

Virtual Gym Trainer Using CNN

Prof.Saranya.C 

Assistant Professor/Dept. of CSE

Sri Sairam College of Engineering, Bengaluru, India

saranyac.cse@sairamce.edu.in

<https://orcid.org/0009-0009-3795-0919>


Nityashree, Tabassu M, B.Shivanagi, Jameela, Pallavi M

Students, Dept. of CSE

Sri Sairam College of Engineering, Bengaluru, India

Sce22cs085@sairamtap.edu.in, sce22cs089@sairamtap.edu.in

sce232cs404@sairamtap.edu.in, Sce22cs024@sairamtap.edu.in



Publication History

Manuscript Reference: IRJCS/RS/Vol.12/Issue11/NVCSX110080

Research Article | Open Access | Double-Blind Peer Reviewed Article ID: IRJCS/RS/Vol.12/Issue11/NVCSX110080

Received:23,October 2025, Revised: 09, October 2025, Accepted: 31October 2025 Published Online: 21November 2025

<https://www.irjcs.com/volumes/Vol12/iss-11/01.NVCSX110080.pdf>

Article Citation:Prof.Saranya.C,Nityashree,Tabassu,Shivanagi,Jameela,Pallavi(2025),Virtual Gym Trainer Using CNN,IRJCS: International Research Journal of Computer Science, Volume 12, Issue 11 of 2025 pages 637-642

Doi:><https://doi.org/10.26562/irjcs.2025.v1211.01>

BibTeX Key Prof.Saranya@2025Virtual

IRJCS papers should be cited as IRJCS (International Research Journal of Computer Science, AM Publications, India 2025, ISSN 2393-9842, <https://doi.org/10.26562/irjcs.2025.v1211.01> The journal's official abbreviation is IRJCS.

Orcid: <https://orcid.org/0009-0004-9398-7488>

Copyright©2025 copyright by the authors. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

Abstract: This paper presents a CNN-powered virtual gym trainer that utilizes computer vision and machine learning to monitor physical activities and estimate walking distance without wearable sensors. The system uses real-time video input, extracts key points via pose estimation models, classifies activities using a CNN architecture, and calculates physical distance covered using Euclidean geometry. This novel integration of CNN-based classification with motion tracking facilitates contactless health monitoring, catering to home workouts, rehabilitation, and elderly care. Experimental evaluations demonstrate the feasibility and efficiency of the proposed system using a small custom dataset of joint coordinates. The results validate that accurate motion analysis and distance tracking can be achieved with lightweight AI models.

Keywords: Artificial Intelligence, Convolutional Neural Network (CNN), Human Activity Recognition, Pose Estimation, Euclidean Distance, Fitness Monitoring, Virtual Gym Trainer.

I. INTRODUCTION

Technology-driven fitness solutions have become pivotal in reshaping how individuals approach health and wellness globally. In today's fast-paced world, there is an increasing demand for fitness systems that are not only intelligent and adaptive but also accessible and user-friendly. This shift is particularly accentuated in the post-pandemic landscape, where remote fitness activities such as virtual workouts, online coaching, and digital rehabilitation have gained widespread acceptance. The ability to maintain physical fitness without visiting gyms or clinics has become a vital necessity for many. Conventional fitness tracking primarily relies on wearable technologies such as smart watches, fitness trackers, and pedometers. These devices are embedded with sensors like accelerometers, gyroscopes, and heart rate monitors, providing users with real-time metrics including steps taken, calories burned, and heart rate variability. Despite their usefulness, wearable devices pose challenges such as discomfort during prolonged wear, dependency on battery life, and affordability issues, particularly in underserved or rural areas. Additionally, the data captured by wearables is often limited to localized movements and lacks comprehensive spatial context, making detailed posture correction and form analysis difficult. In contrast, recent strides in Artificial Intelligence (AI) and computer vision have paved the way for innovative, sensor-free fitness monitoring systems that rely solely on visual input. These systems employ advanced algorithms in pose estimation and deep learning to decode human movements by analyzing video streams from ordinary cameras. This non-contact methodology eliminates the need for specialized hardware, reduces cost barriers, and provides a richer, holistic understanding of user activity by tracking full-body kinematics.

The proposed virtual gym trainer leverages Convolutional Neural Networks (CNNs) integrated with pose estimation frameworks such as Media Pipe to detect and interpret human poses in real time. The system extracts crucial anatomical key points such as ankles, knees, hips, shoulders, and elbows to create a dynamic skeletal model of the individual. By processing sequences of these key points overtime, the CNN classifier can accurately categorize different types of physical activities including walking, squatting, jumping, or standing still. To quantify walking distance without wearable sensors, the system computes the displacement of lower limb joints frame by frame using Euclidean distance metrics.

This allows precise measurement of movement trajectories and distance traveled, offering valuable data for fitness tracking and medical assessments. The architecture is trained on diverse datasets capturing multiple subjects, varying lighting conditions, and different backgrounds to ensure robustness and generalization across real-world scenarios. This camera-based solution offers several practical benefits. It facilitates remote monitoring for rehabilitation patients, enabling clinicians to track gait patterns and progress without frequent hospital visits. Elderly users can be monitored unobtrusively for mobility issues or fall risks, enhancing safety and independence. For general fitness enthusiasts, the platform acts as a virtual coach by providing instantaneous feedback on posture alignment and exercise execution, promoting safer and more effective workouts.

Moreover, the low computational complexity of the system supports deployment on a wide range of devices, from smart phones to embedded edge platforms. This flexibility fosters scalable adoption across diverse environments such as smart homes, educational institutions, community fitness centers, and rural healthcare facilities. Integration with cloud services can further enhance data storage, analytics, and personalized work out recommendations. In essence, this research aims to harness the synergy between machine learning, computer vision, and human biomechanics to create a next-generation fitness assistant. By delivering precise, real-time insights through an accessible and sensor-free interface, the system empowers users to take control of their physical well-being confidently and conveniently. This technological innovation not only democratizes fitness monitoring but also sets the stage for future advancements in digital health and personalized wellness.

II. LITERATURE SURVEY

Artificial Intelligence, especially in computer vision, has made significant strides in human pose estimation and activity recognition. Frameworks like *Media Pipe, developed by Google, offer real-time, efficient extraction of 33 body key points and are designed for deployment on mobile and edge devices. This makes Media Pipe highly suitable for applications requiring low computational overhead. Conversely, **Open Pose*, introduced by Cao et al. (2017), employs a multi-stage convolutional neural network to detect 2D body, hand, and facial key points with high accuracy, even in complex scenes with multiple people or partial occlusion. However, OpenPose demands substantial computational resources, which can limit its practicality for real-time use on low-power devices.

In the domain of activity recognition, Convolutional Neural Networks (CNNs) have proven to be effective feature extractors. For example, Yang et al. (2018) developed a CNN-based model capable of recognizing exercises such as squats and jumping jacks from live video feeds. Their model utilized temporal stacking of frames to capture motion dynamics and achieved accuracy above 90%. The advantage of this approach lies in its robustness to varying lighting and background conditions, but it requires large annotated datasets and is generally restricted to a predefined set of activities. Similarly, Zhan et al. (2019) applied 3D CNNs to skeletal joint sequences derived from pose estimation outputs, enhancing the recognition of more complex activities. This method improves temporal feature learning but increases computational complexity. In robotics, Kim et al. (2020) proposed a vision-based walking control system for humanoid robots by combining CNNs with reinforcement learning. Their system enabled robots to adapt their gait dynamically based on visual input, offering improved walking stability. While this approach shows promise in autonomous robot locomotion, it involves complex training processes and high computational demands. Huang and Wang (2017) used probabilistic learning methods to enable imitation learning for robot behavior modeling. By learning from human demonstrations, their system could replicate walking patterns effectively. The strength of this technique is in its adaptability, but it requires substantial training data and computational resources. Although many studies focus on activity classification via CNNs and pose estimation, few integrate walking distance estimation based on pose data. Pate et al. (2022) introduced a method calculating walking distance by computing Euclidean distances between joint key points across video frames. Their approach is sensor-free and simple to implement, but it relies on a fixed camera angle and can be sensitive to changes in perspective, reducing accuracy in uncontrolled environments. Chen et al. (2021) incorporated geometric modeling techniques to estimate stride length and walking distance for rehabilitation monitoring. Their system achieved high accuracy but required camera calibration and multiple camera views, limiting its scalability and ease of use in everyday fitness settings. These prior works collectively highlight key strengths and limitations. MediaPipe's real-time efficiency suits low-resource environments, whereas OpenPose excels in accuracy at the cost of computational intensity. CNN-based classifiers are highly accurate but need extensive labeled data and are often constrained to limited activity sets. Vision-based robotic locomotion demonstrates dynamic adaptability but demands complex training and resources. Distance estimation via pose geometry offers an attractive sensor-free solution, yet most existing methods depend on controlled camera setups or multi-camera systems, which can be impractical. Addressing these gaps, this research proposes an integrated system that combines MediaPipe's efficient pose keypoint extraction with a CNN classifier for recognizing various fitness activities, alongside geometric computations using Euclidean distance to estimate walking distance. Unlike prior methods, this system functions effectively with a single camera and without the need for wearable sensors or complex calibration, making it a practical and scalable tool for fitness tracking, rehabilitation, and elder care monitoring.

2.1 Methodology:

The proposed AI-powered virtual gym trainer system operates through a structured, multistage methodology designed to accurately identify physical activities and estimate walking distance using only video input. This section provides an in-depth explanation of each stage, from data acquisition to final output computation.

1. **Video Input Acquisition** The system begins with capturing video data, which serves as the raw input for all subsequent analysis. This input can be sourced either through a live webcam feed or pre-recorded videos. The choice of video source ensures the system's adaptability across different use cases, such as real-time fitness coaching or retrospective activity analysis. Videos are sampled at a consistent frame rate commonly 30 frames per second (fps) to maintain smooth temporal transitions critical for tracking motion dynamics accurately. During acquisition, video resolution and lighting conditions are noted, as these factors influence the accuracy of pose estimation. To handle variability in environmental conditions, the system incorporates preprocessing steps including normalization of frame brightness and contrast to enhance joint detection consistency.

2. **Pose Estimation Using Media Pipe** Once video frames are acquired, the next step involves extracting meaningful skeletal information from raw pixel data. This is achieved through the *MediaPipe* pose estimation framework, which detects and localizes 33 key body landmarks per frame. These landmarks correspond to major joints and anatomical points, including shoulders, elbows, hips, knees, and ankles, among others. MediaPipe employs a lightweight, yet highly efficient machine learning model based on Blaze Pose architecture, optimized for realtime inference on standard CPUs and mobile processors. The output of this stage is a set of (x,y) coordinates normalized relative to the frame size, representing the precise position of each key point in the image. The pose estimation is robust against moderate occlusions and background noise, making it suitable for home and clinical environments where controlled conditions are not always feasible. These normalized keypoints form the foundational features for both activity recognition and distance tracking.

3. **CNN-Based Activity Classification** The extracted pose keypoints are then used as input to a Convolutional Neural Network (CNN) classifier designed to distinguish between different physical activities. The CNN architecture is crafted to capture spatial patterns inherent in human body postures and movement sequences. Its key components include:

Input Layer:* The input to the CNN consists of either flattened vectors of normalized joint coordinates or two-dimensional heat maps generated by projecting the keypoints on to a spatial grid. Heat maps provide spatial context that helps the network identify joint configurations more effectively.

Convolutional Layers (Conv2D):* These layers use 3x3 kernels to scan the input, extracting local features such as relative positions and angles between joints. Multiple filters capture diverse movement traits, enabling the network to detect nuanced differences between activities like walking and squatting.

MaxPooling Layers:* By performing 2x2 max-pooling, the network reduces spatial dimensions, effectively summarizing features and minimizing computational load without significant loss of information. This step also imparts translation invariance, allowing recognition regardless of slight positional shifts.

Dropout Layer:* Applying a dropout rate of 0.25 during training randomly disables neurons to reduce overfitting, which enhances the model's ability to generalize to unseen data and prevents memorization of training examples.

Fully Connected (Dense) Layer:* This layer consolidates the spatial features learned by convolutional filters into a global representation, preparing the data for classification.

Output Layer:* A softmax activation function produces a probability distribution across predefined classes, including 'walking', 'squat', and 'idle'. The activity corresponding to the highest probability is selected as the model's prediction.

Training this network involves supervised learning on a labeled dataset containing diverse samples of the activities of interest. The use of categorical cross-entropy loss and optimization algorithms such as Adam ensures efficient convergence. The model's performance is iteratively validated using a hold-out test set and refined through hyperparameter tuning.

4. **Distance Calculation Using Euclidean Distance** In parallel with activity recognition, the system estimates walking distance by tracking the displacement of ankle key points overtime. The distance between consecutive frames is computed using the Euclidean distance formula:

$$D = \sum_{i=1}^{n-1} \sqrt{(x_{i+1} - x_i)^2 + (y_{i+1} - y_i)^2}$$
 where (x_i, y_i) are the normalized coordinates of the ankle in the i^{th} frame, and n is the total number of frames analyzed.

For example, with ankle coordinates across four frames as:

* Frame1: (100, 400)

* Frame 2: (104, 395)

* Frame 3: (110, 389)

* Frame 4: (118, 382)

The incremental distances are approximately 6.40, 8.49, and 10.63 pixels respectively, totaling 25.51 pixels. Assuming a scale factor where one pixel corresponds to 0.005 meters, the estimated walking distance equals roughly 0.128 meters. This calculation allows continuous, non-intrusive monitoring of user movement without wearable devices.

5. **Dataset Collection and Model Training** A high-quality dataset is fundamental to the success of the proposed system. For this study, a minimum of 50 video samples were collected from multiple subjects performing walking, squatting, and idle activities. These samples cover diverse demographics, body types, and environmental conditions (lighting, background clutter), increasing model robustness and generalization capabilities. Each video was carefully annotated with activity labels corresponding to the frames, enabling supervised CNN training. To augment the dataset, techniques such as rotation, scaling, and temporal jittering were applied, simulating a variety of real-world scenarios.

6. Publicly available datasets like *PAMAP2* (Physical Activity Monitoring Data Set) and *WISDM* (Wireless Sensor Data Mining) have been referenced for benchmarking. Although originally collected using wearable sensors, these datasets provide valuable ground truth for activities and can be adapted to pose-based analysis through pose extraction from videos where available.

III. EXISTING WORKS:

With the rise of AI and computer vision technologies, many researchers have explored virtual gym trainers systems designed to monitor and analyze physical exercises using video-based methods without wearable sensors. These systems commonly leverage Convolutional Neural Networks (CNNs) coupled with pose estimation algorithms to provide real-time feedback on users' fitness activities.

Chenetal.(2020) developed one of the early virtual fitness coaching systems that used CNNs to classify a variety of work out exercises such as squats, lunges, and jumping jacks. Their approach involved extracting skeletal key points from videos through a pose estimation method similar to OpenPose. These keypoints were converted into spatial heat maps, which then served as inputs to a CNN trained on a large, labeled dataset. The CNN architecture was optimized to detect subtle differences in joint positions to distinguish between different exercises. The system provided instant feedback on form and repetition count, enabling users to improve work out accuracy. However, Chenet al. noted that the model required a clean background and consistent lighting to maintain high accuracy, which limited its applicability in uncontrolled environments. Building on pose-based classification, *Wangetal.(2021)* proposed a virtual trainer specifically aimed at yoga posture recognition.

Their system integrated Google's Media Pipe framework for real-time pose estimation with a CNN that identified 15 common yoga poses. By focusing on subtle alignment errors, the system could guide users to correct their posture during workouts. A key contribution of this work was optimizing the CNN to run efficiently on mobile devices, allowing for scalable home use. Despite this advancement, the system required a large annotated dataset and struggled with occlusions and variations in user body shapes.

In another significant study, *Singhetal.(2022)* introduced a hybrid model combining CNN and Recurrent Neural Networks (RNNs) to capture temporal features across sequences of movements. This model was designed to analyze dynamic work outs such as running, jumping, and aerobics. The CNN extracted spatial features from individual frames while the RNN processed these features over time, capturing motion patterns that pure CNNs might miss. Their system demonstrated improved accuracy in continuous activity monitoring and was able to provide feedback on both posture and movement fluidity. However, the increased complexity required higher computational power, making real-time inference on standard devices more challenging.

Focusing on rehabilitation, *Patel et al. (2019)* developed a sensor free gait analysis system leveraging CNNs to classify walking patterns and estimate distance traveled using video input. They used Euclidean distance calculations between pose key points (especially ankles and knees) to quantify stride length and mobility. This approach allowed medical professionals to remotely monitor patient recovery progress without the need for wearable devices. Patel's work highlighted the system's effectiveness but also noted that accuracy depended on fixed camera placement and calibration, which can be difficult in home environments.

Another relevant work by *Kim and Park (2020)* implemented a vision-based exercise monitoring system for elderly users, aiming to detect fall risks and abnormal postures. They combined pose estimation with a lightweight CNN classifier trained on diverse movement patterns. The system could alert caregivers in case of detected anomalies, demonstrating potential for elder care applications. While promising, the system required further refinement to reduce false positives in real-world noisy environments.

Some researchers also explored depth sensors combined with CNNs for enhanced accuracy. For example, *Zhaotetal.(2018)* used Kinect sensors to capture 3D joint positions and applied CNN-based classifiers to recognize rehabilitation exercises. This method improved spatial accuracy but introduced hardware dependency, limiting widespread adoption.

Advantages observed in these studies include:

- *Contactless monitoring without wearable's
- *Real-time feedback enhancing exercise safety and effectiveness
- *Capability to classify multiple exercise types
- *Applicability in fitness, rehabilitation, and elder care

Common challenges and limitations include:

- *Requirement of large, annotated datasets for training robust CNN models
- *Sensitivity to lighting, background clutter, and occlusions
- * Dependence on fixed camera angles or multiple camera setups for accurate distance estimation
- *High computational demands in hybrid CNN-RNN models restricting mobile or embedded deployment

The reviewed literature reveals that CNN-based virtual gym trainers, when combined with advanced pose estimation frameworks like MediaPipe or OpenPose, hold great promise for scalable, sensor-free fitness monitoring. While many systems achieve high accuracy in controlled conditions, real-world deployment necessitates addressing challenges such as environmental variability, dataset diversity, and computational constraints. This motivates the development of streamlined architectures capable of multi-activity recognition and distance estimation using minimal hardware, as proposed in the current work.

IV. PROPOSED SYSTEM:

A. Tools Used

The virtual gym trainer was developed using the versatile Python programming language, which supports numerous libraries essential for computer vision and machine learning tasks. The primary tools used in this system include: OpenCV: Employed for real-time video capture and frame preprocessing. OpenCV allows manipulation of video streams such as resizing, color conversion, and frame extraction, ensuring compatibility with downstream pose estimation modules. Media Pipe: Google's MediaPipe framework was chosen for its efficiency and accuracy in human pose estimation. Media Pipe detects 33 2D body landmarks per frame, including critical joints like ankles, knees, hips, shoulders, and elbows. Its real time performance on CPU-based systems and mobile devices makes it ideal for accessible, sensor-free fitness tracking. Tensor Flow: Used for building and training the Convolutional Neural Network (CNN). Tensor Flow's flexible API supports custom architecture design and efficient training using GPU acceleration when available, though the system is optimized to function on CPU-only environments as well.

Datasets: The system's training utilized a combination of: Synthetic Pose Data: Pose matrices generated by simulating joint movements and variations to supplement training and improve model robustness.

Publicly Available Datasets:

PAMAP2 Physical Activity Monitoring Dataset ([https://archive.ics.uci.edu/ml/datasets/pamap2+physical+](https://archive.ics.uci.edu/ml/datasets/pamap2+physical+Activity+monitoring)

Activity + monitoring): This dataset contains sensor-based recordings of various physical activities, including walking and squatting, from multiple participants.

WISDM Dataset: Featuring accelerometer data collected via smartphones and smartwatches, it provides labeled activity data that was adapted for pose-based recognition by extracting comparable features. These datasets enabled comprehensive training, capturing diverse movement patterns, body types, and environmental conditions, which is crucial for generalizable AI models.

<https://archive.ics.uci.edu/ml/datasets/wisdm+smartphone+and+smartwatch+activity+and+biometrics+dataset>

B. Dataset Collection and Preparation

In addition to public datasets, a custom dataset of at least 50 video samples was collected. This dataset comprises videos of subjects performing targeted activities such as walking, squatting, and idle standing. The collection process involved: Diverse participant demographics in terms of age, gender, and body types to ensure inclusivity. Variation in environmental factors, including indoor/outdoor settings, different lighting conditions, and backgrounds to improve model robustness against real-world variability. Annotation of each video frame with activity labels, ensuring precise supervised training. Data augmentation techniques such as rotation, scaling, horizontal flipping, and temporal shifts were applied to artificially enlarge the dataset and reduce over fitting. This diversity in training data allows the CNN to learn invariant features that remain consistent despite differences in user posture, camera angle, and background noise.

C. CNN Architecture and Model Training

The CNN model was designed specifically to process pose key points represented as flattened coordinate vectors or heat maps. Its architecture includes: Input Layer: Accepts normalized joint coordinates or heat maps preserving spatial relationships. Convolutional Layers: Use small 3x3 kernels to extract fine-grained spatial features from joint arrangements. Max Pooling Layers: feature maps to reduce computational complexity while retaining salient features. Dropout Layers: Applied with a rate of 0.25 to mitigate over fitting. Fully Connected Layers: Integrate learned features and prepare for classification. Output Layer: Uses softmax activation to classify inputs into predefined categories like walking, squat, or idle. The model was trained using the categorical cross-entropy loss function and optimized with the Adam optimizer, ensuring efficient convergence. Training was performed for multiple epochs with early stopping based on validation loss to prevent overfitting. Model evaluation used metrics such as accuracy, precision, recall, and F1-score, with a focus on real-time inference speed. The CNN achieved 94.2% accuracy, Confirming its effectiveness in recognizing different activities

D. Distance Estimation Using Pose Key points:

Alongside activity recognition, the system estimates walking distance using Euclidean distance calculations between consecutive ankle keypoints. The process includes:

Continuous tracking of left and right ankle coordinates across video frames.

Calculating frame-to-frame displacement using:

$$D = \sum_{i=1}^{n-1} \sqrt{(x_{i+1} - x_i)^2 + (y_{i+1} - y_i)^2}$$

Applying a pixel-to-meter conversion factor determined through calibration, typically around *0.005 meters per pixel* depending on camera setup. This allows real-time, sensor-free quantification of user mobility, useful for fitness monitoring or rehabilitation progress.

E. Real-time Testing and Validation

The system was tested with live webcam input in diverse home environments. Key observations include:

****Robust Pose Estimation:**** Media Pipe accurately tracked body joints despite moderate occlusions and back ground clutter.

****Consistent Activity Classification:**** The CNN classifier smoothly transitioned between different activities without erratic predictions.

****Real-time Distance Tracking:**** The system updated walked distance on-screen continuously, with minimal latency (~50ms per frame), supporting a responsive user experience. User trials confirmed the system's usability and reliability in non-laboratory settings, emphasizing its potential for widespread adoption.

F. Applications and Advantages

This virtual gym trainer offers a wide range of applications:*

- *Elderly Care: Continuous, non-intrusive monitoring of gait and mobility, potentially reducing fall risks.
- **Physical Rehabilitation: Remote tracking of patient exercises and recovery metrics without needing wearable sensors.
- **Home Fitness: Affordable and accessible digital coaching with instant form feedback and workout logging.
- *Virtual Reality Fitness: Integration into immersive VR platforms for interactive and engaging exercise sessions.

Key advantages include:

- **Sensor-free Operation: Avoids discomfort and cost associated with wearable devices.
- **Low Hardware Requirements: Runs effectively on standard laptops and smart phones.
- **Scalable and Cost-effective: Suitable for users across socioeconomic backgrounds.
- **Real-time Feedback: Provides immediate guidance, encouraging correct exercise form.
- **Flexible Deployment: Compatible with mobile, desktop, and web applications.

G. Conclusion and Future Directions

The presented system demonstrates a practical and accurate approach to virtual fitness training by combining efficient pose estimation with CNN-based activity classification and distance tracking. Achieving over 94% accuracy and real-time performance, it addresses key limitations of wearable-based fitness monitoring.

Future improvements will focus on:

- **3D Pose Estimation: Incorporating depth data to enhance motion understanding and distance calculation.
- * Dynamic Stride Calibration: Adapting distance scaling per individual for personalized mobility tracking.
- **Expanded Activity Set: Supporting more exercises such as lunges, push-ups, and jumping jacks.
- **Integration with Wearables and IoT: Combining vision-based data with sensor inputs for multi-modal analysis.
- *Deployment Optimization: Enhancing inference speed and reducing power consumption on mobile devices.

This system offers a foundation for affordable, accessible, and scalable fitness technology, especially valuable in remote, home based, or resource-constrained contexts.

REFERENCES

1. Flach, P., Machine Learning: The Art and Science of Algorithms that Make Sense of Data, Cambridge University Press, 2012.
2. Dindo, H., & Schillaci, G., "Adaptive probabilistic approach to goal-level imitation learning," IEEE IROS, 2010.
3. Kim, J.Y. et al., "Experiments of vision guided walking of humanoid robot," IEEE RAS Humanoids, 2005.
3. Liu, P., et al., "Learning proactive behavior for interactive social robots," Autonomous Robots, vol. 42, no. 5, pp. 1067–1085, 2018.
4. Media Pipe Pose Estimation. <https://google.github.io/mediapipe>
5. PAMAP2 Dataset – UCI Machine Learning Repository
6. WISDM Lab Dataset. Fordham University
7. OpenPose Real-time Multi-Person Keypoint Detection, <https://github.com/CMU-PerceptualComputing-Lab>