

# Virtual Reality for Action Evaluation

Mario De Lucas Garcia\*  
mdelucasgarcia@hawk.iit.edu  
Illinois Institute of Technology  
Chicago, Illinois, USA

Mark Roman Miller†  
mmiller30@iit.edu  
Illinois Institute of Technology  
Chicago, Illinois, USA

## ABSTRACT

Physical rehabilitation plays a crucial role in restoring functional abilities, but traditional approaches often face challenges in terms of cost, accessibility, and personalized monitoring. Asynchronous physical rehabilitation has gained traction as a cost-effective and convenient alternative, but it lacks real-time monitoring and assessment capabilities. This study investigates the feasibility of using low-cost Virtual Reality (VR) devices for action evaluation in rehabilitation exercises. We leverage state-of-the-art deep learning models and evaluate their performance on three data streams (head and hands) derived from existing rehabilitation datasets that approximate VR headset and hand data. Our results demonstrate that VR tracking data can be effectively utilized for action evaluation, paving the way for more accessible and affordable remote monitoring solutions in physical therapy. By leveraging artificial intelligence techniques and consumer-grade virtual reality technology, this study proposes an approach that could potentially address some of the challenges in asynchronous rehabilitation, such as the need for expensive motion capture systems or in-person sessions.

## CCS CONCEPTS

• **Computing methodologies** → *Boosting; Motion capture; Neural networks.*

## KEYWORDS

Motion data, Kinect, Virtual Reality, Action evaluation, Machine Learning, XGBoost, Feature Engineering, Deep Learning

## 1 INTRODUCTION

In the United States, there are over 250,000 physical therapists, and this number is expected to grow by 15% in the coming years to meet the increasing demands of society<sup>1</sup>.

In-person physical rehabilitation with a professional has long been the standard for therapy intervention and is widely recognized as one of the most effective ways to rehabilitate patients. However, this traditional process involves several challenges, including the need for synchronous sessions, significant time commitments, high costs, and the inconvenience of clinic travel, which may be particularly difficult for some patients [13].

To address these issues, healthcare providers are increasingly adopting asynchronous rehabilitation modalities. In this model, patients perform a prescribed set of exercises at home, following guidelines provided by professionals. This approach has proven effective in musculoskeletal practices and has led to high patient satisfaction [8]. However, monitoring these exercises and providing real-time assessments remain challenging tasks for researchers and companies.

Researchers have employed statistical algorithms like Dynamic Time Warping (DTW) and distance measures like Euclidean and Jacard. Some have explored machine learning techniques like Support Vector Machines (SVM) and Gaussian Mixture Models (GMMs) to classify correct and incorrect movements. More recently, deep learning architectures such as Convolutional Neural Networks (CNNs) and Recurrent Neural Networks (RNNs) have been used to model joint dependencies and assist in movement scoring [16].

Additionally, Virtual Reality (VR) technology, which is widely used in vocational training, entertainment, and medical care, offers promising applications for rehabilitation. Stand-alone head-mounted display devices are becoming more affordable, user-friendly, and accessible to many users simultaneously, providing a low-cost method for motion tracking [9].

This paper aims to combine artificial intelligence techniques with VR technology to evaluate the effectiveness of these devices as cost-effective motion-tracking systems in asynchronous rehabilitation processes. Hence, we proposed a simple baseline model that outputs the average score from the training set and an eXtreme Gradient Boosting (XGBoost) machine learning model that, given a selected set of features, predicts a score. We will provide an in-depth analysis of these models in Section 3.

We summarize the contributions of this paper as follows:

- **Automated Preprocessing Pipeline:** We proposed a comprehensive preprocessing pipeline that automatically splits the data into exercise repetitions, facilitating the training of various models.
- **Model Comparison Framework:** We conducted a thorough comparison of state-of-the-art models using standardized training setups. To the best of our knowledge, this is the first work aimed at democratizing model comparison in this context.
- **Novel XGBoost Model:** We implemented a novel eXtreme Gradient Boosting (XGBoost) machine learning model that performs comparably to the current state-of-the-art models when using our training setup data and parameters.

The rest of this work is organized as follows: The section 2 provides an overview of related research. Section 3 details the data and models used in our experiments, including the novel eXtreme Gradient Boosting (XGBoost) machine learning model. Section 4 presents the results, while Sections 5 and 6 summarize the key findings, draw conclusions, and outline potential future research directions.

<sup>1</sup><https://www.bls.gov/ooh/healthcare/physical-therapists.htm>

## 2 RELATED WORK

### 2.1 Motion Capture Systems

The effectiveness of physical rehabilitation programs depends on patients' adherence and the accurate execution of prescribed exercises. To ensure systematic monitoring of exercise performance, prior research has explored the use of advanced technologies such as video-based systems, wearable devices, and Inertial Measurement Units (IMUs) for measuring and tracking motion [4].

Vision-based strategies for motion tracking can be categorized into marker-based and marker-less techniques. Marker-based techniques involve placing tags or markers on specific body parts, which allows for precise measurement of their positions in space. This approach can be useful for guiding patients through rehabilitation exercises, as it provides accurate tracking of their movements. However, marker-based systems and body-worn sensors, while accurate for motion capture, can be intrusive and disruptive to a patient's daily activities. The need to properly place and wear the sensors or markers can make these systems impractical for successful implementation of home-based rehabilitation programs, where patients are expected to perform exercises independently without direct supervision [5, 7].

On the other hand, marker-less techniques rely solely on computer vision algorithms to track and analyze body movements without the need for any physical markers or sensors attached to the body (Microsoft Kinect). This approach can be less intrusive and more convenient for patients, but may potentially sacrifice some accuracy compared to marker-based systems [5, 18].

### 2.2 Applied Virtual Reality

Virtual Reality (VR) technology has gained significant attention in recent years, with numerous research studies exploring its potential applications across various domains involving physical movement and motion activities. These applications range from learning complex skills, such as couple dancing, to enhancing physiotherapy and rehabilitation processes. By leveraging the immersive and interactive nature of VR, researchers and practitioners have been able to develop innovative approaches that show promising results in terms of improved learning outcomes, patient engagement, and overall effectiveness.

One notable example is the use of VR in learning couple dance. By creating virtual environments that simulate real-world dance settings, VR systems can provide learners with immersive experiences that help them practice and refine their dance moves. These systems can offer real-time feedback, personalized guidance, and the ability to practice at one's own pace, making the learning process more efficient and enjoyable [17].

Similarly, VR technology has been applied to the field of physiotherapy. By incorporating VR into rehabilitation programs, therapists can create engaging and interactive exercises that motivate patients to adhere to their treatment plans. VR-based physiotherapy can also provide real-time monitoring and analysis of patient movements, enabling therapists to track progress, adjust treatment plans, and provide targeted feedback [3, 10, 15].

### 2.3 Action evaluation

In prior studies, movement evaluation is typically conducted by comparing a patient's performance to that of healthy participants. Early researchers employed machine learning models to classify individual movements as either correct or incorrect. They experimented with well-known methods such as the Adaboost classifier, k-nearest neighbours, and Bayesian classifier, which produce binary class outputs. However, these approaches lack the ability to detect varying levels of motion quality or to identify gradual changes in patient performance throughout the rehabilitation program [1, 19].

A line of research has utilized probabilistic approaches for modeling and evaluating rehabilitation actions. Studies employing hidden Markov models (HMMs) [2] and Gaussian Mixture Models (GMMs) [12] generally conduct quality assessments based on the likelihood that individual movement sequences are drawn from a trained model.

In recent years, researchers have been exploring the application of deep learning techniques to model complex dependencies between body joints for the purpose of assessing physical exercises. One notable study by Liao et al. [11] introduced the Spatio-Temporal Neural Network (STNN), a deep learning architecture designed specifically for evaluating the quality of exercise performance. The STNN combines several key components, including temporal pyramids, multi-branch convolution, and recurrent layers, to effectively capture and analyze the spatial and temporal patterns in human motion data. The input to the STNN is typically skeleton data, which represents the positions and movements of 22 to 39 body joints, depending on the sensors used to capture the motion.

Another important contribution in this field comes from Deb et al. [6], who employed Spatio-Temporal Graph Convolution Networks (SGNN) for the task of physical exercise assessment. SGNNs are a type of deep learning model that can operate directly on graph-structured data, such as the skeletal representation of human body joints. By leveraging the inherent graph structure of the body, SGNNs can effectively capture the dependencies and relationships between different joints, enabling a more accurate and robust assessment of exercise performance.

In a different approach, Mottagui et al. [14] proposed a hybrid framework that combines deep learning algorithms with probabilistic models. By integrating these two complementary paradigms into a unified training process, their method aims to leverage the strengths of both deep learning and probabilistic modelling. Deep learning excels at automatically learning complex patterns and representations from raw data, while probabilistic models provide a principled way to handle uncertainty and incorporate prior knowledge. By combining these approaches, Mottagui et al. [14] seek to develop more robust and interpretable models for exercise assessment.

These studies highlight the growing interest in applying advanced machine learning techniques, particularly deep learning, to the problem of modelling and assessing complex human movements in the context of physical exercises.

### 3 METHODS

#### 3.1 Problem formulation

For each subject  $s \in S$  the set of all subjects listed in the dataset, we denote her joint positions and orientations for  $i$ th exercise (i.e. each exercise file may have more than one repetition) as  $P_{s,i}$  and  $O_{s,i}$ , respectively. Each exercise file has a ground-truth score annotation  $y_{s,i} \in [0, 1]$  that represents the quality of the performed exercise based on professional criteria. The higher the score, the better the patient moves. We train each model  $\text{Model}_i$  for each  $i$  distinct exercise with its  $\theta_i$  parameters.

Each model we have tested predicts a continuous score  $\hat{y}_j$  close to the ground-truth assessment score,  $y_j^*$  for a given joint position or orientation.

$$\hat{y} = \text{Model}_i(P_{S,N}; \theta_i), \text{ s.t. } \hat{y}_j \approx y_j^* \quad (1)$$

or using joint orientations as inputs, as follows

$$\hat{y} = \text{Model}_i(O_{S,N}; \theta_i), \text{ s.t. } \hat{y}_j \approx y_j^* \quad (2)$$

#### 3.2 Action Evaluation Public Datasets

The requirements we established for useful datasets were: (1) it must be publicly available; (2) it must have a clearly defined exercise quality measure; and (3) it must contain skeleton data, specifically including head and hands data. Consequently, we thoroughly researched public options for evaluated action data. Our findings indicate that KIMORE is the most reliable data source, as it meets all our requirements. Table 1 provides a comparison of each option.

For running the experiments, we used the KIMORE dataset, which contains data from 78 subjects: 44 in the control group (i.e., experts or non-experts who are healthy) and 34 with chronic motor disabilities. This dataset was recorded using the Kinect v2 depth sensor, capturing 25 distinct body parts, including their skeleton positions and orientations. Each subject performed five exercises:

- Lifting of the arms
- Lateral tilt of the trunk with the arms extended
- Trunk rotation
- Pelvis rotations on the transverse plane
- Deep squats


During preprocessing, as described in the original paper, we applied a 3rd order low-pass Butterworth filter to remove temporary spikes. The dataset consists of a single CSV file containing positions or orientations per person and per exercise. We developed a custom function to split these files into individual repetitions. For example, in the lifting of the arms exercise, we identified that the Y coordinate of hands effectively indicated when the patient started a repetition. Additionally, we applied threshold filtering to remove false peaks that were found. We used the timestamps where the X position was maximal to mark the start of each repetition. See Figure 1

However, a limitation of the KIMORE dataset is that it provides only one score per exercise, meaning all repetitions share the same ground truth score. The given score ranges from 0 to 50, which we scaled to a range of 0 to 1 for our analysis.

#### 3.3 Proposed Method

We have built two baseline models to compare the precision of current state-of-the-art models for the KIMORE dataset. These baseline

Splitting first exercise using y-axis right hand joint position



**Figure 1: The above image shows all the peaks it found. The below image shows the final start timestamps we used for splitting the CSV file into repetitions.**

models serve as a reference point for evaluating the performance of more advanced techniques.

The first baseline model is an aggregator that always outputs the average score of the training repetitions. Since the dataset contains a variety of different exercises, we apply this baseline model to each exercise separately.

The second baseline model is an eXtreme Gradient Boosting (XGBoosting) machine learning model that utilizes a curated list of features. Our goal is to determine whether Virtual Reality data can be effectively used to assess physical rehabilitation exercises. To this end, we calculate the following 44 features:

- Maximum, minimum, mean, and standard deviation of the distance between hands for the three coordinates.
- Maximum, minimum, mean, and standard deviation of the distance between the right/left hand and the head for the three coordinates.
- Maximum, minimum, mean, and standard deviation for each right/left hand and head joint for each coordinate.

Our proposed model demonstrates competitive performance, even outperforming state-of-the-art models like the Spatio-Temporal Neural Network (STNN) in various exercises from the KIMORE dataset (see Figure ??).

## 4 RESULTS

### 4.1 Setup

We first conducted a thorough review of publicly available datasets and chose to work with the KIMORE dataset. Although UI-PRMD and IntelliRehabDS include a greater variety of exercises, they lack a clear methodology for scoring. Specifically, UI-PRMD provides a labeled subset with scores derived from distance functions, while IntelliRehabDS only offers binary labels (i.e., correct or incorrect movements). The KIMORE dataset, by contrast, provides a more

**Table 1: Publicly available datasets.**

Dataset Name	Exercise Quality Measure?	Position and Rotation Data?	Access?	Notes
KIMORE	Yes - rated by a physician.	Yes. Skeleton positions (trajectory positions) and orientations in the format of Kinect.	Yes	78 subjects
IntelliRehabDS (IRDS)	No, the dataset contains a correctness label for each gesture, indicating whether it was performed correctly or incorrectly by the subject.	Yes, the dataset contains the 3D coordinates of 25 body joints extracted by a Kinect sensor, as well as the raw depth map images for each frame.	Yes	29 subjects, out of which 15 were patients and 14 were healthy controls.
UI-PRMD	No, the paper proposes a taxonomy for performance metrics for the evaluation of therapy movements, but only a reduced subset was scored.	Yes, the paper provides both position and angle data for 39 joints of the human body, obtained from a Vicon optical tracker and a Kinect sensor.	Yes	10 people, 10 actions
3D Motion Capture Data	No	Yes, the paper provides 3D joint centre positional data, 3D joint angles, and 3D segment velocity and acceleration data of the head, trunk, upper arms, forearms, pelvis, thighs, shanks, and feet segments.	No	183 Athletes

comprehensive and methodologically sound framework for evaluating rehabilitation exercises, making it the most suitable choice for our study.

Next, we developed a preprocessing pipeline to segment each recording into individual exercise repetitions. This was accomplished by detecting peaks in the position data of key joints. We then applied a cleaning filter to eliminate false peaks. To ensure generality, we standardized each repetition to a fixed length of 104 timesteps. Additionally, we prepared the data so that any model could be trained on either the whole body joints or just the head and hand joints.

Then, we evaluate the performance of each model using three metrics where each model was trained for each exercise separately since the inherent features of each movement were disparate: (1) Root Mean Square Error (RMSE), (2) Mean Absolute Error (MAE), and (3) Mean Absolute Percentage Error (MAPE). Lower scores indicate more precise predictions across all metrics. We use the same training and testing subsets for each model to ensure a fair comparison and eliminate the influence of external factors. Additionally, we modify each model to accept either the whole body joint positions and orientations or only the head and hand data.

We maintain the training procedures described in the original papers for the Spatio-Temporal Neural Network (STNN) and Spatio-Temporal Graph Convolution Network (SGNN) models. Specifically, we train a separate STNN for each exercise, using 500 epochs, a batch size of 10, and a learning rate of 0.0001. Similarly, we train a separate GCN for each exercise, using 1000 epochs, a batch size of 10 and a learning rate of 0.0001.

## 4.2 Position

The results of training the four models with the same data are presented below. In this context, **VR** indicates that only head and hand data were used.

As expected, the Baseline model consistently performs the worst across all exercises, both for the full joint data and the VR (head and hand) data (See Figure 2). This simple averaging approach serves as a reasonable lower bound but is outperformed by more sophisticated models.

Surprisingly, according to the full joints chart, the SGNN model exhibits a worse RMSE compared to the simple baseline model and has an MAE that is equally poor. In contrast, the XGBoost model outperformed the STNN model in exercises 3 and 5, and it matched the STNN model in exercise 4, based on the RMSE metric. When evaluating the MAE metric, XGBoost generally performed worse than the STNN model, except for the last exercise.

This trend persisted when using VR data format (head and hand data only). Interestingly, both models maintained their performance levels even when limited to head and hand joint positions. Notably, the XGBoost model showed improved performance in the third exercise under these conditions. See Figure 2.

## 4.3 Orientation

As mentioned earlier, the baseline model consistently exhibits the poorest performance across all workout scenarios, both when considering the full joint data and the virtual reality (VR) data comprising head and hand orientations. Similarly, the SGNN (Spatio-Temporal Graph Neural Network) model mirrors the unsatisfactory performance of the baseline model, as illustrated in Figure 3.

In contrast, the XGBoost model closely reaches the performance of the STNN (Spatio Temporal Neural Network) model for most exercises. This tendency continues even when employing the VR

Joint position Prediction - RMSE and MAE

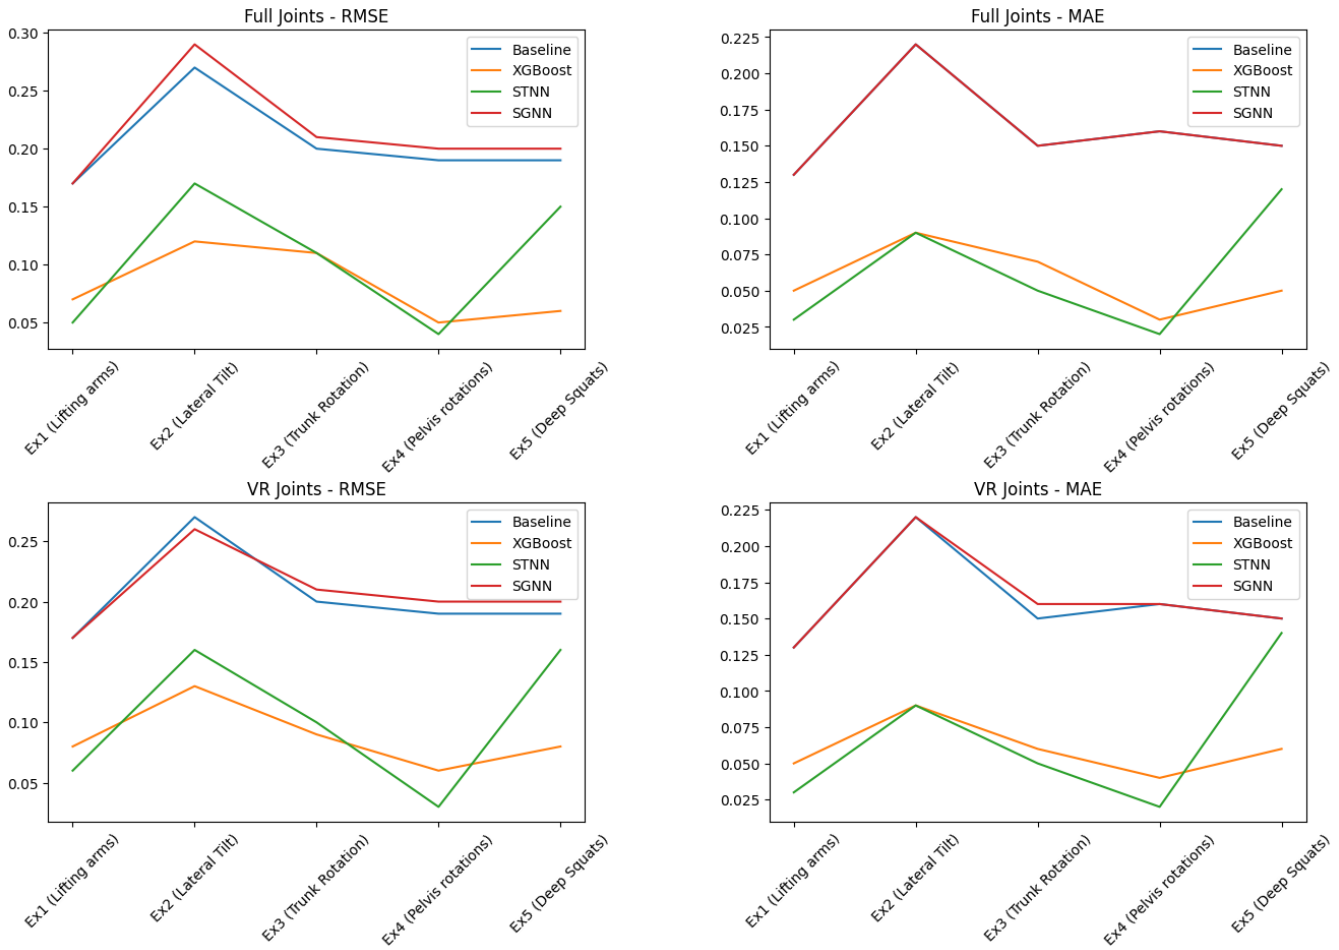


Figure 2: Results of training the four models with the same data. VR means only head and hand data.

data format. Nonetheless, both the XGBoost and STNN models experience a significant decline in performance when working with limited joint orientation data, as evidenced by Figure 3.

#### 4.4 Position VS Orientation

To optimize performance, it is essential to determine which data input format yields superior results. Focusing on the full joint data analysis, as depicted in Figure 4, we observe that the Spatio-Temporal Neural Network (STNN) model consistently outperforms across all exercise routines when leveraging orientation data as input compared to when using position data alone. This behavior aligns with the expected outcome, as the STNN architecture was specifically designed to capitalize on orientation data, considering the spatial relationships and orientations of the body joints during movement.

Interestingly, our findings also reveal that the XGBoost model, a powerful ensemble learning technique, performs significantly better when incorporating orientation data as an additional input

feature. This suggests that the inclusion of joint orientation information, in addition to positional data, provides valuable insights that enable the XGBoost model to make more accurate predictions and classifications across exercise routines.

Shifting our focus to the head and hand data analysis, as illustrated in Figure 5, we observe that the trend discovered in the full joint data analysis persists. Specifically, the Spatio-Temporal Neural Network (STNN) model continues to exhibit superior performance when leveraging orientation data as input compared to position data alone. However, it is essential to note that exercise 4, which involves pelvis rotations on the transverse plane, stands as an exception to this trend. For this particular exercise, the STNN model surprisingly achieved better results when working with position data rather than orientation data.

This discovery is intriguing because prior research studies in this domain did not explicitly explore the implications of employing either position or orientation data as input features for their respective analyses. Most previous works focused on leveraging a

Joint orientation Prediction - RMSE and MAE

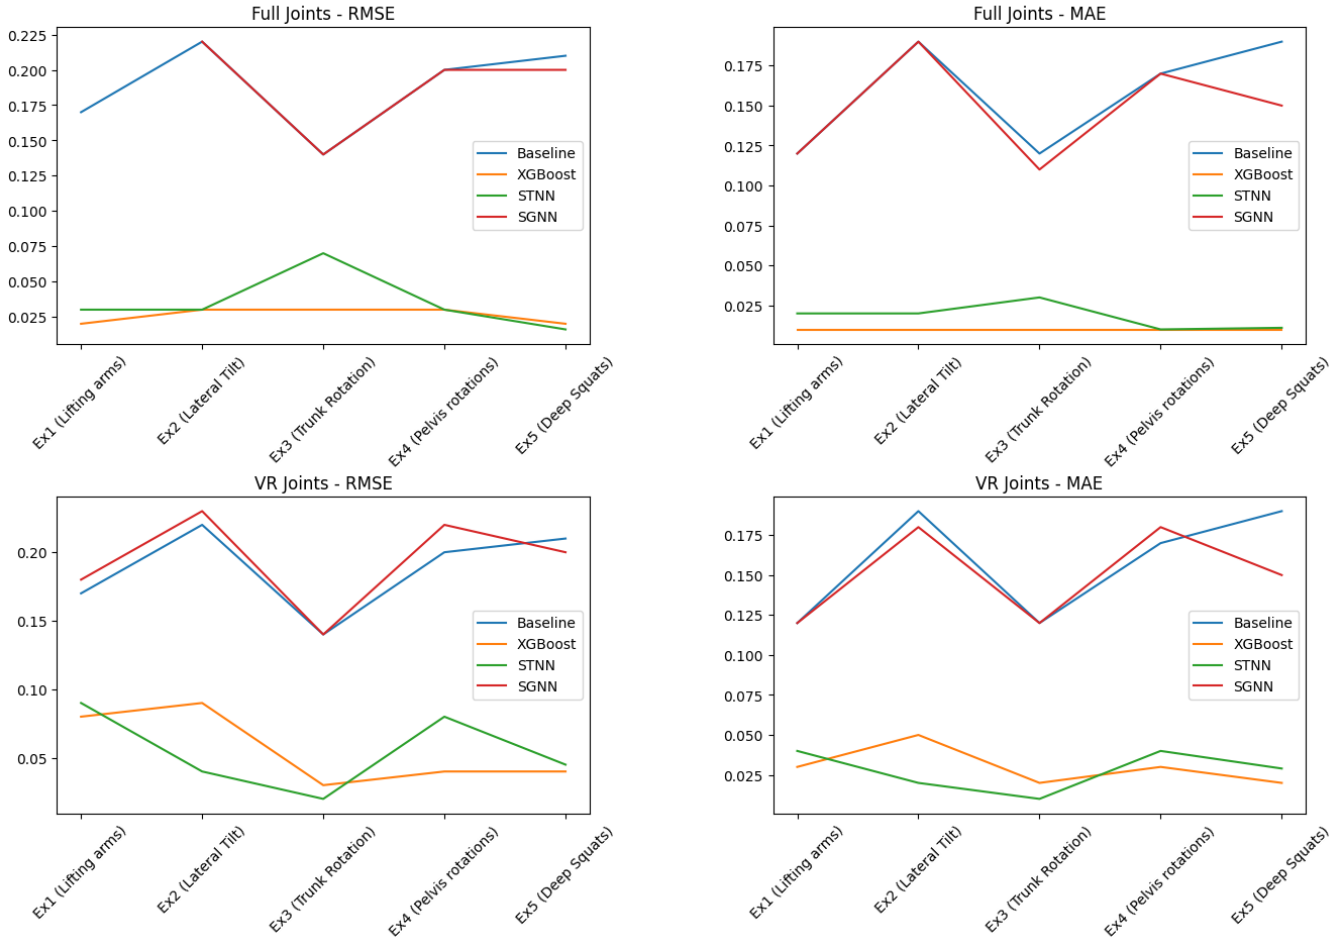


Figure 3: Results of training the four models with the same data. VR means only head and hand data.

specific data format without directly comparing the performance implications of using alternative input representations, such as position versus orientation information.

### 5 DISCUSSION

Advances in machine learning and computer vision techniques have sparked increased interest in the automated evaluation of rehabilitation exercises. However, despite significant progress, several open questions and numerous challenges remain before these asynchronous systems can be widely deployed.

To encourage adherence to rehabilitation programs, it is crucial to design real-time assessment applications that guide patients through each exercise. However, this remains a notable challenge.

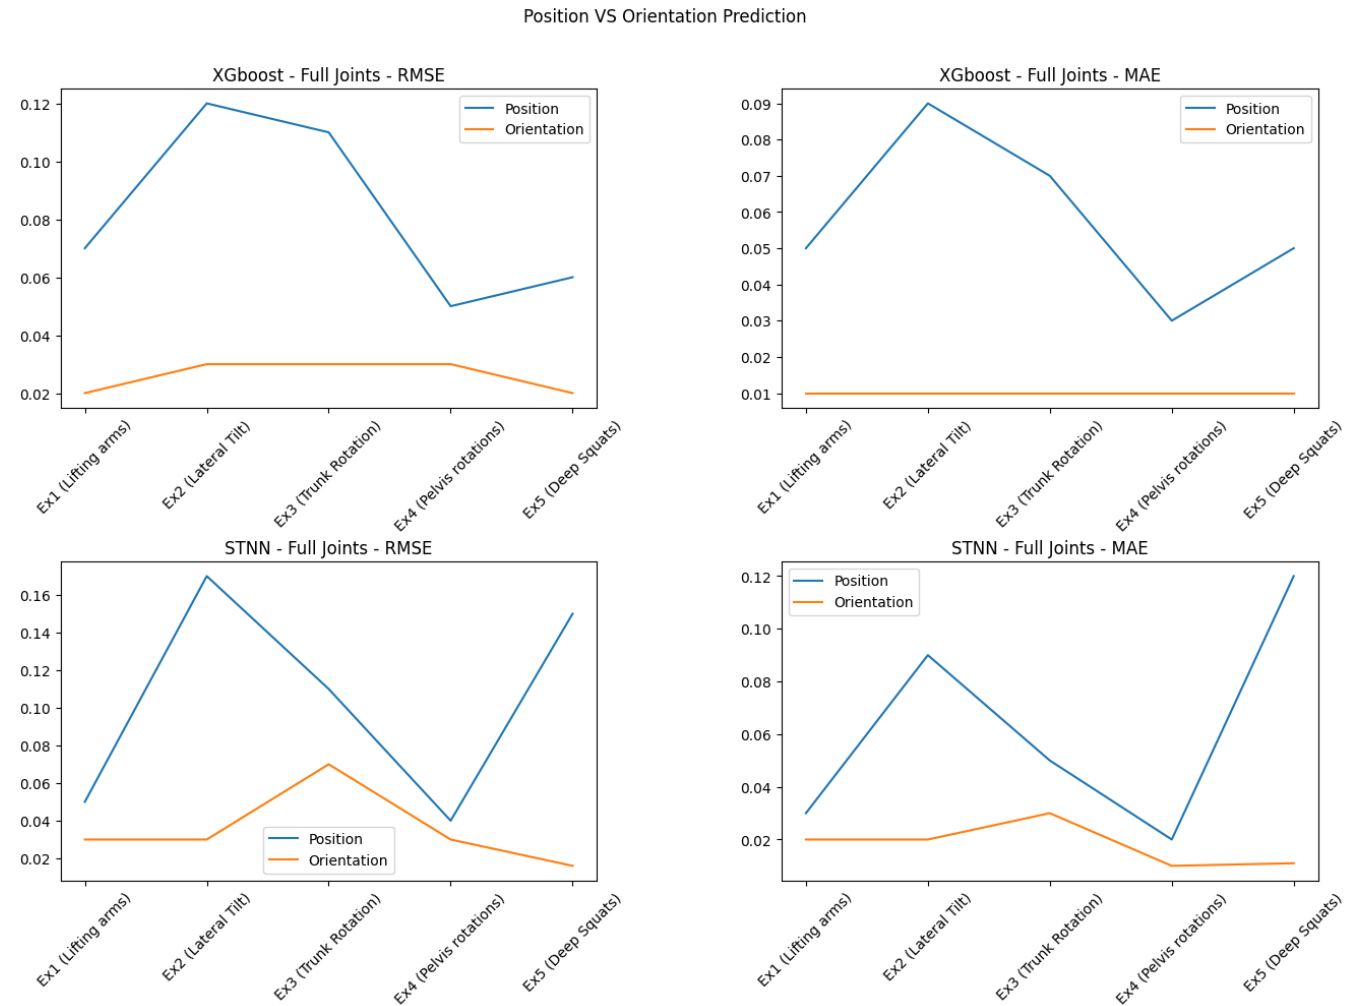
In this work, we present a novel approach for handling skeleton data using an ensemble learning method. Specifically, we developed an eXtreme Gradient Boosting (XGBoost) model that achieves performance comparable to state-of-the-art models on the KIMORE dataset. Additionally, we provide a common framework that allows

researchers to easily integrate and evaluate new models against existing ones using our open preprocessing and splitting pipeline.

Initially, we analyzed various public datasets relevant to this problem. After a thorough comparison, we decided to focus exclusively on the KIMORE dataset. We developed an automatic approach to segment the data into exercise repetitions by detecting the peaks in certain joint movements.

Our results indicate that the Spatio-Temporal Neural Network (STNN) generally outperforms other models, even surpassing the results reported by its original authors. However, as shown in Figures 4 and 5, our XGBoost model achieves performance comparable to the STNN across all exercises. This finding suggests that machine learning models can remain competitive even as more complex deep learning models emerge. Sometimes, a simpler yet well-tuned approach can be more effective than a more intricate solution.

Furthermore, we observed that using only head and hand data results in a slight performance drop but still remains close to state-of-the-art values. This suggests that head-mounted virtual reality



**Figure 4: Comparison between the XGBoost and STNN models using the full joint positions and orientations.**

devices could potentially be used in rehabilitation programs as a low-cost method for tracking and assessing exercises.

All the code will be available from <https://github.com/mdelucasg/VR-for-Action-Evaluation>.

## 6 FUTURE WORK

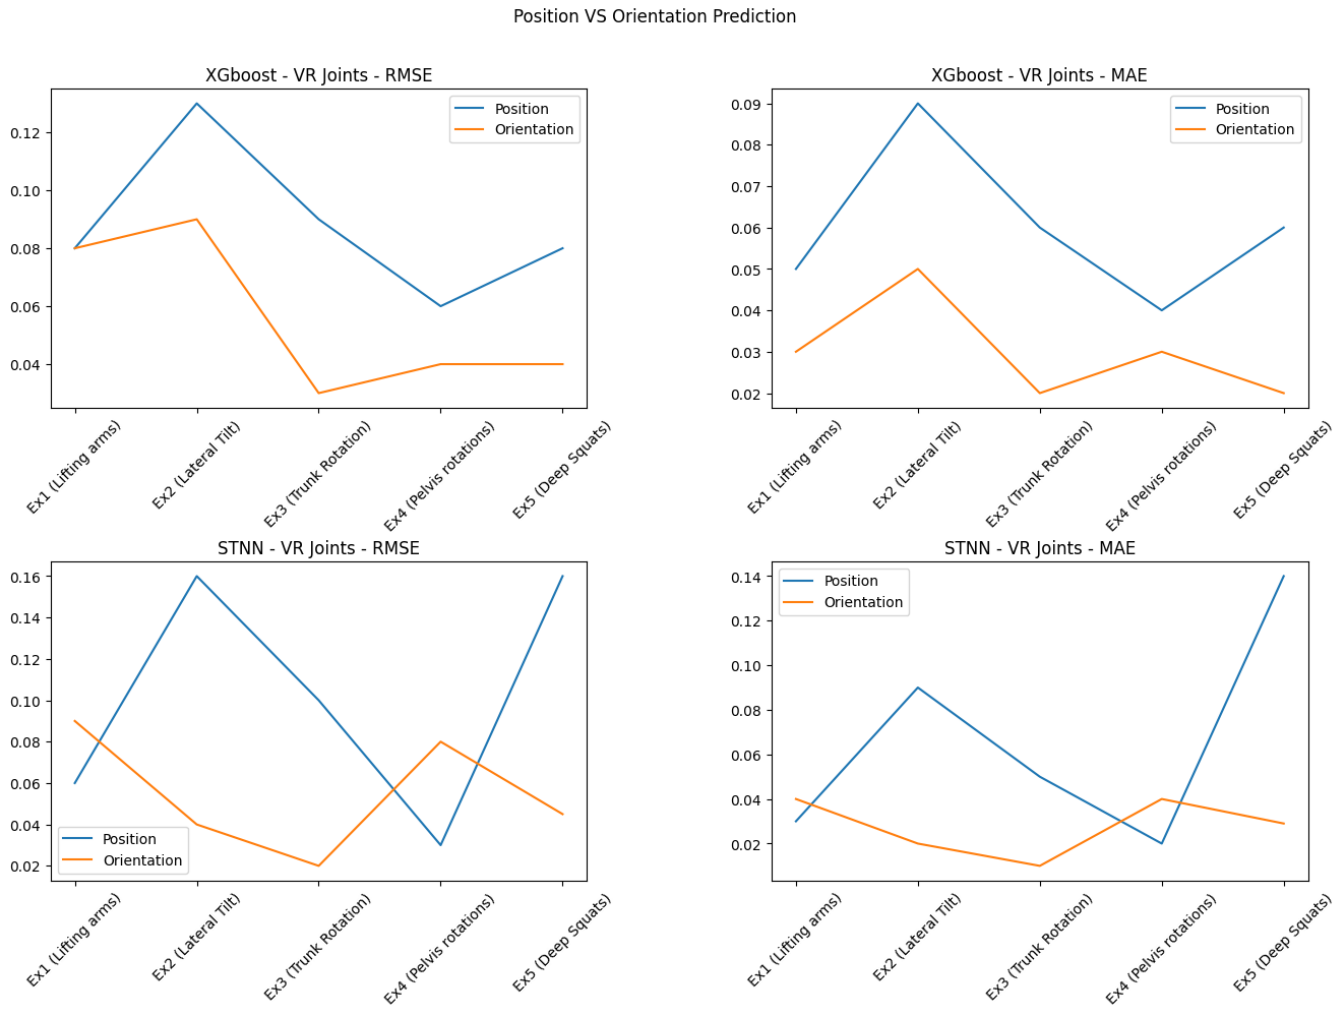
To further validate and corroborate our findings, we have planned to conduct additional data collection using head-mounted virtual reality (VR) devices. While our current analysis has demonstrated that, in theory, even with a loss in performance, the use of limited joint orientation data remains relatively affordable and feasible, it is crucial to verify these observations with real-world data gathered from actual VR devices. This step is essential to fully verify our claims and ensure the robustness of our conclusions.

Furthermore, we intend to revisit the implementation of the Spatio-Temporal Graph Neural Network (STGNN) model. Despite its promising theoretical foundations, the performance we have observed from this model in our current analysis falls short of

state-of-the-art (SOTA) levels. By re-implementing and training the SGNN architecture, we aim to unlock its full potential and push its performance closer to the cutting edge in this domain.

## REFERENCES

- [1] Ilkhan Ar and Yusuf Sinan Akgul. 2014. A computerized recognition system for the home-based physiotherapy exercises using an RGBD camera. *IEEE Transactions on Neural Systems and Rehabilitation Engineering* 22, 6 (2014), 1160–1171.
- [2] Marianna Capecci, Maria Gabriella Ceravolo, Francesco Ferracuti, Sabrina Iarlori, Ville Kyrki, Andrea Monteriu, Luca Romeo, and Federica Verdini. 2018. A Hidden Semi-Markov Model based approach for rehabilitation exercise assessment. *Journal of biomedical informatics* 78 (2018), 1–11.
- [3] Andrea Chirico, Tania Giovannetti, Pietro Neroni, Stephanie Simone, Luigi Gallo, Federica Galli, Francesco Giancamilli, Marco Predazzi, Fabio Lucidi, and Antonio Giordano. 2020. Virtual reality for the assessment of everyday cognitive functions in older adults: An evaluation of the virtual reality action test and two interaction devices in a 91-year-old woman. *Frontiers in Psychology* 11 (2020), 493493.
- [4] Emily E Cust, Alice J Sweeting, Kevin Ball, and Sam Robertson. 2019. Machine and deep learning for sport-specific movement recognition: A systematic review of model development and performance. *Journal of sports sciences* 37, 5 (2019), 568–600.
- [5] Kishor Das, Thiago de Paula Oliveira, and John Newell. 2023. Comparison of markerless and marker-based motion capture systems using 95% functional limits



**Figure 5: Comparison between the XGBoost and STNN models using the head and hand joint positions and orientations.**

of agreement in a linear mixed-effects modelling framework. *Scientific Reports* 13, 1 (2023), 22880.

[6] Swakshar Deb, Md Fokhrul Islam, Shafin Rahman, and Sejuti Rahman. 2022. Graph convolutional networks for assessment of physical rehabilitation exercises. *IEEE Transactions on Neural Systems and Rehabilitation Engineering* 30 (2022), 410–419.

[7] Habib M Fardoun, Pedro Gamito, Victor MR Penichet, and Daniyal M Alghazawi. 2015. *Proceedings of the 3rd 2015 Workshop on ICTs for improving Patients Rehabilitation Research Techniques*. ACM.

[8] Stacey L Grona, Brenna Bath, Angela Busch, Thomas Rotter, Catherine Trask, and Elizabeth Harrison. 2018. Use of videoconferencing for physical therapy in people with musculoskeletal conditions: a systematic review. *Journal of telemedicine and telecare* 24, 5 (2018), 341–355.

[9] Fangqiuzi He, Yong Liu, Weiwen Zhan, Qingjie Xu, and Xiaoling Chen. 2022. Manual operation evaluation based on vectorized spatio-temporal graph convolutional for virtual reality training in smart grid. *Energies* 15, 6 (2022), 2071.

[10] Kate E Laver, Belinda Lange, Stacey George, Judith E Deutsch, Gustavo Saposnik, and Maria Crotty. 2017. Virtual reality for stroke rehabilitation. *Cochrane database of systematic reviews* 11 (2017).

[11] Yalin Liao, Aleksandar Vakanski, and Min Xian. 2020. A deep learning framework for assessing physical rehabilitation exercises. *IEEE Transactions on Neural Systems and Rehabilitation Engineering* 28, 2 (2020), 468–477.

[12] Jonathan Feng-Shun Lin, Michelle Karg, and Dana Kulić. 2016. Movement primitive segmentation for human motion modeling: A framework for analysis. *IEEE Transactions on Human-Machine Systems* 46, 3 (2016), 325–339.

[13] Steven R Machlin, Julia Chevan, William W Yu, and Marc W Zodet. 2011. Determinants of utilization and expenditures for episodes of ambulatory physical therapy among adults. *Physical therapy* 91, 7 (2011), 1018–1029.

[14] Elham Mottaghi and Mohammad-R Akbarzadeh-T. 2022. Automatic evaluation of motor rehabilitation exercises based on deep mixture density neural networks. *Journal of Biomedical Informatics* 130 (2022), 104077.

[15] Octavian Postolache, Luís Teixeira, João Cordeiro, Luisa Lima, Patrícia Arriaga, Mariana Rodrigues, and Pedro Girão. 2019. Tailored virtual reality for smart physiotherapy. In *2019 11th International Symposium on Advanced Topics in Electrical Engineering (ATEE)*. IEEE, 1–6.

[16] Sara Sardari, Sara Sharifzadeh, Alireza Daneshkhab, Bahareh Nakisa, Seng W Loke, Vasile Palade, and Michael J Duncan. 2023. Artificial Intelligence for skeleton-based physical rehabilitation action evaluation: A systematic review. *Computers in Biology and Medicine* (2023), 106835.

[17] Simon Senecal, Niels A Nijdam, Andreas Aristidou, and Nadia Magnenat-Thalmann. 2020. Salsa dance learning evaluation and motion analysis in gamified virtual reality environment. *Multimedia Tools and Applications* 79, 33 (2020), 24621–24643.

[18] Rosanna Maria Vigliani, Sara Condino, Giuseppe Turini, Marina Carbone, Vincenzo Ferrari, and Marco Gesi. 2019. Review of the augmented reality systems for shoulder rehabilitation. *Information* 10, 5 (2019), 154.

[19] Zhe Zhang, Qiang Fang, Liuping Wang, and Peter Barrett. 2011. Template matching based motion classification for unsupervised post-stroke rehabilitation. In *International Symposium on Bioelectronics and Bioinformatics 2011*. IEEE, 199–202.