-24261-3\\_7](https://doi.org/10.1007/978-3-319-24261-3_7)
9. Hu, L., Liu, S., Feng, W.: Skeleton-based action recognition with local dynamic spatial-temporal aggregation. *Expert Systems with Applications* **232**, 120683 (2023). <https://doi.org/10.1016/j.eswa.2023.120683>
10. Liao, Y., Vakanski, A., Xian, M.: A deep learning framework for assessing physical rehabilitation exercises. *IEEE Transactions on Neural Systems and Rehabilitation Engineering* **28**(2), 468–477 (2020). <https://doi.org/10.1109/TNSRE.2020.2966249>
11. Microsoft Azure Kinect, <https://learn.microsoft.com/en-us/azure/kinect-dk/>, accessed July 2024
12. Mourchid, Y., Slama, R.: D-stgcnt: A dense spatio-temporal graph conv-gru network based on transformer for assessment of patient physical rehabilitation. *Computers in Biology and Medicine* **165**, 107420 (2023). <https://doi.org/10.1016/j.combiomed.2023.107420>
13. Ogata, R., Simo-Serra, E., Iizuka, S., Ishikawa, H.: Temporal distance matrices for squat classification. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops (June 2019)
14. 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 (ICCV). pp. 6331–6340 (2019). <https://doi.org/10.1109/ICCV.2019.00643>
15. Pandit, S., Tran, S., Gu, Y., Saraee, E., Jansen, F., Singh, S., Cao, S., Sadeghi, A., Shandelman, E., Ellis, T., et al.: ExerciseCheck: A scalable platform for remote physical therapy deployed as a hybrid desktop and web application. In: Proceedings of the 12th ACM International Conference on Pervasive Technologies Related to Assistive Environments. pp. 101–109 (2019). <https://doi.org/10.1145/3316782.3321537>
16. Parmar, P., Gharat, A., Rhodin, H.: Domain knowledge-informed self-supervised representations for workout form assessment. In: European Conference on Computer Vision (ECCV). pp. 105–123. Springer (2022). [https://doi.org/10.1007/978-3-031-19839-7\\_7](https://doi.org/10.1007/978-3-031-19839-7_7)
17. 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 Confer-

- ence on Computer Vision and Pattern Recognition (CVPR). pp. 304–313 (2019). <https://doi.org/10.1109/CVPR.2019.00039>
18. Parmar, P., Tran Morris, B.: Learning to Score Olympic Events. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW). pp. 20–28 (2017). <https://doi.org/10.48550/arXiv.1611.05125>
  19. Presti, L.L., La Cascia, M.: 3d skeleton-based human action classification: A survey. *Pattern Recognition* **53**, 130–147 (2016). <https://doi.org/10.1016/j.patcog.2015.11.019>
  20. Réby, K., Dulau, I., Dubrasquet, G., Aimar, M.B.: Graph transformer for physical rehabilitation evaluation. In: 2023 IEEE 17th International Conference on Automatic Face and Gesture Recognition (FG). pp. 1–8. IEEE (2023). <https://doi.org/10.1109/FG57933.2023.10042778>
  21. Saraee, E., Gu, Y., Pandit, S., Tran, S., Shandelman, E., Singh, S., Nordahl, T.J., Ellis, T., Betke, M.: ExerciseCheck: Data analytics for a remote monitoring and evaluation platform for home-based physical therapy. In: Proceedings of the 12th ACM International Conference on Pervasive Technologies Related to Assistive Environments. pp. 110–118 (2019). <https://doi.org/10.1145/3316782.3316790>
  22. Sardari, F., Paiement, A., Hannuna, S., Mirmehdi, M.: Vi-net—view-invariant quality of human movement assessment. *Sensors* **20**(18), 5258 (2020). <https://doi.org/10.3390/s20185258>
  23. Schroff, F., Kalenichenko, D., Philbin, J.: FaceNet: A unified embedding for face recognition and clustering. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 815–823 (2015)
  24. 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 (CVPR). pp. 9839–9848 (2020). <https://doi.org/10.1109/CVPR42600.2020.00986>
  25. Vakanski, A., Jun, H.P., Paul, D., Baker, R.: A data set of human body movements for physical rehabilitation exercises. *Data* **3**(1) (2018). <https://doi.org/10.3390/data3010002>, 15 pages
  26. Wang, J., Song, Y., Leung, T., Rosenberg, C., Wang, J., Philbin, J., Chen, B., Wu, Y.: Learning fine-grained image similarity with deep ranking. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (June 2014)
  27. Wang, X., Li, J., Hu, H.: Skeleton-based action quality assessment via partially connected lstm with triplet losses. In: Chinese Conference on Pattern Recognition and Computer Vision (PRCV). pp. 220–232. Springer (2022). [https://doi.org/10.1007/978-3-031-18913-5\\_17](https://doi.org/10.1007/978-3-031-18913-5_17)
  28. Xin, W., Liu, R., Liu, Y., Chen, Y., Yu, W., Miao, Q.: Transformer for skeleton-based action recognition: A review of recent advances. *Neurocomputing* **537**, 164–186 (2023). <https://doi.org/10.1016/j.neucom.2023.03.001>
  29. Yao, X., She, D., Zhang, H., Yang, J., Cheng, M.M., Wang, L.: Adaptive deep metric learning for affective image retrieval and classification. *IEEE Transactions on Multimedia* **23**, 1640–1653 (2020). <https://doi.org/10.1109/TMM.2020.3001527>
  30. Zhao, Q., Zheng, C., Liu, M., Wang, P., Chen, C.: PoseFormerV2: Exploring Frequency Domain for Efficient and Robust 3D Human Pose Estimation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 8877–8886 (2023). <https://doi.org/10.48550/arXiv.2303.17472>
  31. Zhou, F., de la Torre, F.: Canonical time warping for alignment of human behavior. *Advances in Neural Information Processing Systems (NIPS)* **22** (2009), 9 pages.