Prompting Large Language Models for Tailored Exercise Recommendations in Office Spaces

Gaetano Dibenedetto[1,*], Marco Polignano[1], Pasquale Lops[1] and Giovanni Semeraro[1]

[1]University of Bari Aldo Moro, Via Orabona 4, Bari, 70125, Italy

Abstract

In the digital era, *Recommender Systems* are a crucial component, commonly used in services such as music and movie streaming. Despite their widespread adoption, surprisingly little attention has been devoted to developing systems that can positively impact users' *well-being* and *health*. In an effort to combat the negative effects of a sedentary way of life on people's health and the subsequent rise in healthcare expenses, we introduce an encouraging approach, i.e. a recommender system that, through webcam-based monitoring of subject postures, suggests *personalized exercise breaks* to do directly near users' desks.

Our system captures users' postures during work hours and employs 3D pose estimation to calculate key angles between shoulders, hips, and head. By identifying postural imbalances, we generate exercise recommendations using a *Large Language Model* (LLM). The system flags potential postural issues when angle thresholds are exceeded and prompts the LLM to provide tailored exercise suggestions. Our method's effectiveness is assessed by experts in the field. While the results are still preliminary, our approach deserves further investigation, with future developments likely to focus on enriching the data and refining the detection methods. The full-reproducible code is available at the following link: https://github.com/GaetanoDibenedetto/healthrecsys24

Keywords

Health Recommender Systems, 3D Pose Estimation, Explainability, Large Language Model

1. Introduction

Extended periods of sitting have become unremarkable in modern lifestyles, but recent research highlights the detrimental effects it can have on our health. Studies have shown that prolonged sedentary behavior is linked to a higher risk of obesity and metabolic syndrome, characterized by elevated blood pressure, high blood sugar, excessive abdominal fat, and unhealthy cholesterol levels [1]. Furthermore, investigations by Marras et al. [2] suggest that prolonged static sitting postures may compromise the nutrition of intervertebral discs. More severe outcomes have also emerged from further investigation. Dunstan et al. [3] observed a strong relation between sedentariety and premature mortality. Hamilton et al. [4] found relations with chronic illnesses, and Inoue et al. [5] with obesity. The implications of maintaining inadequate sitting postures over prolonged periods are indeed alarming. Conditions such as cervical spondylosis, lumbar diseases, and other ailments commonly known as "chair diseases" have been linked to poor sitting habits [6, 7, 8]. These musculoskeletal disorders can have a significant impact on an individual's quality of life, leading to chronic pain, reduced mobility, and decreased productivity.

Building on our previous work focused on posture correction for office workers [9], this research aims to extend that approach by developing a recommender system that suggests targeted physical activities to improve posture. We recognize that simply offering posture correction and, for this work, physical activity recommendations, does not guarantee that users will follow the system feedback aimed to help themselves. Hence, our goal, starting with this preliminary work, is to provide expert-level advice in a convenient and accessible format. By delivering daily, personalized recommendations, we

HealthRecSys'24: The 6th Workshop on Health Recommender Systems co-located with ACM RecSys 2024

*Corresponding author.
gaetano.dibenedetto@uniba.it (G. Dibenedetto); marco.polignano@uniba.it (M. Polignano); pasquale.lops@uniba.it (P. Lops); giovanni.semeraro@uniba.it (G. Semeraro)
0000-0001-6083-3600 (G. Dibenedetto); 0000-0002-3939-0136 (M. Polignano); 0000-0002-6866-9451 (P. Lops); 0000-0001-6883-1853 (G. Semeraro)

hope to enhance user engagement and motivation, ultimately leading to gradual improvements in posture and overall wellness. In this context, the role of technology in promoting physical activity and exercise becomes even more critical, particularly in office environments where sedentary behavior is prevalent. To address these issues, we propose an approach that leverages webcam-based posture analysis and 3D pose estimation to monitor users' postures during working hours. By analyzing key angles between the shoulders, hips, and head, our system detects postural imbalances and provides personalized exercise recommendations, generated through a Large Language Model (LLM).

The main contributions of this work are summarized below:

- **3D Keypoint Extraction Model**: we propose a 3D Human Pose Estimation (HPE) keypoint extraction module able to detect the human pose from a single laptop camera;
- **Pose Classification Model**: we develop a system able to detect wrong poses while sitting at a desk efficiently;
- **Angles Computation**: we develop a vector-based angles computation module that calculates the angles between key body parts (shoulders, hips, and head) using 3D pose data. This helps detect postural imbalances by identifying deviations from normal alignment;
- **Recommendation Generation Module**: a module that utilizes a LLM to generate personalized exercise recommendations based on detected posture issues. When key angle thresholds are exceeded, the system prompts the LLM to suggest corrective exercises tailored to the user's specific posture;
- **Dataset publicly released**: the dataset with the 3D keypoints extracted from frames, used for training and evaluating our model will be publicly released for replicability purposes on Zenodo[1].

2. RELATED WORK

2.1. Posture Classification

As highlighted in our previous work [9], several studies share our goal of posture classification [10, 11, 12]. However, many of these approaches rely on data collected under strict constraints, e.g., users are often required to position themselves directly in front of the camera or assume predefined sitting postures; are required specialized accessories that are impractical and expensive such as multiple cameras or sensors installed in chairs.

These limitations make it difficult to apply such systems in a typical daily routine. In our previous work, we introduced a posture classification approach tailored for office workers, which avoids the constraints imposed by related works. Specifically, we allowed users to record themselves while working, without imposing restrictions. In that study, we collected our own dataset and developed a system capable of posture classification using a multi-layer perceptron. The system offered explanations based on data statistics and featured a personalized feedback module designed to correct user posture. Building on this, our current research aims to **enhance the system by leveraging 3D HPE** while maintaining the same dataset and classification technique. This approach shifts the focus towards creating a recommendation system to provide personalized exercise suggestions.

2.2. Physical Activity Recommendation

In today's information-overloaded world, recommender systems have emerged as essential tools for navigating countless options. Whether we're shopping, streaming, or socializing, these intelligent systems are constantly working behind the scenes to personalize our experiences. Generally, these systems work on a daily basis based on a user's previous behavior in a system.

[1] https://zenodo.org/records/13498794

As discussed in [13, 14], mostly health recommender systems aim to improve the general well-being of users, such as recommending diets and exercise plans. However, developing recommender systems specifically for physical activity has proven challenging due to practical limitations. In particular, the scarcity of ratings for exercise activities, a cornerstone of traditional recommender systems. This is largely attributed to the difficulty in quantifying individual interest in specific exercises. Therefore, in our search for closely related works, we came across a limited number of works similar to the one we have proposed. We will discuss them below.

A Physical Activity Recommender aims to recommend a daily routine of physical activities and workouts to the user, this could be based on data as user's characteristics, his health status or other demographic information such as age and gender. One of the first work analyzed by us is RecFit [15], a context-aware recommender system, which systematically suggests physical activities based on the user's context, e.g. risk tolerance, budget, location, weather, but not on user's physical characteristics. However, it does not consider the user's physical characteristics and is not specifically designed to address health concerns. Advancements in the field have incorporated more physical details into the recommendation process. For instance, in 2020, Ferretto et al. [16] developed a system for patients with arterial hypertension, generating personalized recommendations based on the patient's age, gender, and physical condition. In 2021, Sengan et al. [17] proposed a system to prevent respiratory diseases, utilizing data on physical activity levels, heart rate, and respiratory rate. The most recent work we reviewed, DEEP-CARDIO [18], proposed in 2024, is a content-based system offering physical and dietary recommendations for cardiac patients. Using Density-Based Clustering (DBSCAN), it categorizes patient data, such as class predictions for cardiovascular diseases, alongside age and gender. However, the recommendations are presented in a categorical format, which may be challenging for non-expert users to interpret, e.g., instructions like "Follow Type 1 diet plan. Exercise Tip: Walking stationary, cycling, rowing, or water aerobics".

We observed that while these systems produce positive outcomes for patients, **they often present categorical outputs**, which can be **difficult for non-expert users to interpret**. In contrast, our goal is to develop a system that not only delivers specific recommendations but also includes brief descriptions. This approach will help non-expert users easily understand the recommendations and the rationale behind them.

2.3. Recommendation through LLMs

In recent years, following the introduction of Transformers [19], Large Language Models (LLMs) based on this architecture have been widely adopted across various machine-learning applications. Transformers have gained popularity primarily due to their attention mechanism, which improves the model's ability to represent text by focusing on the semantics of surrounding words.

Similar works have been found with our aim in the area of recommendations generated by LLMs, but only based on movies or books, which are the most common applications in the recommendation area, thanks to the large availability of datasets. W.-C. Kang et al. [20] propose one of the studies that we explored, which analyzed multiple LLMs across different settings, i.e., zero-shot, few-shot, and fine-tuning, by feeding the LLM with a prompt representing the user profile based on their past item ratings. Another work proposed by Sanner et al. [21], focuses on sequential recommendations by incorporating item descriptions and user preferences. While these approaches, like ours, leverage LLMs for generating recommendations through prompting, **these differs from our approach, which does not rely on an item-based dataset**, due to the lack of data in our scenario. Instead, our system's core lies in the *Angles Computation module* (Sec. 3.4), which detects anomalies in the user's posture. These anomalies are prompt to the system to generate personalized physical exercise recommendations aimed at improving the user's well-being.

3. PROPOSED APPROACH

3.1. Data Collection/Dataset

Up to our knowledge, there is a noticeable gap in the literature related to the availability of a specific public dataset within the domain of sitting poses. In order to bridge this gap, in our previous research, we created our own dataset, where workers in office spaces have been recorded while sitting. No constraints have been fixed in terms of camera angles or perspectives, minimum or maximum video lengths, or number of videos. We gathered videos from 10 subjects, processed them to extract frames and finally, they were annotated by humans. Some statistics of the dataset are reported in Tab. 1, all the details about, collection, filtering and processing of data are already provided in [9].

Table 1
Dataset Data statistic

Subject	Frames	INcorrect	Correct	correct%
ap	500	411	89	18%
gc	35	30	5	14%
gd	174	111	63	36%
gs	145	139	6	4%
k	49	49	0	0%
mc	31	30	1	3%
ms	47	42	5	11%
nl	44	41	3	7%
rt	23	23	0	0%
vr	99	99	0	0%
TOTAL	1147	975	172	15%

3.2. Keypoint Extraction Module

Building on our previous 2D-based approach [9], we shift to a 3D environment. This transition is made possible using MotionBERT [22], a 3D HPE model, implemented via the MMPose framework [23]. MotionBERT operates through 2D-to-3D lifting, where, during the pretraining phase, a motion encoder learns to reconstruct the underlying 3D motion from incomplete 2D observations by integrating geometric, kinematic, and physical insights about human movement. We selected MotionBERT over other HPE models due to its performance, since it achieves the lowest 3D pose estimation error to date on the Human3.6M dataset [24]. An example of the Keypoint Exraction Module is shown in Fig. 2b. The keypoints extracted on our Dataset are available on Zenodo[1] to guarantee reproducibility.

3.3. Pose Classification Model

To guarantee consistency, we used the same architecture with same hyperparameters proposed with our 2D-based approach We used a Multilayer Perceptron architecture (Fig. 1), Adam optimiser and the BCE (Binary Cross Entropy) loss function, weighted with the class distribution in the training fraction. Training was conducted with a batch size of 20 on 1000 epochs with a learning rate of 0.001. The high number of training epochs is due to a slow learning process due to the high variability of the keypoints position. The best performing model based on the training loss was selected and saved to be used in the recommendation step.

The training phase of this architecture uses the keypoints extracted with MotionBERT previously described (Sec. 3.2). A data augmentation step is performed, considering the imbalance of the dataset, inserting the vertical flip of the correct postures. It is worth to note that only two subjects exhibited a quite balanced amount of correct posture data, leading to the consideration of a binary split: training and testing. Among these two subjects, the one with a more balance volume of data was designated as

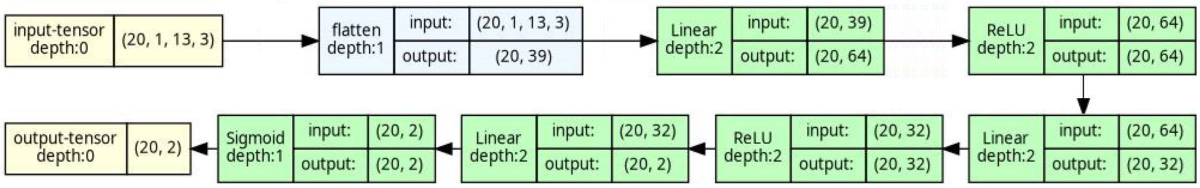

Figure 1: MLP architecture

the sole subject for the test split, i.e. *gd*, while the remaining subjects were allocated to the training split. The results of the classification are shown in the Tab. 2.

Table 2
Classification Report

	precision	recall	f1-score	support
INcorrect Posture	0.9135	0.8559	0.8837	111
Correct Posture	0.8797	0.9286	0.9035	126
accuracy			0.8945	237
macro avg	0.8966	0.8922	0.8936	237
weighted avg	0.8955	0.8945	0.8942	237

3.4. Angles Computation

We focus on computing specific angles related to the user's posture, which are crucial for detecting imbalances and misalignments. The primary angles of interest are the inclination of the shoulders and hips, as well as the tilt of the head in both sideways and forward directions.

To compute these angles, we employ vector-based calculations. Using the 3D skeleton representation provided by MotionBERT, the position of each body part is defined as a point in 3D coordinate space. The vectors formed between these points allow us to determine their inclination relative to the coordinate axes. The angles are then derived using the dot product between these vectors and a reference axis, as detailed below:

Given two points representing different body parts, $P_1(x_1, y_1, z_1)$ and $P_2(x_2, y_2, z_2)$, we compute the vector between these points as $\mathbf{v} = P_2 - P_1$. The general formula for computing the angle θ between two vectors $\mathbf{A}$ and $\mathbf{B}$ is:

$$\theta = \cos^{-1}\left(\frac{\mathbf{A} \cdot \mathbf{B}}{|\mathbf{A}||\mathbf{B}|}\right)$$

where $\mathbf{A} \cdot \mathbf{B}$ is the dot product, and $|\mathbf{A}|$ and $|\mathbf{B}|$ represent the magnitudes of the respective vectors. The direction of rotation can be inferred using the cross product of the two vectors.

In our case, vector $\mathbf{A}$ corresponds to $\mathbf{v}$, the vector connecting two body parts, and vector $\mathbf{B}$ is either the x-axis $[1, 0, 0]$ or the z-axis $[0, 0, 1]$, depending on the body part and the direction of interest.

- **Shoulder Inclination**: The shoulder inclination is computed relative to the x-axis. The shoulders are represented by a vector connecting the left and right shoulder points. Using the dot product, we calculate the angle between this vector and the horizontal axis to assess any unevenness in shoulder alignment.

- **Hip Inclination**: Similarly, hip inclination is measured with respect to the x-axis. A vector formed between the left and right hip points is analyzed to detect imbalances in the lower body, providing insight into hip misalignment.

- **Head Tilt**: To evaluate head posture, two distinct angles are computed:

- Sideways Tilt: This angle is measured relative to the x-axis and provides information on lateral imbalances in head and neck posture.
- Forward Tilt: The angle of forward head tilt is computed with respect to the z-axis, allowing us to assess how far the head is leaning forward compared to a neutral, upright position.

An example of angle detection is shown in Fig. 2.

By analyzing these angles using vector-based methods and dot products, we aim to detect deviations from normal alignment that could indicate issues with posture. Unfortunately, we were unable to establish specific angle thresholds to definitively define "anomaly" for various body parts, also with the support of experts in the field. To address this, we propose a data-driven approach. After calculating all angles in our dataset, we will determine the average angle for each specific body part. This average will serve as a threshold to flag any angles that deviate significantly from the norm. This threshold-based approach will serve as the foundation for generating personalized recommendations in the subsequent stage.

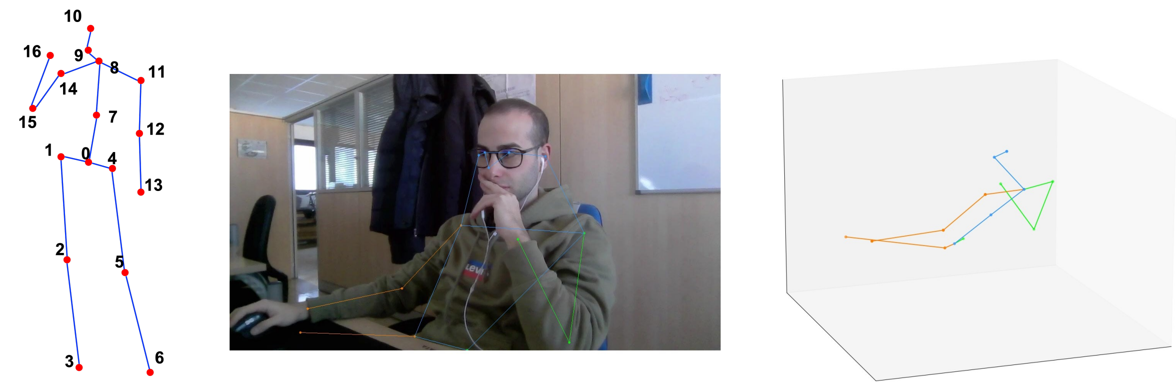

(a) Human3.6M keypoints representation

(b) MMPose Extraction Module with the 3D HPE MotionBERT

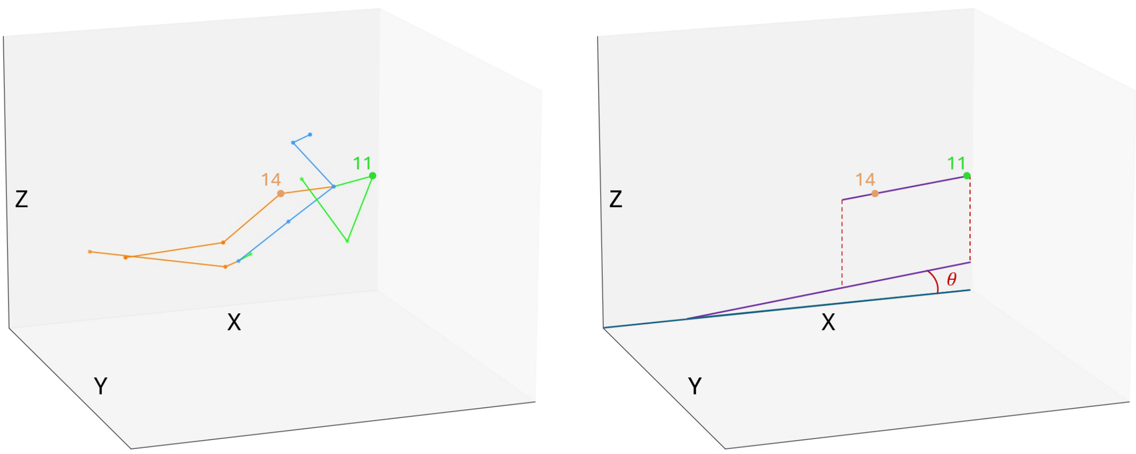

(c) Identify the 3D points of the shoulders

(d) Example of the angle θ for the shoulder angle

Figure 2: The figure illustrates an example of keypoints extracted from a frame, showcasing the computation of the shoulder inclination angle θ relative to the x-axis. Subfigure (a) shows the representation of the keypoints in the HPE model used; Subfigure (b) presents the output of the Keypoint Extraction Module (Sec. 3.2); Subfigure (c) highlights the 3D pose with a focus on the shoulder keypoints; Subfigure (d) details the computation of the shoulder inclination angle θ (Sec. 3.4)

3.5. Recommendation Generation

The angles, obtained as described previously, are used to detect anomalies in the user's posture. These detected anomalies are inserted in a textual prompt to be used as input of an LLM. The prompt starts with the sentence *"I've an incorrect posture caused by"*, than each anomaly is joint to the prompt with an "and" (except the first), and it ends with the sentence *", what exercise or stretching should I do?"*. If no anomalies are detected the prompt will be *"If I had an incorrect posture what exercise or stretching should I do?"*. The system prompt provided to the model is *"You are a medical assistant able to suggest a specific physical exercise for users with uneven body parts to prevent an increase in their unease"*. An example of the complete prompt provided to the LLM is shown in Listing 1.

Listing 1: LLM prompt example

```
<|begin_of_text|><|start_header_id|>system<|end_header_id|>

You are a medical assistant able to suggest a specific phisical exercise for users
    with uneven body parts to prevent an increase in their unease.<|eot_id|><|
    start_header_id|>user<|end_header_id|>

I've an incorrect posture caused by Shoulder Inclination, what exercise or
    stretching should I do?<|eot_id|><|start_header_id|>assistant<|end_header_id|>
```

Our recommendation generation process is powered by *Llama-3.1-8B-Instruct* [25]. We deploy the model on a machine using 2 NVIDIA A16 16GB VRAM GPUs, which required from 50 to 70 seconds for the text generation. Using the transformers library[2], we utilize the *"text-generation"* pipeline[3] to generate the recommendations. The default hyperparameters from the pipeline are used, i.e. temperature=0.6, top-p=0.9, instead we set for the text generation length the *"max_new_tokens=1024"*. An output example of the recommendation is shown in the Fig. 3.

4. EXPERIMENTAL RESULTS

Due to the limited availability of data, we were unable to conduct a quantitative evaluation. Instead, we conducted a preliminary evaluation by administering questionnaires to domain experts, with at least 5 years of background experience in the field, specifically two physiotherapists. While the number of enrolled experts was small, their insights were crucial for understanding the pipeline's potential. They were asked to assess 10 different scenarios in which the user is seated in an incorrect posture. These scenarios were either manually selected by the evaluators from provided example images or uploaded by them using specific frames they found relevant. To facilitate this evaluation, we developed a website using a Gradio[4] interface, allowing users to interact directly with our system. The interface presents a chat-based interaction where users can either upload their own image or select one from a preloaded list of examples. Upon submission of an image, the system processes it through the keypoint extraction module and displays the results in the chat. The system also identifies any posture anomalies, providing the respective inclination angles for each anomaly detected. Following this, the system generates personalized exercise recommendations based on the detected anomalies. After reviewing the recommendations, users are prompted to answer a set of five questions on a 1 to 5 scale, followed by an open-ended feedback question. An example of interaction with the interface is shown in Fig. 3. The expert feedback collected through this process forms the basis of our evaluation. The questionnaire provided to the experts consists of the following questions:

- **Q1** - *Do you agree with the recommended exercises based on the user's posture analysis? Rating: 1 (Totally Disagree) to 5 (Totally Agree)*

[2]https://huggingface.co/docs/transformers/index
[3]https://huggingface.co/meta-llama/Meta-Llama-3.1-8B-Instruct
[4]www.gradio.app

- **Q2** - *Do the recommended exercises adequately address the user's posture issue and imbalances? Rating: 1 (Totally Inadequate) to 5 (Totally Adequate)*
- **Q3** - *Are the exercises clearly explained and easy to understand? Rating: 1 (Very Unclear) to 5 (Very Clear)*
- **Q4** - *In your opinion, how effective would these exercises be in improving the user's posture? Rating: 1 (Not Effective at All) to 5 (Highly Effective)*
- **Q5** - *Are the recommended exercises simple enough for the user to perform without expert supervision? Rating: 1 (Not Simple at All) to 5 (Very Simple)*
- **Q6** - *What improvements would you suggest for this recommendation?*

With these questions, we evaluate the correctness (Q1, Q2, Q4), accuracy(Q1), adequacy(Q2), clarity(Q3), effectiveness(Q4) and the simplicity of the execution (Q5) of the recommendation. The average rating for each question, on a scale from 1 to 5, is presented in Tab. 3. Overall, while none of the average results were negative (i.e., below 3), they were also not exceptionally high, except for Q3, which exceeded a rating of 4. This indicates that the clarity of the exercise recommendations was appreciated, but there is significant room for improvement in other areas. In particular, the average score for the "correctness" criterion was 3.4. After discussions with the evaluators, it became clear that, also if they appreciated the concept behind the system, they noted that it was difficult to determine whether a recommended exercise is entirely appropriate or not, starting with limited user characteristics, such as a single image. They emphasized that many more variables need to be considered for an accurate recommendation.

From the expert open-feedback, particularly the insights provided by physiotherapists, we identified several areas for improvement; these are discussed below to better illustrate limitations and future work.

Table 3
Evaluation Report

Q	Evaluation Criteria	Avg Rating
1	Accuracy	3.5
2	Adequacy	3.2
3	Clarity	4.2
4	Effectiveness	3.6
5	Simplicity	3.8
1,2,4	Correctness	3.4

5. Conclusion, Limitation and Future Works

This study has introduced a system for physical activity recommendation based on assessment of sitting postures utilizing accessible cameras, vector-based angles computation, and a LLM. Despite the promising results, several limitations were identified through open feedback evaluations:

- When the system detects uneven shoulders, it lacks to recommend exercises targeting the cervical region, which plays a significant role in connecting shoulder alignment and head posture;
- When the subject is sliding down in their chair (example in Fig. 3), the system may only detect misalignment in the head's inclination relative to the body;
- The angles computation module does not recognize shoulder lifting and excessive shoulder closure;

- The system should suggest to perform exercises in front of a mirror, which could help users maintain correct form.

As future work, based on these limitations, several ways for improvement have been identified. First of all, we need to improve the angle calculation module to detect a wider range of postural anomalies, enabling more accurate and varied exercise recommendations. Additionally, there is a need to explore the use of different LLMs to evaluate their performance in this recommendation task, and last but not least, a change of module that suggests the exercises, which could be based on a collected historical user data. These data are taken from the interactions of the current system, specifically combining the recommendations and the expert feedback: by storing the data and ratings provided by experts, the system could offer better recommendations to users with similar postural issues, leveraging the best-rated advice from prior interactions. Although the experimental results presented in this paper are preliminary, we can see a great potential to improve them. The exercise descriptions generated by the LLMs were particularly appreciated for their clarity, highlighting the possibility of using LLMs to improve users' understanding. In the future, by treating past recommendations and detected physical anomalies as items of a recommender system, we can refine the system in terms of correctness, adequacy, accuracy, and effectiveness.

Acknowledgments

The research is partially funded by PNRR - Mission 4 ("Education and research") – Component 2 ("From research to business"), Investment 3.3 ("Introduction of innovative doctorates that respond to the innovation needs of companies and promote the hiring of researchers by companies") D.M.n. 117/2023 - CUP: H91I23000170007 and is supported by the co-funding of the European Union - Next Generation EU: NRRP Initiative, Mission 4, Component 2, Investment 1.3 – Partnerships extended to universities, research centers, companies, and research D.D. MUR n. 341 del 15.03.2022 – Next Generation EU (PE0000013 – "Future Artificial Intelligence Research – FAIR" - CUP: H97G22000210007).
We extend my sincere gratitude to Naps Lab S.r.l.s.[5] for their support and collaboration in the realisation of this research.
We would also like to thank Ilaria Vianello[6] and Michele Dilillo[7], who provided invaluable assistance in the evaluation phase of this research.

[5]www.napslab.it
[6]www.linkedin.com/in/ilaria-vianello
[7]www.linkedin.com/in/michele-dilillo

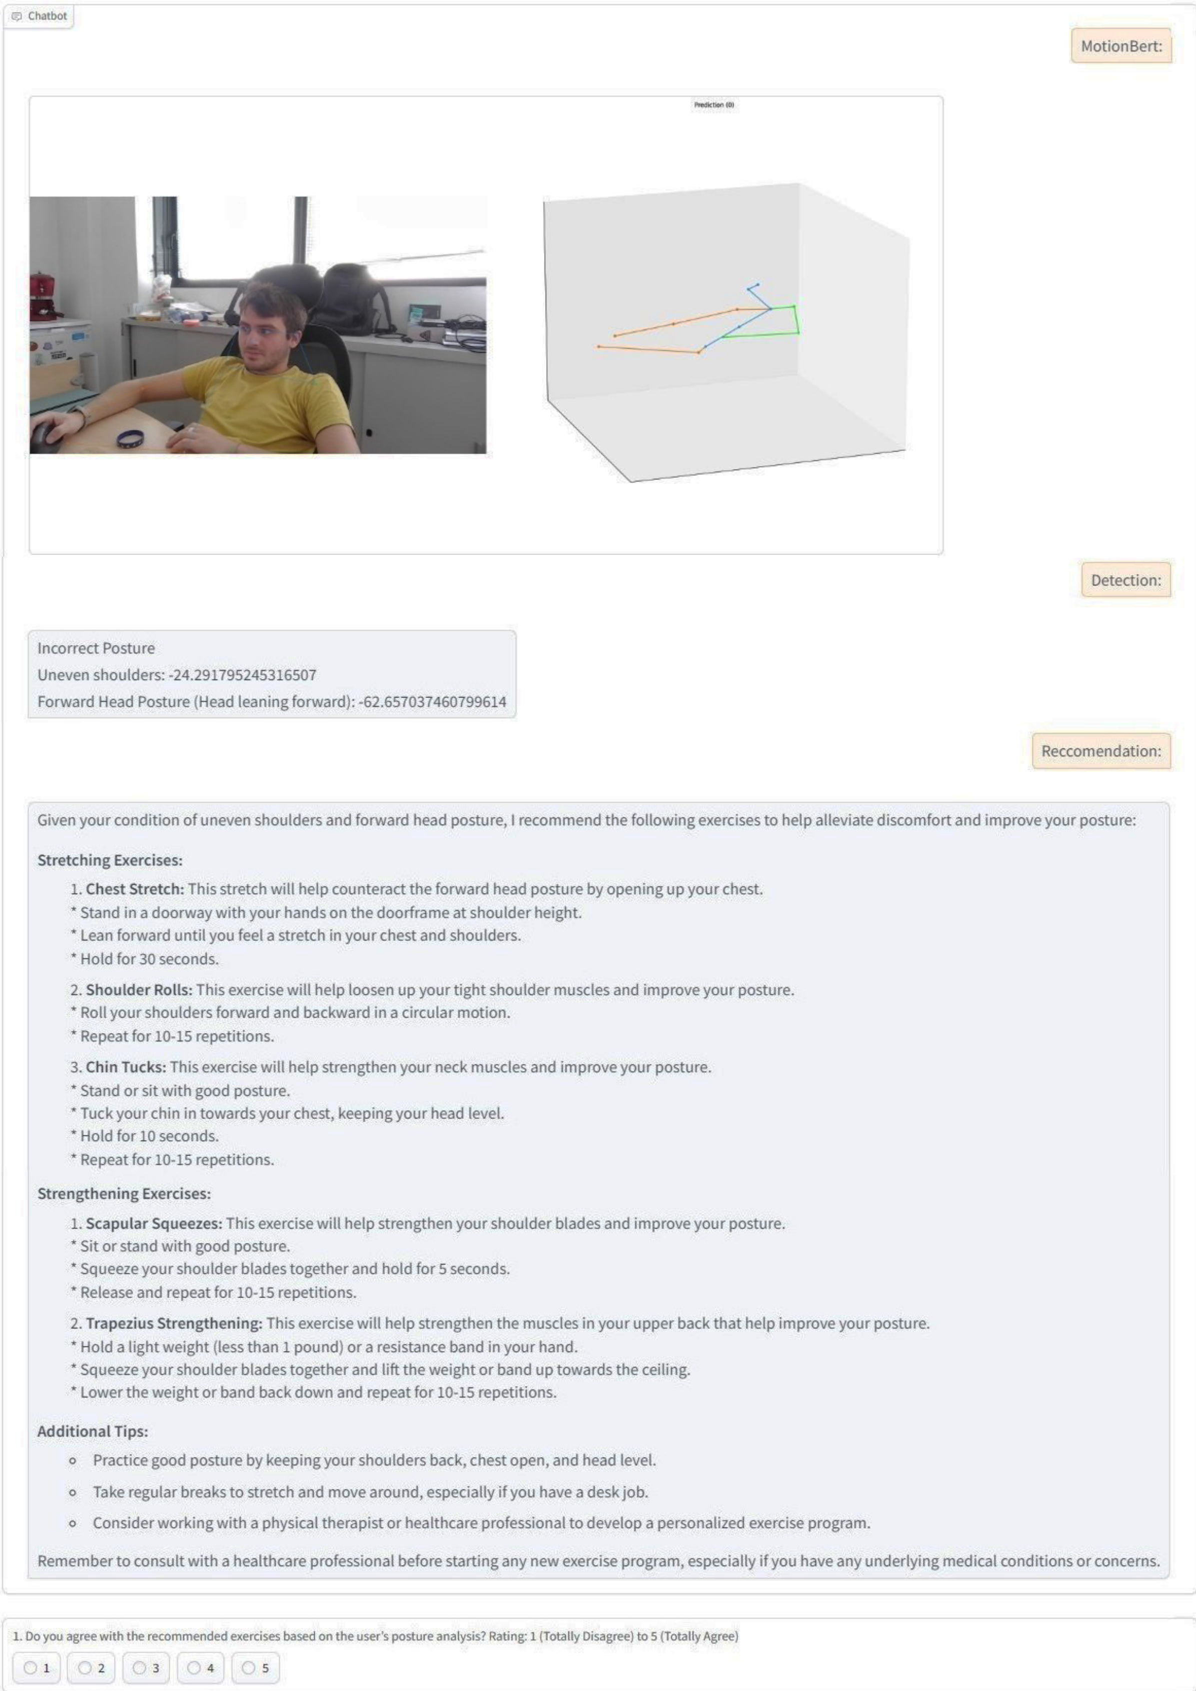

Figure 3: Our system with Gradio Interface

References

[1] M. Edward R. Laskowski, What are the risks of sitting too much?, https://www.mayoclinic.org/healthy-lifestyle/adult-health/expert-answers/sitting/faq-20058005, 2022. Accessed: 19-04-2024.

[2] W. S. Marras, S. A. Lavender, S. E. Leurgans, F. A. Fathallah, S. A. Ferguson, W. Gary Allread, S. L. Rajulu, Biomechanical risk factors for occupationally related low back disorders, Ergonomics 38 (1995) 377–410. doi:https://doi.org/10.1080/00140139508925111.

[3] D. Dunstan, E. Barr, G. Healy, J. Salmon, J. Shaw, B. Balkau, D. Magliano, A. Cameron, P. Zimmet, N. Owen, Television viewing time and mortality, Circulation 121 (2010) 384–391. doi:10.1161/CIRCULATIONAHA.109.894824.

[4] M. T. Hamilton, D. G. Hamilton, T. W. Zderic, Role of Low Energy Expenditure and Sitting in Obesity, Metabolic Syndrome, Type 2 Diabetes, and Cardiovascular Disease, Diabetes 56 (2007) 2655–2667. URL: https://doi.org/10.2337/db07-0882. doi:10.2337/db07-0882.

[5] M. Inoue, S. Yamamoto, N. Kurahashi, M. Iwasaki, S. Sasazuki, S. Tsugane, for the Japan Public Health Center-based Prospective Study Group, Daily Total Physical Activity Level and Total Cancer Risk in Men and Women: Results from a Large-scale Population-based Cohort Study in Japan, American Journal of Epidemiology 168 (2008) 391–403. URL: https://doi.org/10.1093/aje/kwn146. doi:10.1093/aje/kwn146.

[6] I. Daian, A. van Ruiten, A. Visser, S. Zubic, Sensitive chair: a force sensing chair with multimodal real-time feedback via agent, in: W. Brinkman, D. Ham, B. L. W. Wong (Eds.), Proceedings of the 14th European Conference on Cognitive Ergonomics: invent! explore!, ECCE 2007, London, UK, August 28-31, 2007, volume 250 of *ACM International Conference Proceeding Series*, ACM, 2007, pp. 163–166. URL: https://doi.org/10.1145/1362550.1362583. doi:10.1145/1362550.1362583.

[7] A. M. Lis, K. M. Black, H. Korn, M. Nordin, Association between sitting and occupational lbp, European spine journal 16 (2007) 283–298. doi:https://doi.org/10.1007/s00586-006-0143-7.

[8] O. Okunribido, M. Magnusson, M. Pope, Low back pain in drivers: The relative role of whole-body vibration, posture and manual materials handling, Journal of Sound and Vibration 298 (2006) 540–555. URL: https://www.sciencedirect.com/science/article/pii/S0022460X06004779. doi:https://doi.org/10.1016/j.jsv.2006.06.007, special Issue on the Third International Conference on Whole-body Vibration Injuries.

[9] G. Dibenedetto, M. Polignano, P. Lops, G. Semeraro, Human pose estimation for explainable corrective feedbacks in office spaces, in: Adjunct Proceedings of the 32nd ACM Conference on User Modeling, Adaptation and Personalization, UMAP Adjunct 2024, Cagliari, Italy, July 1-4, 2024, ACM, 2024. URL: https://doi.org/10.1145/3631700.3665184. doi:10.1145/3631700.3665184.

[10] S. Han, H. Kim, H. Choi, Rehabilitation posture correction using deep neural network, in: 2017 IEEE International Conference on Big Data and Smart Computing, BigComp 2017, Jeju Island, South Korea, February 13-16, 2017, IEEE, 2017, pp. 400–402. URL: https://doi.org/10.1109/BIGCOMP.2017.7881743. doi:10.1109/BIGCOMP.2017.7881743.

[11] J. C. T. Mallare, D. F. G. Pineda, G. M. Trinidad, R. D. Serafica, J. B. K. Villanueva, A. R. Dela Cruz, R. R. P. Vicerra, K. K. D. Serrano, E. A. Roxas, Sitting posture assessment using computer vision, in: 2017IEEE 9th International Conference on Humanoid, Nanotechnology, Information Technology, Communication and Control, Environment and Management (HNICEM), 2017, pp. 1–5. doi:10.1109/HNICEM.2017.8269473.

[12] Y. M. Kim, Y. Son, W. Kim, B. Jin, M. H. Yun, Classification of children's sitting postures using machine learning algorithms, Applied Sciences 8 (2018). URL: https://www.mdpi.com/2076-3417/8/8/1280. doi:10.3390/app8081280.

[13] Z. Zhao, A. Arya, R. Orji, G. Chan, Physical activity recommendation for exergame player modeling using machine learning approach, in: 8th IEEE International Conference on Serious Games and Applications for Health, SeGAH 2020, Vancouver, BC, Canada, August 12-14, 2020, IEEE, 2020, pp. 1–9. URL: https://doi.org/10.1109/SeGAH49190.2020.9201820. doi:10.1109/SEGAH49190.2020.9201820.

[14] J. Pincay, L. Terán, E. Portmann, Health recommender systems: A state-of-the-art review, 2019

Sixth International Conference on eDemocracy & eGovernment (ICEDEG) (2019) 47–55. URL: https://api.semanticscholar.org/CorpusID:189824028. doi:10.1109/ICEDEG.2019.8734362.
[15] Q. He, E. Agu, D. M. Strong, B. Tulu, Recfit: a context-aware system for recommending physical activities, in: S. K. S. Gupta, A. Banerjee (Eds.), Proceedings of the 1st Workshop on Mobile Medical Applications, MMA '14, Memphis, Tennessee, USA, November 3-6, 2014, ACM, 2014, pp. 34–39. URL: https://doi.org/10.1145/2676431.2676439. doi:10.1145/2676431.2676439.
[16] L. R. Ferretto, E. A. Bellei, D. Biduski, L. C. P. Bin, M. M. Moro, C. R. Cervi, A. C. B. D. Marchi, A physical activity recommender system for patients with arterial hypertension, IEEE Access 8 (2020) 61656–61664. URL: https://doi.org/10.1109/ACCESS.2020.2983564. doi:10.1109/ACCESS.2020.2983564.
[17] S. Sengan, S. V, R. H. Jhaveri, V. Varadarajan, R. Setiawan, L. Ravi, A secure recommendation system for providing context-aware physical activity classification for users, Security and Communication Networks 2021 (2021) 1–15. doi:https://doi.org/10.1155/2021/4136909.
[18] A. Yashudas, D. Gupta, G. C. Prashant, A. Dua, D. AlQahtani, A. S. K. Reddy, Deep-cardio: Recommendation system for cardiovascular disease prediction using iot network, IEEE Sensors Journal 24 (2024) 14539–14547. doi:10.1109/JSEN.2024.3373429.
[19] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, I. Polosukhin, Attention is all you need, in: I. Guyon, U. von Luxburg, S. Bengio, H. M. Wallach, R. Fergus, S. V. N. Vishwanathan, R. Garnett (Eds.), Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA, 2017, pp. 5998–6008. URL: https://proceedings.neurips.cc/paper/2017/hash/3f5ee243547dee91fbd053c1c4a845aa-Abstract.html.
[20] W. Kang, J. Ni, N. Mehta, M. Sathiamoorthy, L. Hong, E. H. Chi, D. Z. Cheng, Do llms understand user preferences? evaluating llms on user rating prediction, CoRR abs/2305.06474 (2023). URL: https://doi.org/10.48550/arXiv.2305.06474. doi:10.48550/ARXIV.2305.06474. arXiv:2305.06474.
[21] S. Sanner, K. Balog, F. Radlinski, B. Wedin, L. Dixon, Large language models are competitive near cold-start recommenders for language- and item-based preferences, in: J. Zhang, L. Chen, S. Berkovsky, M. Zhang, T. D. Noia, J. Basilico, L. Pizzato, Y. Song (Eds.), Proceedings of the 17th ACM Conference on Recommender Systems, RecSys 2023, Singapore, Singapore, September 18-22, 2023, ACM, 2023, pp. 890–896. URL: https://doi.org/10.1145/3604915.3608845. doi:10.1145/3604915.3608845.
[22] W. Zhu, X. Ma, Z. Liu, L. Liu, W. Wu, Y. Wang, Motionbert: A unified perspective on learning human motion representations, in: IEEE/CVF International Conference on Computer Vision, ICCV 2023, Paris, France, October 1-6, 2023, IEEE, 2023, pp. 15039–15053. URL: https://doi.org/10.1109/ICCV51070.2023.01385. doi:10.1109/ICCV51070.2023.01385.
[23] M. Contributors, Openmmlab pose estimation toolbox and benchmark, https://github.com/open-mmlab/mmpose, 2020.
[24] C. Ionescu, D. Papava, V. Olaru, C. Sminchisescu, Human3.6m: Large scale datasets and predictive methods for 3d human sensing in natural environments, IEEE Trans. Pattern Anal. Mach. Intell. 36 (2014) 1325–1339. URL: https://doi.org/10.1109/TPAMI.2013.248. doi:10.1109/TPAMI.2013.248.
[25] A. Dubey, A. Jauhri, A. Pandey, A. Kadian, A. Al-Dahle, A. Letman, A. Mathur, A. Schelten, A. Yang, A. Fan, et al., The llama 3 herd of models, CoRR abs/2407.21783 (2024). URL: https://doi.org/10.48550/arXiv.2407.21783. doi:10.48550/ARXIV.2407.21783. arXiv:2407.21783.