Mobile Application for Optimizing Exercise Posture Through Machine Learning and Computer Vision in Gyms

Kendall Contreras-Salazar, Paulo Costa-Mondragon and Willy Ugarte[a]
Universidad Peruana de Ciencias Aplicadas (UPC), Lima, Peru

Keywords: Pose Estimation, Machine Learning, Computer Vision, LSTM, MediaPipe, Ionic, Exercise, Gym, Injury, Mobile Application, Posture.

Abstract: This paper introduces a mobile application that aims to improve exercise posture analysis in gym environments using machine learning and computer vision. The solution processes user-uploaded videos to detect posture errors, utilizing Long Short-Term Memory (LSTM) networks and MediaPipe for precise pose estimation. The trained model achieved high accuracy in classifying exercise postures, demonstrating reliable performance across different user scenarios. Traditional posture correction methods, such as personal trainers and wearable devices, often lack accessibility and precision. In contrast, our application offers a scalable, user-friendly tool that delivers actionable feedback, helping users optimize their workouts and reduce injury risks. The study highlights the potential of combining machine learning with mobile technology to enhance exercise safety and performance, setting a foundation for future improvements.

1 INTRODUCTION

The fitness industry is continuously evolving, with more people becoming aware of the importance of exercise for physical and mental well-being. However, with this growing awareness comes an increase in the risk of injury, especially in unsupervised gym settings. Poor posture during exercises like squats and deadlifts can lead to serious injuries, hindering progress and long-term health. Recent studies in The Netherlands reveal that 73.1% of gym-related injuries occur during unsupervised sessions, often due to improper posture (Kemler et al., 2022). Addressing this issue requires innovative solutions that can provide posture correction without the need for expensive personal trainers. This work presents a mobile application designed to assist gym-goers in maintaining proper posture during exercises.

The app uses a combination of machine learning and computer vision to analyze user movements and provide feedback on posture accuracy. By focusing on user-uploaded videos, the system offers an accessible and scalable solution to a widespread problem in fitness training. The core of this project lies in the integration of two powerful technologies: Long Short-Term Memory (LSTM) networks and the MediaPipe framework. LSTM networks, which excel at analyzing sequential data, are particularly well-suited for dynamic gym exercises where movements are fluid.

MediaPipe, an open-source framework for pose estimation, allows for precise detection of key body points during exercises. These two components work together to deliver accurate, actionable feedback to users after their workout sessions. Traditional solutions for posture correction, such as in-person trainers or wearable devices, come with significant drawbacks. Trainers, while effective, are costly and not always accessible. Wearable devices, on the other hand, can track basic metrics but often lack the precision needed to assess complex, multi-joint movements like those involved in strength training (Vali et al., 2024). Our mobile application addresses these limitations by providing a cost-effective alternative that can be used by anyone with a smartphone.

Several recent studies have explored the use of machine learning for posture recognition. For example, Mallick et al. employed LSTM networks and Hidden Markov Models to recognize postures in Bharatanatyam dance sequences, demonstrating the effectiveness of these models in capturing temporal dynamics (Mallick et al., 2022). Similarly, a study on yoga posture recognition using LSTM networks and pose estimation achieved high accuracy in classifying static postures (Palanimeera and Ponmozhi,

[a] https://orcid.org/0000-0002-7510-618X

Contreras-Salazar, K., Costa-Mondragon, P. and Ugarte, W.
Mobile Application for Optimizing Exercise Posture Through Machine Learning and Computer Vision in Gyms.
DOI: 10.5220/0013439300003938

In *Proceedings of the 11th International Conference on Information and Communication Technologies for Ageing Well and e-Health (ICT4AWE 2025)*, pages 360-367
ISBN: 978-989-758-743-6; ISSN: 2184-4984

2023). These works highlight the potential of machine learning in movement analysis, further validating the approach taken. However, each of these works faces specific limitations that are addressed by our solution. For instance, Mallick et al.'s approach to Bharatanatyam posture recognition was limited by the complexity of the dance movements and the need for synchronization with music. Their method, which relied heavily on Hidden Markov Models, struggled with the temporal variability of the movements and was highly domain-specific (Mallick et al., 2022). In contrast, our solution avoids these constraints by focusing on gym exercises, where the movements are more standardized and easier to track.

The use of LSTM networks allows our system to handle the dynamic nature of gym exercises while providing feedback without the need for synchronization with external factors like music. Similarly, the YAP-LSTM study on yoga posture recognition achieved high accuracy, but it was primarily focused on static postures (Palanimeera and Ponmozhi, 2023). Yoga, by nature, involves slower and more controlled movements compared to gym exercises, making it easier to track and classify. Our work, on the other hand, tackles the challenge of highly dynamic, multi-joint movements in gym exercises. By leveraging LSTM networks, which excel at processing sequential data, we are able to analyze and provide feedback on these complex movements. Furthermore, MediaPipe's pose estimation ensures that even minor deviations in form are detected and corrected, something that the yoga study did not fully address due to its focus on static positions.

In (Kaewrat et al., 2024), the augmented reality (AR) for exercise monitoring also faced limitations related to the type of exercises being monitored and the technology used. While AR provided an innovative approach to offering feedback, it was primarily focused on simple movements like marching in place, which do not capture the complexity of exercises typically performed in the gym. Our solution focuses on providing feedback based on pre-recorded videos, allowing users to concentrate fully on their workout without interruptions. Additionally, our system's ability to handle more complex movements like squats and deadlifts sets it apart from the simpler movements monitored in AR-based systems. Physiotherapy assistance systems, like the one developed by Dudekula et al., are designed to help patients maintain proper form during rehabilitation exercises using pose estimation technologies such as MediaPipe (Vali et al., 2024). However, these systems are often tailored to slower, more controlled physiotherapy movements, limiting their applicability to the fast-paced, dynamic nature of gym exercises. Our application builds upon the strengths of pose estimation in physiotherapy by adapting it to handle the speed and complexity of gym movements, ensuring that even subtle errors in posture are detected. To demonstrate that our solution meets its objectives, we will employ a comprehensive evaluation methodology. The first step will involve gathering a dataset of gym exercises performed by users of varying experience levels.

This dataset will include both correct and incorrect executions of exercises like squats, benchpress, and deadlifts. These videos will be annotated with ground truth labels indicating the correctness of the posture, which will serve as the benchmark for evaluating the system's performance. The system's performance will be evaluated based on its accuracy in detecting posture errors, the clarity of the feedback provided, and user satisfaction. To measure accuracy, we will compare the system's feedback with the ground truth labels, calculating metrics such as precision, recall, and F1-score. We will also conduct user studies to assess how effectively the system's feedback helps users correct their posture and improve their form over time. Additionally, the usability of the system will be evaluated through user experience surveys, focusing on factors such as ease of use, clarity of instructions, and overall satisfaction. These surveys will provide valuable insights into how well the system integrates into users' workout routines and whether the feedback is intuitive and actionable.

In conclusion, our mobile application offers a robust solution to the problem of posture correction in gym exercises, addressing the limitations faced by previous approaches while introducing new capabilities for handling dynamic, multi-joint movements. By leveraging LSTM networks and MediaPipe's pose estimation, we provide users with a powerful tool to improve their form, reduce the risk of injury, and enhance their overall workout experience. Through rigorous evaluation and user testing, we will demonstrate that our solution not only meets but exceeds the needs of gym-goers seeking to optimize their exercise performance.

This article is distributed in the following sections: first, we review related works on posture detection for exercises in Section 2. Then, we discuss classification algorithms and their effectiveness in our research in Section 3 and describe our main contribution in more detail. Additionally, we will explain the procedures carried out and the experiments conducted in this work in Section 4. Finally, we will show our main conclusions in Section 5.

2 RELATED WORKS

This section highlights related work that employs advanced machine learning techniques, particularly LSTM networks and pose estimation, to recognize and classify human postures in different contexts. These articles showcase the versatility of these approaches in handling both static and dynamic movements, while also addressing the limitations and challenges associated with each application domain.

In the article (Mallick et al., 2022), the authors develop a method to analyze Bharatanatyam dance by segmenting video sequences to identify and recognize key postures using Convolutional Neural Networks (CNNs). They further enhance the system with Hidden Markov Models (HMMs) and Long Short-Term Memory (LSTM) networks to capture the temporal sequence of dance movements. Unlike our work, which focuses on using LSTM models and PoseNet to classify and correct gym exercise postures, this work emphasizes the recognition and sequencing of dance postures for cultural and educational purposes, integrating audio cues to enhance accuracy.

In (Palanimeera and Ponmozhi, 2023), the authors present a method that integrates pose estimation with LSTM models to classify yoga asanas from real-time video data. The system uses OpenPose to extract body key points, which are then input into an LSTM network to capture the temporal dynamics of the yoga poses, achieving high accuracy in asana recognition. Unlike their approach, which is tailored to the static and structured nature of yoga poses, our work focuses on classifying dynamic gym exercises, which present unique challenges due to the complexity and variability of movements, making the application of LSTM and computer vision techniques specifically adapted to handle these challenges.

In (Müller et al., 2024), the authors propose a mobile AR application for exercise monitoring that leverages pose estimation and AR technologies to provide real-time feedback on exercise form. Unlike traditional methods that rely heavily on wearable devices or in-person assessments, this approach uses RGB cameras and LiDAR sensors to track key anatomical landmarks during exercises like marching-in-place. The application utilizes MediaPipe for 2D pose estimation and ARFoundation for 3D depth sensing, calculating joint angles to determine exercise correctness. Visual and auditory feedback is provided to users through AR overlays, helping them adjust their posture in real-time. Unlike our work, which centers on developing a mobile application using the Ionic framework to upload and classify posture of the exercises, this work leverages AR to provide real-time feedback during the exercise.

In (Kemler et al., 2022), the authors presents a descriptive epidemiological study focusing on gym-based fitness-related injuries among 494 Dutch participants, emphasizing the significant role of unsupervised activities and poor posture in injury occurrence. The study found that 73.1% of injuries happened during unsupervised gym-based activities, with strength training and individual cardio exercises being the most common. The shoulder, leg, and knee were the most frequently injured body parts, often due to overuse, incorrect posture, or improper movement. The findings highlight the need for injury prevention strategies that emphasize proper technique and possibly increased supervision during complex exercises to reduce injury risks in unsupervised settings. The findings underscore the importance of developing injury prevention strategies that prioritize proper technique and increased supervision, particularly for complex exercises, to mitigate injury risks in unsupervised settings. Our work seeks to address this by classifying and supervising exercises to proactively prevent such injuries using videos recorded by the same user.

In (Vali et al., 2024), the authors discusses the use of MediaPipe for human pose estimation in a physiotherapy assistance system integrated with Raspberry Pi. MediaPipe's real-time pose estimation capabilities play a crucial role in monitoring and correcting patient postures during physiotherapy exercises. By accurately identifying body key points, MediaPipe allows the system to detect and correct improper postures, which is essential for preventing further injuries and ensuring effective rehabilitation. This approach is especially beneficial in remote or unsupervised settings, where traditional supervision might not be possible. Our work leverages MediaPipe's pose estimation to classify and supervise exercises, aiming to prevent incorrect posture and related injuries, thereby enhancing the safety and efficacy of rehabilitation.

3 MAIN CONTRIBUTION

This section outlines the theoretical framework, which allows our system to learn and improve posture analysis in exercise.

3.1 Preliminary Concepts

Our work, relies on key concepts from machine learning and computer vision. We also cover Long Short-Term Memory (LSTM) networks, crucial for processing sequential exercise data, and computer vision, which enables the system to interpret visual inputs to assess and correct posture. Technologies like MediaPipe play a central role in motion perception, enabling the accurate real-time analysis required for our approach to enhancing workout safety and effectiveness in Lima's gyms.

Definition 1 (Long Short-Term Memory (LSTM) (Bairaktaris and Levy, 1993))**.** *The Long Short-Term Memory (LSTM) model in machine learning is a recurrent neural network architecture specifically designed to address the vanishing gradient problem that affects standard networks.*

This model has the ability to learn long-term dependencies in data due to its unique structure, which includes input, output, and forget gates.

Example 1. *Fig. 1 shows the internal workings of an LSTM cell, highlighting the flow of information through the forget, input, and output gates, along with the cell state and hidden state transitions over time.*

Definition 2 (Computer Vision (Gionfrida et al., 2024))**.** *Computer vision is a field of artificial intelligence that focuses on enabling computers to understand visual information from images or videos by developing algorithms to extract relevant patterns.*

Applications of this technology range from image classification to object detection, recognition, and semantic segmentation (Gionfrida et al., 2024).

Example 2. *As shown in Fig. 2, the computer vision system is structured into acquisition, processing, and visualization modules, which work together to detect and classify visual data efficiently.*

Definition 3 (MediaPipe (Lugaresi et al., 2019))**.** *MediaPipe is an open-source framework designed for building and running perception pipelines.*

It provides an efficient platform for real-time processing of visual data, such as video and audio, with compatibility across multiple devices.

Example 3. *Fig. 3 illustrates the key body landmarks detected by MediaPipe, which are used for pose estimation and motion analysis in our system.*

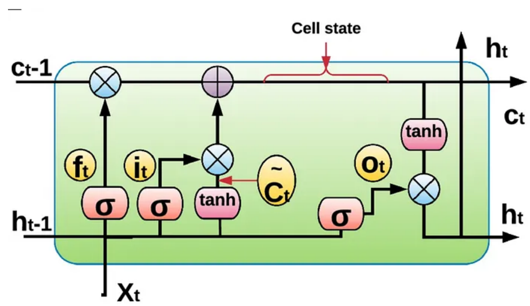

Figure 1: Key components of LSTM (Ghojogh and Ghodsi, 2023).

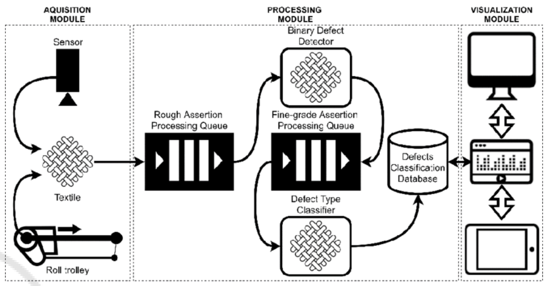

Figure 2: Architecture of the computer vision (Adão et al., 2022).

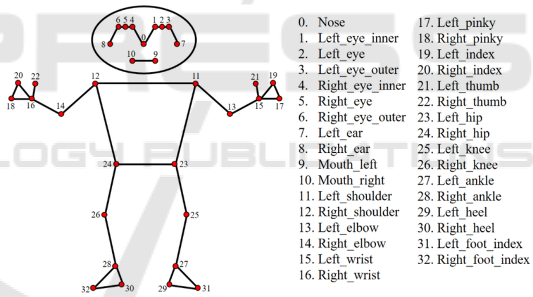

Figure 3: Key body landmarks detected by MediaPipe (Chen et al., 2022).

3.2 Method

Now, we detail the main methods of our proposal, based on web development and machine learning techniques for pose detection while exercising.

3.2.1 Physical Architecture

The physical architecture of the Gym Pose mobile application is designed to ensure the scalability, security, and efficiency of the system. This architecture deploys the different components of the system on specific infrastructures: the backend and database are hosted on Digital Ocean, while the Machine Learning microservice runs on Google Cloud. The backend, developed with NestJS, manages business logic, user authentication, and communication with the MySQL database, where critical data such as users, exercises, goals, and precision records are stored. Meanwhile, the Machine Learning microservice, implemented in Python, processes exercise videos uploaded by users using MediaPipe and LSTM models, returning a precision percentage.

Ionic Framework: Ionic is an open-source UI toolkit for building cross-platform mobile, web, and desktop applications, enabling developers to create applications using web technologies like HTML,

Figure 4: Ionic Architecture[1].

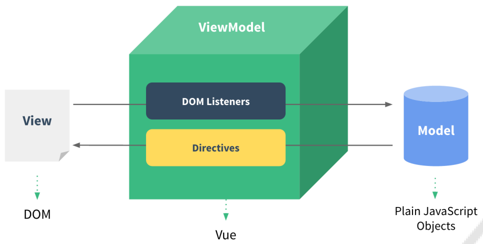

Figure 5: Vue Concepts[2].

CSS, and JavaScript. Additionally, it provides a set of pre-designed UI components that make it easier to build interactive and high-performance user interfaces, making it an efficient option for mobile application development [3]. Fig. 4 shows the architecture of the Ionic framework, which integrates web technologies, UI controls, native access through Capacitor, and multiple distribution platforms.

Vue.js: Vue.js is a progressive framework for building user interfaces, known for its simplicity and ease of integration with other projects. It helps us efficiently manage the front-end components of the application, ensuring optimal performance and scalability for our app's user interface [4]. Fig. 5 illustrates the Vue.js architecture, where the ViewModel manages the interaction between the View (DOM) and the Model (JavaScript objects), using directives and DOM listeners to synchronize data efficiently. The mobile application, built with Ionic and Vue.js, serves as the primary user interaction point, allowing video uploads and results viewing (see Fig. 6). Distributed through the Play Store and App Store, it ensures accessibility across a wide range of Android and iOS devices, providing a seamless and secure experience for users. This architecture not only distributes the workload but also ensures that the system can scale efficiently to handle an increasing number of users and videos without compromising performance.

[1]M. Lynch, "Announcing Capacitor 1.0," Ionic Blog, Oct. 16, 2020. https://ionic.io/blog/announcing-capacitor-1-0

[2]Getting started - Vue.js." https://012.vuejs.org/guide/

[3]The Ionic Platform - Ionic Documentation. - https://ionic.io/docs/platform

[4]The Progressive JavaScript Framework - Vue.js. - https://vuejs.org/

3.2.2 Logical Architecture

The logical architecture of Gym Pose is organized into layers, providing a clear separation of responsibilities that facilitates system maintenance, security, and scalability (see Fig. 7). The presentation layer consists of the mobile application, which offers an intuitive and accessible interface for users to interact with the system, upload videos, set goals, and view their progress. The business services layer includes the backend, which acts as an intermediary between the mobile application and the data and processing services. This layer handles user authentication, exercise and goal management, and ensures secure communication with the database and the Machine Learning microservice. Finally, the data layer manages the storage of all user-generated information, from personal settings to records of their exercises and goals. This logical architecture allows the various components of the system to operate in a coordinated manner, ensuring that data flows correctly and that each user request is handled efficiently and securely. This structure ensures that Gym Pose can deliver an optimized and reliable experience, promoting the continuous improvement of users' postures through precise and personalized analysis, supported by a robust and well-integrated physical and logical architecture.

3.3 Machine Learning Model Flow Diagram

In Fig. 8, the diagram represents the flow of the Machine Learning model used in the Gym Pose mobile application, highlighting each step from video input to posture evaluation score output. This flow is essential to understanding how the system processes user videos and assesses exercise posture, adding significant value to the user experience. The process begins with a user-uploaded video, recorded directly from the mobile application. MediaPipe analyzes the video to estimate the user's 2D body pose, identifying key points that create a virtual skeleton. The coordinates of key body parts are extracted, capturing the specific positions of limbs such as shoulders, elbows, and knees. The LSTM model, designed to handle temporal sequences, processes the extracted coordinates to evaluate the posture. The model outputs a score reflecting the accuracy of the exercise performed.

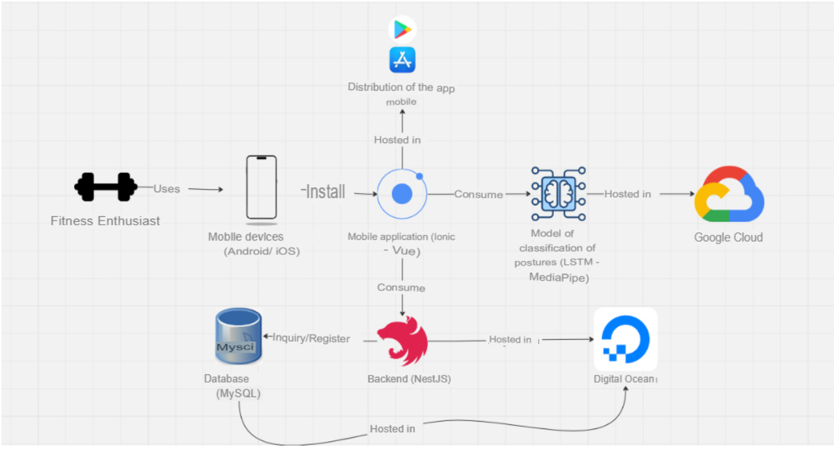

Figure 6: Physical Architecture.

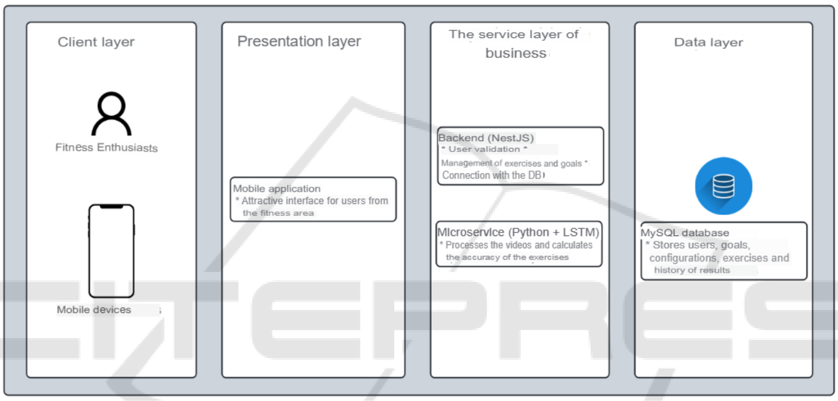

Figure 7: Logical Architecture.

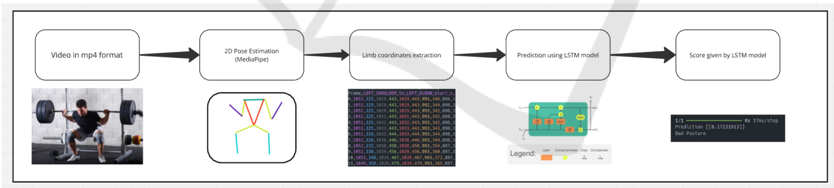

Figure 8: Machine Learning Model Flow Diagram.

4 EXPERIMENTS

4.1 Experimental Protocol

In this subsection, the setup required to develop and evaluate of our proposal is detailed. We have two main components: the machine learning model for posture classification and the mobile application. The machine learning model was developed and trained on a laptop with the following specifications: Arch Linux x86_64, Intel i7-10750H (12) @ 5.000GHz, NVIDIA GeForce GTX 1650 Mobile / Max-Q and 32GB RAM @ 2700MHz. The dataset used for training the machine learning model consists of videos of individuals performing squats, sourced from the following dataset: https://hi.cs.waseda.ac.jp/~ogata/Dataset.html.

The mobile application was developed on a PC with the following specifications: Windows 11, Intel i5 10400F, NVIDIA RTX 2060 and 32GB RAM @ 3200MHz. The mobile application was built using Ionic and Vue 3, using TypeScript for front-end development. The backend was developed with NestJS and Prisma, with dependencies managed through Node.js. All the source code for is available at https://github.com/orgs/P20242083-GymPose/repositories.

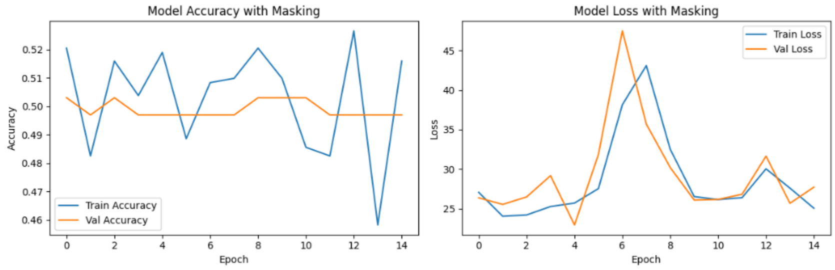

(a) First iteration.

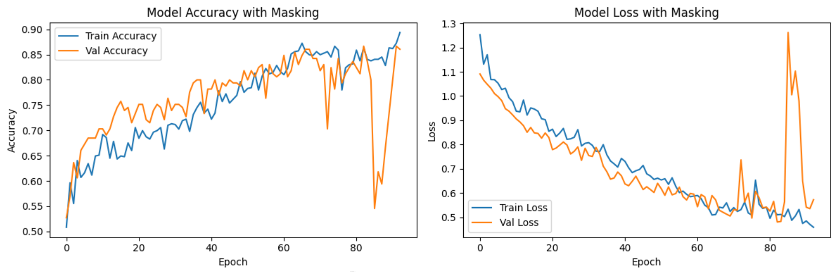

(b) Final iteration.

Figure 9: Models' Accuracy and Loss.

4.2 Results

In this section, we present the results of training our machine learning model on the ”Maseda Squats Dataset” to optimize exercise posture detection. By leveraging data from diverse workout scenarios, varied lighting conditions, and multiple poses, our model achieved high accuracy in identifying and analyzing key body positions across different exercise repetitions. This precision enables the system to reliably fetch scores for posture quality, ensuring accurate and context-sensitive feedback for users. These results highlight the model's robustness and adaptability, underscoring its potential for real-world application in gym environments. The training iteration results for the model were as follows:

Fig. 9a shows the results of the first training iteration. The model, configured with an LSTM architecture, applied masking for padded values, L2 regularization, batch normalization, and dropout layers to enhance stability. However, the results indicate significant issues in learning and generalization, as evidenced by erratic fluctuations in accuracy and a validation accuracy plateauing around 50%. These trends suggest underfitting, highlighted by a low test accuracy of 47.41% and an F1 score of 0.00. This iteration exposed the need for further refinements in the model architecture and hyperparameter tuning.

Fig. 9b presents the results of the final iteration, showcasing the significant improvements achieved after optimizing the model architecture and hyperparameters. The model displayed stable and steady learning, with accuracy reaching between .85 and .90. Both the training and validation loss curves show consistent decreases, with minor fluctuations in validation loss, suggesting effective learning and minimal overfitting. The final configuration, with a reduced dropout rate of .4 and L2 regularization adjusted to 0.002, resulted in a robust test accuracy of 87% and an F1 score of 0.87. This demonstrates the model's capacity to generalize well across different classes, achieving near-optimal performance for this task. This final training session demonstrates significant model improvement, with steady learning and generalization due to updated architecture and hyperparameters. The accuracy curve reaches .85 to .90, indicating the model effectively learns patterns, while the training and validation loss curves decrease consistently, showing stable learning with minor fluctuation in validation loss. This updated configuration lowered dropout to 0.4 and adjusted L2 regularization to 0.002, enhancing generalization without overfitting. The final test accuracy of 87% and F1 score of 0.87 indicate balanced performance across classes, and the best model was saved at epoch 83. These adjustments, alongside the stable learning rate of 0.0001, make this setup highly effective and close to optimal for this task.

5 CONCLUSIONS

In conclusion, this study contributes a meaningful tool to the fitness industry, offering an accessible and effective means of posture correction for gym enthusiasts. Iterative improvements in model accuracy and stability reinforce the model's practical applicability, while the final results demonstrate a reliable solution for exercise optimization. The potential impact on reducing injuries and enhancing exercise efficacy positions this application as a valuable asset for individuals and fitness institutions aiming to foster safer and more effective workout environments.

The application of LSTM networks for sequential data processing has proven effective in handling the complex and dynamic nature of gym exercises. Initial training iterations revealed challenges related to model accuracy and stability, including fluctuations and underfitting. However, by refining the model architecture—using techniques such as L2 regularization, dropout adjustments, and lowering the learning rate—subsequent iterations showed marked improvements. The final model achieved a test accuracy of 87% and an F1 score of 0.87, reflecting robust learning and effective generalization.

While the model performed well in posture analysis, the reliance on 2D pose estimation limits its ability to fully capture depth-related details in complex movements. This limitation may affect feedback accuracy in exercises that involve multiple joint movements. (Lozano-Mejía et al., 2020) The current model's performance could benefit from a more diverse dataset that includes a wider range of body types, exercise intensities, and environments. (Cornejo et al., 2021) Expanding the dataset would enhance the model's generalization across various user demographics and workout conditions, contributing to more consistent feedback accuracy. (Ysique-Neciosup et al., 2022)

Future research could focus on integrating 3D pose estimation and conducting longitudinal studies to evaluate the application's long-term impact on users' exercise habits, injury rates, and performance improvements. Additionally, implementing personalized feedback based on user-specific goals could further tailor the fitness experience, making it more engaging and effective.

REFERENCES

Adão, T., Gonzalez, D., Castilla, Y. C., Pérez, J., Shahrabadi, S., Sousa, N., Guevara, M., and Magalhães, L. G. (2022). Using deep learning to detect the presence/absence of defects on leather: on the way to build an industry-driven approach. *Journal of Physics: Conference Series*, 2224(1):012009.

Bairaktaris, D. and Levy, J. (1993). Using old memories to store new ones. In *IJCNN*, volume 2, pages 1163–1166 vol.2.

Chen, K., Shin, J., Hasan, M. A. M., Liaw, J., Okuyama, Y., and Tomioka, Y. (2022). Fitness movement types and completeness detection using a transfer-learning-based deep neural network. *Sensors*, 22(15):5700.

Cornejo, L., Urbano, R., and Ugarte, W. (2021). Mobile application for controlling a healthy diet in peru using image recognition. In *FRUCT*, pages 32–41. IEEE.

Ghojogh, B. and Ghodsi, A. (2023). Recurrent neural networks and long short-term memory networks: Tutorial and survey. *CoRR*, abs/2304.11461.

Gionfrida, L., Wang, C., Gan, L., Chli, M., and Carlone, L. (2024). Computer and robot vision: Past, present, and future [TC spotlight]. *IEEE Robotics Autom. Mag.*, 31(3):211–215.

Kaewrat, C., Khundam, C., and Thu, M. (2024). Enhancing exercise monitoring and guidance through mobile augmented reality: A comparative study of RGB and lidar. *IEEE Access*, 12:95447–95460.

Kemler, E., Noteboom, L., and van Beijsterveldt, A.-M. (2022). Characteristics of fitness-related injuries in the netherlands: A descriptive epidemiological study. *Sports*, 10(12).

Lozano-Mejía, D. J., Vega-Uribe, E. P., and Ugarte, W. (2020). Content-based image classification for sheet music books recognition. In *EIRCON*, pages 1–4. IEEE.

Lugaresi, C., Tang, J., Nash, H., McClanahan, C., Uboweja, E., Hays, M., Zhang, F., Chang, C., Yong, M. G., Lee, J., Chang, W., Hua, W., Georg, M., and Grundmann, M. (2019). Mediapipe: A framework for building perception pipelines. *CoRR*, abs/1906.08172.

Mallick, T., Das, P. P., and Majumdar, A. K. (2022). Posture and sequence recognition for *Bharatanatyam* dance performances using machine learning approaches. *J. Vis. Commun. Image Represent.*, 87:103548.

Müller, P. N., Müller, A. J., Achenbach, P., and Göbel, S. (2024). Imu-based fitness activity recognition using cnns for time series classification. *Sensors*, 24(3):742.

Palanimeera, J. and Ponmozhi, K. (2023). Yap_lstm: yoga asana prediction using pose estimation and long short-term memory. *Soft Computing*.

Vali, D., Venkata Chalapathi, M., Yellapragada, V., Purna Prakash, K., Challa, P., Gangishetty, D., Solanki, M., and Singhu, R. (2024). Physiotherapy assistance for patients using human pose estimation with raspberry pi. *ASEAN Journal of Scientific and Technological Reports*, 27:e251096.

Ysique-Neciosup, J., Chavez, N. M., and Ugarte, W. (2022). Deephistory: A convolutional neural network for automatic animation of museum paintings. *Comput. Animat. Virtual Worlds*, 33(5).