

3D Pose Based Feedback For Physical Exercises

Ziyi Zhao^{1,2} Sena Kiciroglu¹ Hugues Vinzant^{3*} Yuan Cheng^{4*},
Isinsu Katircioglu^{5*}, Mathieu Salzmann^{1,6} Pascal Fua¹

¹CVLab, EPFL, Switzerland ²Tongji University, China

³a-gO, France ⁴KTH, Sweden

⁵SDSC, Switzerland ⁶Clearspace, Switzerland

ziyi.zhao@epfl.ch sena.kiciroglu@epfl.ch

Unsupervised self-rehabilitation exercises and physical training can cause serious injuries if performed incorrectly. We introduce a learning-based framework that identifies the mistakes made by a user and proposes corrective measures for easier and safer individual training. Our framework does not rely on hard-coded, heuristic rules. Instead, it learns them from data, which facilitates its adaptation to specific user needs. To this end, we use a Graph Convolutional Network (GCN) architecture acting on the user's pose sequence to model the relationship between the the body joints trajectories. To evaluate our approach, we introduce a dataset with 3 different physical exercises. Our approach yields 90.9% mistake identification accuracy and successfully corrects 94.2% of the mistakes. Our code and dataset are available at <https://github.com/Jacoo-Zhao/3D-Pose-Based-Feedback-For-Physical-Exercises>.

1 Introduction

Being able to perform exercises without requiring the supervision of a physical trainer is a convenience many people enjoy, especially after the COVID-19 pandemic. However, the lack of effective supervision and feedback can end up doing more harm than good, which may include causing serious injuries. There is therefore a growing need for computer-aided exercise feedback strategies.

A few recent works have addressed this problem [1,2,3,4,5,6]. However, they focus only on identifying whether an exercise is performed correctly or not [1,6], or they rely on hard-coded rules based on joint angles that cannot easily be extended to new exercises [2,3,4]. In this work, we therefore leverage recent advances in the fields of pose estimation [7,8,9], action recognition [10,11] and motion prediction [12,13,14] to design a framework that provides automated and personalized feedback to supervise physical exercises.

Specifically, we developed a method that not only points out mistakes but also offers suggestions on how to fix them without relying on hard-coded, heuristic rules to define what a successful exercise sequence should be. Instead, it learns from data. To this end, we use a two-branch deep network. One branch is an action classifier that tells users what kind of errors they are making. The other proposes corrective measures. They both rely on Graph Convolutional Networks (GCNs) that can learn to exploit the

*Former members of CVLab, EPFL.

relationships between the trajectories of individual joints. Fig. 1 depicts the kind of output our network produces.

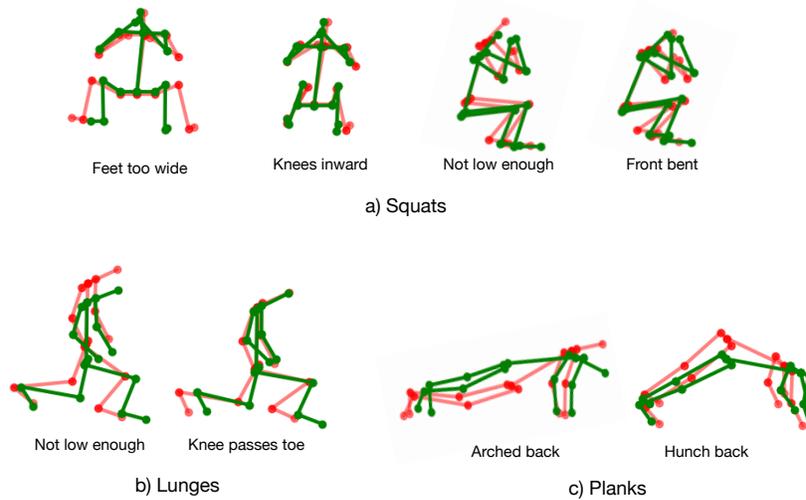


Fig. 1. Example results from our framework depicting frames from the a) squat, b) lunge, and c) plank classes. The red poses correspond to the exercises performed incorrectly while the green poses correspond to our corrections. Note that although we display a single pose from each mistake type, our framework operates on entire sequences.

To showcase our framework’s performance, we recorded a physical exercise dataset with 3D poses and instruction label annotations. Our dataset features 3 types of exercises; squats, lunges and planks. Each exercise type is performed correctly and with mistakes following specific instructions by 4 different subjects. Our approach achieves 90.9% mistake recognition accuracy on a test set. Furthermore, we use the classification branch of our framework to evaluate the performance of the correction branch, considering the correction to be successful if the corrected motion is classified as “correct”. Under this metric, our approach successfully corrects 94.2% of users’ mistakes. We will make our code and dataset publicly available upon acceptance.

2 Related Work

Our work is at the intersection of several sub-fields of computer vision: (i) We draw inspiration from GCN based **human motion prediction** architectures; (ii) we identify the users’ mistakes in an **action recognition** fashion; and (iii) we address the task of **physical exercise analysis**. We therefore discuss these three topics below.

2.1 Human Motion Prediction

Human motion prediction is a complex task due to the inherent uncertainty in forecasting into the future. In recent years, many deep learning methods based on recurrent neural networks (RNNs) [15,16,17,18,19,20,14], variational auto encoders (VAEs) [21,22,23,24], transformers [12], and graph convolutional networks (GCNs) [25,26,13,27] have been proposed. We focus our discussion on GCN based ones, as we also exploit the graph-like connections of human joints with GCNs in our approach.

In [25], Mao *et al.* proposed to exploit a GCN to model the relationships across joint trajectories by representing them with Discrete Cosine Transform (DCT) coefficients. The approach was in [13] by integrating an attention module, and in [27,28] by using cross-subject attention for multi-person motion prediction. In [26], the input coefficients to the GCN were extracted via an inception module instead of the DCT. Our motion correction branch is inspired by [25], but instead of forecasting future motion, we predict correctly performed exercises.

2.2 Action Recognition

Although there is a vast literature on image-based action recognition, here we focus on its skeleton-based counterpart, as our approach also processes 3D poses. Early deep learning based approaches to skeleton-based action recognition mostly relied on RNNs [29,30,31,32,33]. Li *et al.* [10] used convolutional neural networks (CNNs) to extract features hierarchically by first finding local point-level features and gradually extracting global spatial and temporal features. Zhang *et al.* [34] designed CNN and RNN networks that are robust to viewpoint changes.

Recently, [35,36,11] employed GCNs for action recognition. Specifically, Tang *et al.* [35] designed a reinforcement learning scheme to select the most informative frames and feed them to a GCN. Li *et al.* [36] developed a GCN framework that not only models human joint connections, but also learns to infer “actional-links”, which are joint dependencies learned from the data.

Zhang *et al.* [11] designed a two-module network, consisting of a first GCN-based module that extracts joint-level information and a second frame-level module capturing temporal information via convolutional layers and spatial and temporal max-pooling. Our classification branch borrows ideas from Mao *et al.*'s [25] and Zhang *et al.*'s [11] architectures. It is composed of graph convolutional blocks as proposed by Mao *et al.* [25] combined with the frame-level module architecture proposed by Zhang *et al.* [11].

2.3 Physical Exercise Analysis

Physical exercise analysis aims to prevent injuries that may arise when a person performs motions incorrectly. In its simplest form, such an analysis amounts to detecting whether the subject performs the exercise correctly or not. This was achieved several works [1,6,5] by exploiting 2D poses extracted from the input images. In particular, Dittakavi *et al.* [5] detected which joints need to be fixed by finding the overall joint angle distribution of the dataset and detecting poses in which a joint angle is an anomaly. This framework operates on single frames, as opposed to our method which operates

on entire sequences. In [37], Zell *et al.* represented the human body as a mass-spring model and analyzed the extension torque on certain joints, allowing them to classify whether a motion is performed correctly or not. While useful, such classification-based approaches offer limited information to the user, as they do not provide them with any feedback about the specific type of mistakes they made. Moreover, most of existing works operate on 2D pose inputs [1,2,3,6,5]. Similar to [4], we also design our framework to work with 3D poses enabling us to be robust to ambiguities found in 2D poses.

While a few works took some steps toward giving feedback [2,3,4], this was achieved in a hard-coded fashion, by thresholding angles between some of the body joints. As such, this approach relies on manually defining such thresholds, and thus does not easily extend to new exercises. Furthermore, it does not provide the user personalized corrective measures in a visual manner, by demonstrating the correct version of their performance. We address this by following a data driven approach able to automatically learn the different “correct” forms of an exercise, and that can easily extend to different types of exercises and mistakes. To the best of our knowledge, our framework is the first to both identify mistakes and suggest personalized corrections to the user.

3 Methodology

Before we introduce our framework in detail, let us formally define the tasks of motion classification and correction. Motion classification seeks to predict the action class c of a sequence of 3D poses from $t = 1$ to $t = N$, denoted as $\mathbf{P}_{1:N}$. We can write this as

$$c = F_{\text{class}}(\mathbf{P}_{1:N}) ,$$

where F_{class} is the classification function.

We define motion correction as the task of finding the “correct” version of a sequence, which can be written as

$$\hat{\mathbf{P}}_{1:N} = F_{\text{corr}}(\mathbf{P}_{1:N}) ,$$

where F_{corr} is the correction function and $\hat{\mathbf{P}}_{1:N}$ is the corrected sequence. Ideally, the corrected sequence should be of class “correct”. We can use the classification function to verify that this is the case, i.e.,

$$c_{\text{correct}} = F_{\text{class}}(\hat{\mathbf{P}}_{1:N}) ,$$

where c_{correct} is the label corresponding to a correctly performed exercise.

Given these definitions, we now describe the framework we designed to address these tasks and discuss our training and implementation details.

3.1 Exercise Analysis Framework

Our framework for providing exercise feedback relies on GCNs and consists of two branches: One that predicts whether the input motion is correct or incorrect, specifying the mistake being made in the latter case, and one that outputs a corrected 3D pose

sequence, providing a detailed feedback to the user. We refer to these two branches as the “classifier” and “corrector” models, respectively.

Inspired by Mao *et al.* [25], we use the DCT coefficients of joint trajectories, rather than the 3D joint positions, as input to our model. This allows us to easily process sequences of different lengths. The corrector model outputs DCT coefficient residuals, which are then summed with the input coefficients and undergo an inverse-DCT transform to be converted back to a series of 3D poses.

To reduce the time and space complexity of training the classifier and the corrector separately and to improve the accuracy of the model, we combine the classification and correction branches into a single end-to-end trainable model. Figure 2 depicts our overall framework. It takes the DCT coefficients of each joint trajectory as input. The first layers are shared by the two models, and the framework then splits into the classification and correction branches.

Furthermore, we feed the predicted action labels coming from the classification branch to the correction branch. We depict this in Figure 2 as the “Feedback Module”. Specifically, we first find the label with the maximum score predicted by the classification branch, convert this label into a one-hot encoding, and feed it to a fully-connected layer. The resulting tensor is concatenated to the output of the first graph convolutional blocks (GCB) of the correction branch. This process allows us to explicitly provide label information to the correction module, enabling us to further improve the accuracy of the corrected motion.

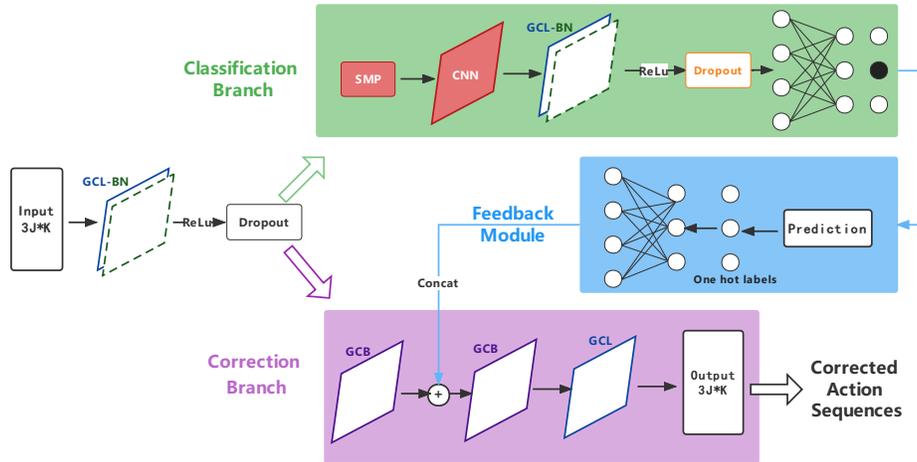


Fig. 2. Our framework consists of a classification and a correction branch. They share several graph convolutional layers are then split such that the classification branch identifies the type of mistakes made by the user and the correction branch outputs a corrected pose sequence. The result of the classification branch is fed to the correction branch via a feedback module.

Implementation and Training Details. We primarily use GCB similar to those presented in [25] in our network architecture, depicted in Figure 3. These modules allow us to learn the connectivity between different joint trajectories. Each graph convolutional layer is set to have 256 hidden features. Additionally, our classification branch borrows ideas from Zhang *et al.*'s [11] action recognition model. It is a combination of GCB modules and the frame-level module architecture of [11] consisting of convolutional layers and spatial max-pooling layers.

We train our network in a supervised manner by using pairs of incorrectly performed and correctly performed actions. However, it is not straightforward to find these pairs of motions. The motion sequences are often of different lengths, and we face the task of matching incorrectly performed actions to the closest correctly performed action from the same actor. To do so, we make use of Dynamic Time Warping (DTW) [38], which enables us to find the minimal alignment cost between two time series of different lengths, using dynamic programming. We compute the DTW loss between each incorrect and correct action pair candidate and select the pair with the smallest loss value.

We use the following loss functions to train our model.

- E_{corr} : The loss of the correction branch, which aims to minimize the soft-DTW [39] loss between the corrected output sequence and the closest correct motion sequence, determined as described previously. The soft-DTW loss is a differentiable version of the DTW loss, implemented by replacing the minimum operation by a soft minimum.
- E_{smooth} : The smoothness loss on the output of the correction branch, to ensure the produced motion is smooth and realistic. It penalizes the velocities of the output motion by imposing an L2 loss on them.
- E_{class} : The loss of the classification branch, which aims to minimize the cross entropy loss between the predicted logits and the ground-truth instruction label.

We combine these losses into

$$E_{\text{loss}} = w_{\text{corr}}E_{\text{corr}} + w_{\text{class}}E_{\text{class}} + w_{\text{smooth}}E_{\text{smooth}}, \quad (1)$$

where E_{loss} is the overall loss and w_{corr} , w_{class} , w_{smooth} are the weights of the correction, classification, and smoothness losses, respectively. For our experiments we set $w_{\text{corr}} = 1$, $w_{\text{class}} = 1$, and $w_{\text{smooth}} = 1e - 3$.

During training, we use curriculum learning in the feedback module: Initially the ground-truth instruction labels are given to the correction branch. We then use a scheduled sampling strategy similar to [40], where the probability of using the ground-truth labels instead of the predicted ones decreases from 1 to 0 linearly as the epochs increase. In other words, the ground-truth labels are progressively substituted with the labels predicted by the classification branch, until only the predicted labels are used. During inference, only the predicted labels are given to the correction branch.

We use Adam [41] as our optimizer. The learning rate is initially set to 0.01 and decays according to the equation $\text{lr} = 0.01 \cdot 0.9^{i/s}$, where lr is the learning rate, i is the epoch and s is the decay step, which is set to 5. To increase robustness and avoid overfitting, we also use drop-out layers with probability 0.5. We use a batch size of 32 and train for 50 epochs.

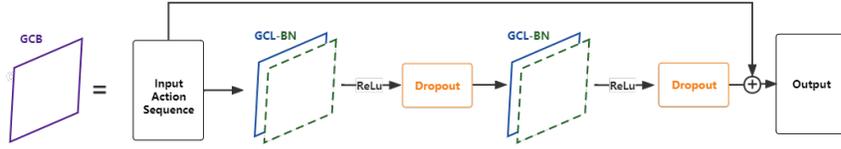


Fig. 3. Graph Convolutional Block (GCB) consisting of graph convolutional layers, batch normalization layers, ReLUs and drop-outs.

4 EC3D Dataset

Existing sports datasets such as Yoga-82 [42], FineGym [43], FSD-10[44], and Diving48 [45] often include correct performances of exercises but do not include incorrect sequences. They are also not annotated with 3D poses. Therefore to evaluate our approach, we recorded and processed a dataset of physical exercises performed both correctly and incorrectly, and named the “EC3D” (**E**xercise **C**orrection in **3D**) dataset.

Specifically, this dataset contains 3 types of actions, each with 4 subjects who repeatedly performed a particular correct or incorrect motion as instructed. We show the number of sequences per action and the instructions for each subject in Table 1. The dataset contains a total of 132 squat, 127 lunge, and 103 plank action sequences, split across 11 instruction labels.

The videos were captured by 4 GoPro cameras placed in a ring around the subject, using a frame rate of 30 fps and a 1920×1080 image resolution. Figure 4 depicts example images taken from the dataset with their corresponding 2D and 3D skeleton representation. The cameras’ intrinsics were obtained by recording a chessboard pattern and using standard calibration methods implemented in OpenCV [46].

Exercise	Instruction Label	Subject 1	Subject 2	Subject 3	Subject 4	Total (per instruction)	Total (per action)
Squats	Correct	10	10	11	10	41	132
	Feet too wide	5	8	5	5	23	
	Knees inward	6	7	5	5	23	
	Not low enough	5	7	5	4	21	
	Front bent	5	6	6	7	24	
Lunges	Correct	12	11	11	12	46	127
	Not low enough	10	10	10	10	40	
	Knee passes toe	10	10	11	10	41	
Planks	Correct	7	8	11	7	33	103
	Arched back	5	5	11	9	30	
	Hunch back	10	10	11	9	40	

Table 1. The EC3D dataset with the number of sequences per instruction of each subject, the total number of sequences per instruction and the total number of sequences per action. We reserve Subjects 1, 2, and 3 for training and 4 for testing.

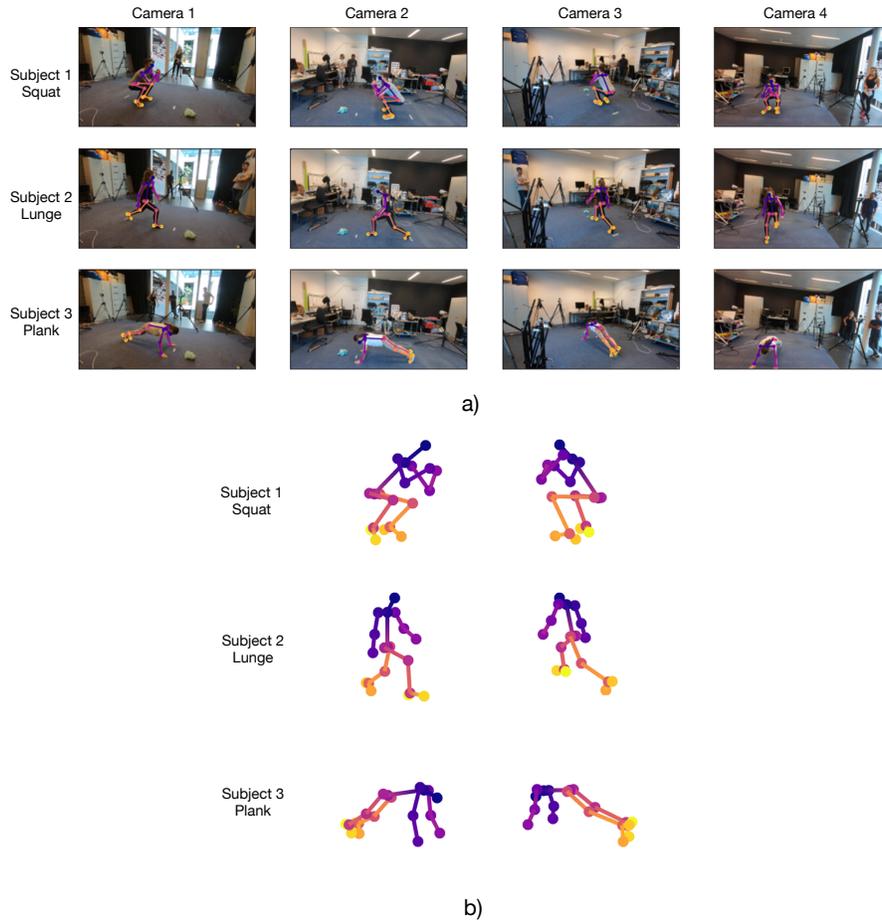


Fig. 4. Examples images from the EC3D dataset, depicting images from the SQUAT, lunge, and plank classes with their corresponding 3D pose visualizations. a) Images for each exercise type from the dataset from each camera viewpoint, with the 2D poses overlaid. b) The corresponding 3D poses, visualized from two different viewpoints.

We annotated the 3D poses in an automated manner, whereas the action and instruction labels were annotated manually. Specifically, the 3D pose annotation was performed as follows: First, the 2D joint positions were extracted from the images captured by each camera using OpenPose [47], an off-the-shelf 2D pose estimation network. We then used bundle adjustment to determine the cameras' extrinsics. For the bundle adjustment algorithm to converge quickly and successfully, additional annotations were made on static landmarks in 5 frames. Since the cameras were static during recording, for each camera, we averaged the extrinsics optimized for each of these frames. After-

wards, these values were kept constant, and we triangulated the 2D poses to compute the 3D poses.

During the triangulation process, we detected whenever any joint had a high reprojection error to catch mistakes in the 2D pose estimates. Such 2D pose annotations were discarded to prevent mistakes in the 3D pose optimization. The obtained 3D pose values were afterwards smoothed using a Hamming filter to avoid jittery motion. Finally, we manually went through the extracted 3D pose sequences in order to ensure that there are no mistakes and that they are consistent with the desired motion.

To make the resulting 3D poses uniform, we further normalized, centred and rotated them. As the different heights and body sizes of the different subjects cause differences in skeletal lengths, a random benchmark was selected to normalize the skeletal lengths while maintaining the connections between joints. Furthermore, we centered all poses on their hip joint and rotated them so that the spine was perpendicular to the ground and all movements performed in the same direction.

5 Evaluation

5.1 Dataset and Metrics

We use the EC3D dataset to evaluate our model performance both quantitatively and qualitatively. We use subjects 1, 2, and 3 for training and subject 4 for evaluation.

We use top-1 classification accuracy to evaluate the results of the instruction classification task, as used by other action classification works [10,11]. For the motion correction task, we make use of the action classifier branch: If the corrected motion is classified as “correct” by our classification branch, we count the correction as successful. We report the percentage of successfully corrected motions as the evaluation metric for this task.

5.2 Quantitative Results

We achieve an average mistake recognition accuracy of 90.9% when classifying sequences in EC3D, as shown by the detailed results for each specific exercise instruction in Table 2. In the same table, we also show that 94.2% of the corrected results are classified as “correct” by our classification model. The high classification accuracy and correction success show that our framework is indeed capable of analyzing physical exercises and giving useful feedback.

As no existing works have proposed detailed correction strategies, we compare our framework to a simple correction baseline consisting of retrieving the closest “correct” sequence from the training data. The closest sequence is determined as the sequence with the lowest DTW loss value to the input sequence. In Table 3, we provide the DTW values between the incorrectly performed input and the corrected output. The DTW loss acts as an evaluation of the accuracy of joint positions, as it is an L2 loss on the time aligned sequences. For this metric, the lower, the better, i.e., the output motion should be as close as possible to the original one while being corrected as necessary. Our framework yields a high success rate of correction together with a lower DTW loss

Exercise	Mistake Label	Classification Accuracy (%)	Correction Success (%)
Squats	Correct	90.0	100
	Feet too wide	100	100
	Knees inward	100	100
	Not low enough	100	100
	Front bent	57.1	85.7
Lunges	Correct	66.7	100
	Not low enough	100	60.0
	Knee passes toe	100	90.0
Planks	Correct	85.7	100
	Arched back	100	100
	Hunch back	100	100
Average		90.9	94.2

Table 2. Results of our classification and correction branches on the EC3D dataset. We achieve 90.9% recognition accuracy on average and successfully correct 94.2% of the mistakes.

than the baseline, thus supporting our claims. Note that we do not evaluate the baseline’s correction success percentage because it retrieves the same sequences that were used to train the network, to which the classification branch might have already overfit.

5.3 Qualitative Results

In Figure 5, we provide qualitative results corresponding to all the incorrect motion examples from each action category. Note that the incorrect motions are successfully corrected, yet still close to the original sequence. This makes it possible for the user to easily recognize their own motion and mistakes.

5.4 Ablation Studies

We have tried various versions of our framework and recorded our results in Table 4. In this section, we present the different experiments, also depicted in Figure 6, and the discussions around these experiments.

Separated models. We first analyze the results of separated classification networks. According to Table 4, our separated classification branch architecture is denoted as “separated classification.” We have also evaluated a simpler, fully GCN based separated action classifier branch, denoted as “separated classification (simple)”. We show that the results of the classification branch degrade slightly when separated from the correction branch. This indicates that the classification branch also sees a minor benefit from being part of a combined model. The simpler classification network performs worse than our architecture inspired by [11], showing that the pooling module improves the classification accuracy.

Afterwards, we analyze the results of a separated correction network, denoted as “separated correction”. Here the difference is quite profound; we see that separating

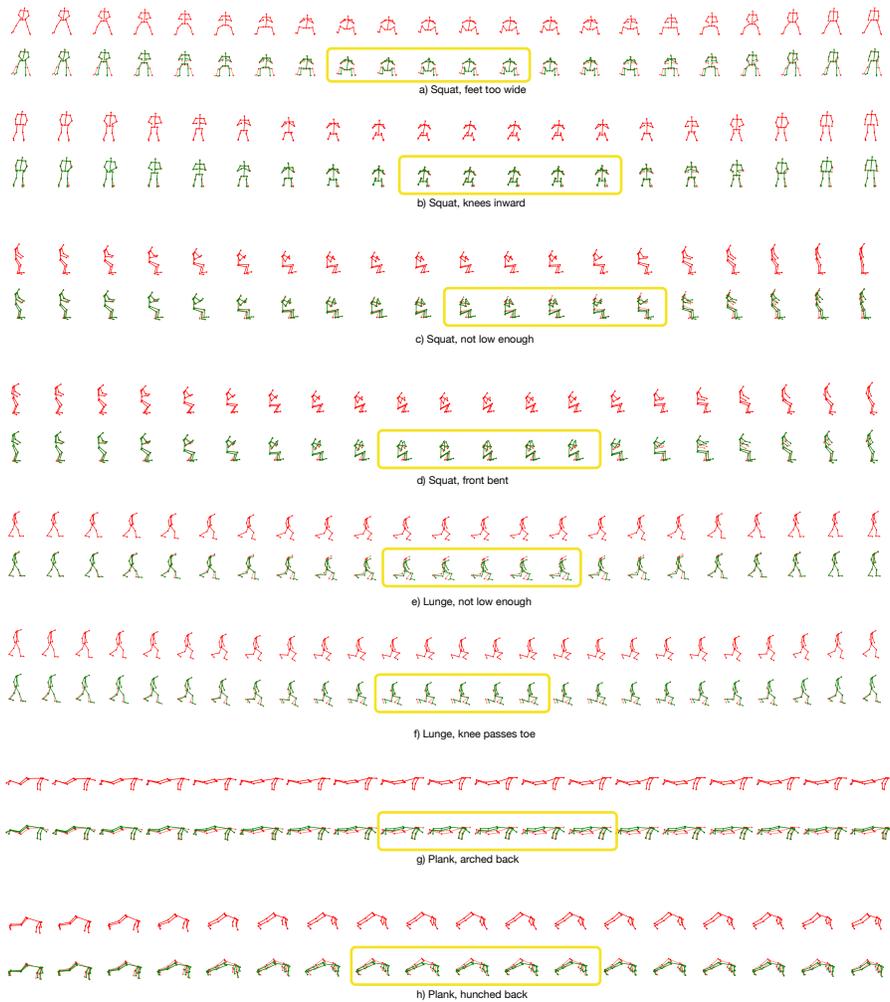


Fig. 5. Qualitative results from our framework depicting incorrect input motions and corrected output motions from categories a-d) squats, e-f) lunges, g-h) planks. We present the incorrect input sequences (red) in the top row. The corrected sequences (green) overlaid on top of the incorrect input sequences (red) are presented in the bottom row. The most significant corrections are highlighted with a yellow bounding box. We find that our proposals are successful in correcting the incorrect sequences. This figure is best viewed in color and zoomed in on a screen.

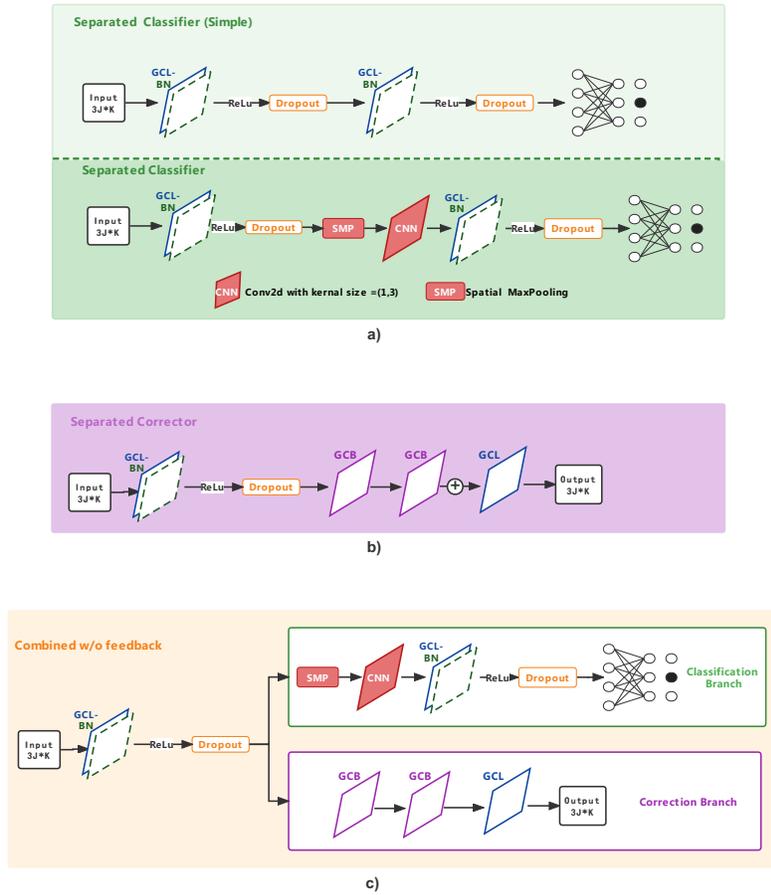


Fig. 6. Ablation study frameworks. We depict the different architectures we evaluated for the ablation studies: a) Separated classifier (simple) and separated classifier. b) Separated corrector. c) Combined model without feedback. This model does not include a “Feedback Module”, the classification branch’s results are not explicitly fed to the correction branch.

Exercise	Mistake Label	Retrieval Baseline	Our Framework
Squats	Correct	1.28	0.56
	Feet too wide	4.23	1.46
	Knees inward	1.61	0.66
	Not low enough	1.83	0.61
	Front bent	4.74	2.53
Lunges	Correct	1.94	1.82
	Not low enough	1.86	1.31
	Knee passes toe	2.27	1.48
Planks	Correct	2.41	1.79
	Arched back	12.20	1.53
	Hunch back	4.10	1.09
Average		3.49	1.35

Table 3. DTW results of the correction branch. We compare our framework to a simple baseline retrieving the best matching “correct” sequence from the training dataset depending on the classification label. We report the DTW loss between the input and the output sequences (lower is better). Our framework successfully corrects the subject’s mistakes, while not changing the input so drastically that the subject would not be able to recognize their own performance.

the correction model from the classification model degrades correction success significantly. We note that 50 epochs was not enough for the separated corrector framework to converge, therefore we trained it for a total of 150 epochs.

Combined models. We train our framework without the feedback module (“combined w/o feedback”), and without the smoothness loss (“combined w/o smoothness”). We find that these perform worse than our model with the feedback module and with the smoothness loss in terms of correction success. This shows that using the classification results as feedback as well as the smoothness loss for training is useful for more successful corrections. We notice that our framework trained without smoothness loss has higher classification accuracy, despite having a lower correction success. We believe this is due to the fact that the smoothness loss acts as a regularizer on the framework, therefore causing slight performance losses to the classification branch. However, the results of the correction success are significantly higher with the smoothness loss. We also evaluate our trained model by passing random incorrect instruction labels to the correction branch instead of the labels predicted by the classification branch (“combined with random incorrect feedback”). The correction success drops significantly, showing that the classification results are indeed very useful for the correction branch.

Further qualitative and quantitative results are presented in the supplementary material.

5.5 Limitations and Future Work

While our current results are quite impressive, there is still room for improvement in terms of performance. In particular, our framework struggles in correcting specific types of motions, such as not low enough lunges. We plan to explore different additions to our framework, such as an attention module, to better correct these types of mistakes.

		Classification Accuracy (%)	Correction Success (%)
Separated	Separated classification (simple)	88.6	-
	Separated classification	89.8	-
	Separated correction	-	83.5
Combined	Combined w/o feedback	82.3	85.3
	Combined with random incorrect feedback	90.9	87.3
	Combined w/o smoothness	93.4	87.5
	Ours	90.9	94.2

Table 4. Results of the ablation studies with several variations of our framework. We report the classification accuracy (%) and the correction success rate (%), where higher is better for both metrics. Our framework benefits greatly from combining the two tasks in a end-to-end learning fashion, from using a feedback module, and from using a pooling layer in the classification branch. The smoothness loss causes slight degradation in classification accuracy but is greatly beneficial for the correction success.

Our future work will consist of expanding the dataset to include more action sequences and more types of mistakes, performed by a larger set of subjects. We believe that this will allow us to further improve our framework and add components that address the shortcomings we will discover using such a dataset.

6 Conclusion

We have presented a 3D pose based feedback framework for physical exercises. We have designed this framework to output feedback in two branches; a classification branch to identify a potential mistake and a correction branch to output a corrected sequence. Through ablation studies, we have validated our network architectural choices and presented detailed experimental results, making a strong case for the soundness of our framework design. We have also introduced a dataset of physical exercises, on which we have achieved 90.9% classification accuracy and 94.2% correction success.

Acknowledgements. This work was supported in part by the Swiss National Science Foundation.

References

1. Chen, S., Yang, R.R.: Pose Trainer: Correcting Exercise Posture Using Pose Estimation, arXiv Preprint (2020)
2. Yang, L., Li, Y., Zeng, D., Wang, D.: Human Exercise Posture Analysis based on Pose Estimation. In: 2021 IEEE 5th Advanced Information Technology, Electronic and Automation Control Conference (IAEAC). (2021)
3. Kanase, R., Kumavat, A., Sinalkar, R., Somani, S.: Pose Estimation and Correcting Exercise Posture. Volume 40. (2021) 03031

4. Fieraru, M., Zanfir, M., Pirlea, S.C., Olaru, V., Sminchisescu, C.: AIFit: Automatic 3D Human-Interpretable Feedback Models for Fitness Training. In: Conference on Computer Vision and Pattern Recognition. (2021)
5. Dittakavi, B., Bavikadi, D., Desai, S.V., Chakraborty, S., Reddy, N., Balasubramanian, V.N., Callepalli, B., Sharma, A.: Pose Tutor: An Explainable System for Pose Correction in the Wild. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops. (2022)
6. Rangari, T., Kumar, S., Roy, P., Dogra, D., Kim, B.G.: Video Based Exercise Recognition and Correct Pose Detection. *Multimedia Tools and Applications* (2022)
7. Kocabas, M., Athanasiou, N., Black, M.J.: VIBE: Video Inference for Human Body Pose and Shape Estimation. In: Conference on Computer Vision and Pattern Recognition. (2020) 5252–5262
8. Zheng, C., Zhu, S., Mendieta, M., Yang, T., Chen, C., Ding, Z.: 3D Human Pose Estimation With Spatial and Temporal Transformers. In: International Conference on Computer Vision. (2021) 11656–11665
9. Gong, K., Li, B., Zhang, J., Wang, T., Huang, J., Mi, M., Feng, J., Wang, X.: PoseTriplet: Co-Evolving 3D Human Pose Estimation, Imitation, and Hallucination Under Self-Supervision. In: Conference on Computer Vision and Pattern Recognition. (2022)
10. Li, Y., Yuan, L., Vasconcelos, N.: Co-Occurrence Feature Learning from Skeleton Data for Action Recognition and Detection with Hierarchical Aggregation. In: International Joint Conference on Artificial Intelligence. (2018)
11. Zhang, P., Lan, C., Zeng, W., Xing, J., Xue, J., Zheng, N.: Semantics-Guided Neural Networks for Efficient Skeleton-Based Human Action Recognition. In: Conference on Computer Vision and Pattern Recognition. (2020)
12. Aksan, E., Kaufmann, M., Cao, P., Hilliges, O.: A Spatio-temporal Transformer for 3D Human Motion Prediction. In: International Conference on 3D Vision (3DV). (2021)
13. Mao, W., Liu, M., Salzmänn, M.: History Repeats Itself: Human Motion Prediction via Motion Attention. In: European Conference on Computer Vision. (2020)
14. Kiciroglu, S., Wang, W., Salzmänn, M., Fua, P.: Long Term Motion Prediction Using Key-poses. In: International Conference on 3D Vision (3DV). (2022)
15. Fragkiadaki, K., Levine, S., Felsen, P., Malik, J.: Recurrent Network Models for Human Dynamics. In: International Conference on Computer Vision. (2015)
16. Jain, A., Zamir, A., and Saxena, S.S.A.: Structural-Rnn: Deep Learning on Spatio-Temporal Graphs. In: Conference on Computer Vision and Pattern Recognition. (2016)
17. Ghosh, P., Song, J., Aksan, E., Hilliges, O.: Learning Human Motion Models for Long-Term Predictions. In: International Conference on 3D Vision. (2017)
18. Martinez, J., Black, M., Romero, J.: On Human Motion Prediction Using Recurrent Neural Networks. In: Conference on Computer Vision and Pattern Recognition. (2017)
19. Wang, B., Adeli, E., Chiu, H.K., Huang, D.A., Niebles, J.C.: Imitation Learning for Human Pose Prediction. In: International Conference on Computer Vision. (2019) 7123–7132
20. Barsoum, E., Kender, J., Liu, Z.: HP-GAN: Probabilistic 3D Human Motion Prediction via GAN. In: Conference on Computer Vision and Pattern Recognition. (2017)
21. Bütepage, J., Black, M., Kragic, D., Kjellstrom, H.: Deep Representation Learning for Human Motion Prediction and Classification. In: Conference on Computer Vision and Pattern Recognition. (2017)
22. Bütepage, J., Kjellström, H., Kragic, D.: Anticipating Many Futures: Online Human Motion Prediction and Generation for Human-Robot Interaction. In: International Conference on Robotics and Automation. (2018)
23. Bütepage, J., Kjellström, H., Kragic, D.: Predicting the What and How - A Probabilistic Semi-Supervised Approach to Multi-Task Human Activity Modeling. In: Conference on Computer Vision and Pattern Recognition. (2019) 2923–2926

24. Aliakbarian, S., Saleh, F.S., Salzmann, M., Petersson, L., Gould, S.: A Stochastic Conditioning Scheme for Diverse Human Motion Prediction. In: Conference on Computer Vision and Pattern Recognition. (2020)
25. Mao, W., Liu, M., Salzmann, M., Li, H.: Learning Trajectory Dependencies for Human Motion Prediction. In: International Conference on Computer Vision. (2019)
26. Lebailly, T., Kiciroglu, S., Salzmann, M., Fua, P., Wang, W.: Motion Prediction Using Temporal Inception Module. In: Asian Conference on Computer Vision. (2020)
27. Katircioglu, I., Georgantas, C., Salzmann, M., Fua, P.: Dyadic Human Motion Prediction, arXiv Preprint (2022)
28. Guo, W., Bie, X., Alameda-Pineda, X., Moreno-Noguer, F.: Multi-Person Extreme Motion Prediction. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). (2022)
29. Du, Y., Wang, W., Wang, L.: Hierarchical Recurrent Neural Network for Skeleton Based Action Recognition. In: Conference on Computer Vision and Pattern Recognition. (2015)
30. Shahroudy, A., Liu, J., Ng, T., Wang, G.: NTU RGB+D: A Large Scale Dataset for 3D Human Activity Analysis. In: Conference on Computer Vision and Pattern Recognition. (2016) 1010–1019
31. Liu, J., a. Shahroudy, Xu, D., Wang, G.: Spatio-Temporal LSTM with Trust Gates for 3D Human Action Recognition. Volume 9907. (2016)
32. Song, S., Lan, C., Xing, J., Zeng, W., Liu, J.: An End-to-End Spatio-Temporal Attention Model for Human Action Recognition from Skeleton Data. In: AAAI Conference on Artificial Intelligence. (2017)
33. Liu, J., Wang, G., Hu, P., Duan, L., Kot, A.: Global Context-Aware Attention LSTM Networks for 3D Action Recognition. In: Conference on Computer Vision and Pattern Recognition. (2017)
34. Zhang, P., Lan, C., Xing, J., Zeng, W., Xue, J., Zheng, N.: View Adaptive Neural Networks for High Performance Skeleton-Based Human Action Recognition. (2019)
35. Tang, Y., Tian, Y., Lu, J., Li, P., Zhou, J.: Deep Progressive Reinforcement Learning for Skeleton-Based Action Recognition. In: Conference on Computer Vision and Pattern Recognition. (2018)
36. Li, M., Chen, S., Chen, X., Zhang, Y., Wang, Y., Tian, Q.: Actional-Structural Graph Convolutional Networks for Skeleton-Based Action Recognition. In: Conference on Computer Vision and Pattern Recognition. (2019)
37. Zell, P., Wandt, B., Rosenhahn, B.: Joint 3D Human Motion Capture and Physical Analysis from Monocular Videos. In: 2017 IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW). (2017)
38. Sakoe, H., Chiba, S.: Dynamic Programming Algorithm Optimization for Spoken Word Recognition. In: IEEE Trans. on Acoustics, Speech, and Sig. Proc. (1978)
39. Cuturi, M., Blondel, M.: Soft-DTW: a Differentiable Loss Function for Time-Series. In: International Conference on Machine Learning. (2017)
40. Bengio, S., Vinyals, O., Jaitly, N., Shazeer, N.: Scheduled Sampling for Sequence Prediction with Recurrent Neural Networks. In: Advances in Neural Information Processing Systems. (2015)
41. Kingma, D.P., Ba, J.: Adam: A Method for Stochastic Optimisation. In: International Conference on Learning Representations. (2015)
42. Verma, M., Kumawat, S., Nakashima, Y., Raman, S.: Yoga-82: A New Dataset For Fine-Grained Classification of Human Poses. In: The IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops. (2020)
43. Shao, D., Zhao, Y., Dai, B., Lin, D.: FineGym: A Hierarchical Video Dataset for Fine-grained Action Understanding. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). (2020)

44. Shenglan, L., Xiang, L., Gao, H., Hong, Q., Lianyu, H., Dong, J., Aibin, Z., Yang, L., Ge, G.: FSD-10: A Fine-Grained Classification Dataset For Figure Skating. Volume 413. (2020) 360–367
45. Yingwei, L., Yi, L., Nuno, V.: RESOUND: Towards Action Recognition without Representation Bias. In: European Conference on Computer Vision (ECCV). (2018)
46. : (Open Source Computer Vision Library) <http://opencv.org>.
47. Cao, Z., Simon, T., Wei, S., Sheikh, Y.: Realtime Multi-Person 2D Pose Estimation Using Part Affinity Fields. In: Conference on Computer Vision and Pattern Recognition. (2017) 1302–1310