

Politecnico di Torino

Mass General Brigham
Spaulding Rehabilitation



HARVARD
MEDICAL SCHOOL

Politecnico di Torino

Biomedical Engineering

A.y. 2025/2026

Graduation Session July 2026

This thesis was conducted at the Motion Analysis Laboratory, Spaulding Rehabilitation Hospital, Harvard Medical School.

A Machine Learning Approach for Video-Based Movement Quality Assessment in Tele-Rehabilitation

Supervisors:

Prof. Danilo Demarchi
Prof. Paolo Bonato
Dr. Giulia Corniani

Candidate:

Alessia Giussani

A mamma e papà

Abstract

Objective: This study aims to develop a framework for the automatic assessment of Tai Chi performance in older adults using video-based analysis, with the goal of enabling scalable, objective monitoring in home-based tele-rehabilitation settings.
Background: Aging is associated with a progressive decline in balance and an increased fall risk, a leading cause of injury and lost independence in older adults. While Tai Chi is proven to mitigate these risks, access is often constrained by a scarcity of structured programs. Home-based interventions could address this issue, yet their efficacy is limited by the lack of objective performance assessment, which in turn prevents the delivery of meaningful feedback. Existing motion analysis systems typically rely on specialized equipment such as marker-based motion capture or wearable sensors, which are impractical for home use. To address this problem, we developed a video-based framework to automatically assess Tai Chi movement quality and proficiency and provide feedback for practitioners at home.
Methods: The dataset consists of 42 video recordings of older adults performing Tai Chi in home and laboratory settings, collected prior to this study. Two exercises were analyzed: Raising the Power and Golden Rooster Stands on One Leg. A Tai Chi expert annotated execution errors for each exercise, defining six and seven specific motor behaviors respectively (e.g., excessive wrist elevation). For pose estimation, we used a parametric 3D Human Body Model (SMPLest-X) to reconstruct full-body kinematics from standard monocular RGB videos. Based on these annotations, we extracted kinematic features capturing joint angles, positions, and displacements. We used UMAP projections to visually inspect the feature space and verify the discriminative power of the extracted features, prior to classification. For each exercise, we developed individual classification models for each motor behavior to detect execution errors, and trained a separate final model to estimate overall proficiency (low, medium, and high) based on expert-provided scores. Performance was evaluated using macro-averaged F1-score, chosen to account for class imbalance.
Results: The proposed approach enabled the extraction of meaningful kinematic features from video data and supported automated detection of movement errors and proficiency levels across both home and laboratory recordings. Error detection, formulated as a binary classification task, achieved good performance across most motor behaviors, with mean macro-averaged F1-scores of 0.81 and 0.85 for the two exercises respectively. Proficiency classification proved more challenging, as it involves distinguishing multiple skill levels with inherently overlapping characteristics, compounded by limited training data.
Conclusions: This work demonstrates the feasibility of objective, video-based Tai

Chi assessment using markerless motion reconstruction, supporting accessible remote monitoring of home-based practice. The proposed system is intended as a complementary tool for physiotherapists and instructors, providing quantitative information on movement quality to support, rather than replace, clinical supervision during tele-rehabilitation sessions. Future work should focus on expanding the dataset and standardizing acquisition conditions to improve robustness and classification performance.

Acknowledgements

My deepest thanks go to Professor Danilo Demarchi for giving me the opportunity to live such an incredible research experience in the United States.
My sincere thanks go to Professor Paolo Bonato for welcoming me into his lab, the Motion Analysis Laboratory, and for his continuous support, encouragement, and thoughtful mentorship throughout these months in Boston. I would also like to thank Dr. Giulia Corniani for the help she offered me along the way. I am truly grateful to all of them.
I would also like to thank everyone who took part in the different stages of the Tai Chi study, and in particular Dr. Peter Wayne, a key reference figure for the project, for his constant availability, support, and guidance throughout this journey.

This thesis marks the end of a long and challenging, yet beautiful journey, one I am deeply grateful for. One thing is certain: I could never have made it without the unwavering support of the people who love me.
Il primo ringraziamento va ai miei amici del MAL. Poco dopo il mio arrivo a Boston ho capito di essere circondata da persone stimolanti, intraprendenti e su cui poter contare. In poco tempo siete diventati la mia famiglia lontano da casa. Vi ringrazio per avermi supportata, aiutata e consigliata nei momenti di difficoltà e fragilità. Vi sono grata per tutto il tempo passato assieme dentro e fuori dal Lab. Siete una parte fondamentale di ciò che ha reso questa esperienza così incredibile. Porterò questi momenti nel mio cuore per sempre.
Ringrazio le mie amiche del Poli: Sara, Mile, Ale, Caro, Marzia, Sophie e Giulia. Ci siamo sostenute a vicenda durante cinque lunghi anni, condividendo i momenti di gioia come quelli di sconforto. Siete state un punto fermo dei miei anni a Torino, che non sarebbero stati altrettanto belli senza questa nostra, speciale, amicizia. Sarò sempre grata al Poli per avermi permesso di incontrarvi.

Un grazie di cuore va alle mie amiche di casa, da sempre, Ilaria, Greta e Sofia. Vi ringrazio per esserci state per me in ogni momento, anche durante i miei mille trasferimenti a Torino, in Belgio e in U.S.A. Mi avete costantemente sostenuta e ho sempre sentito la vostra presenza al mio fianco, anche a migliaia di chilometri di distanza.
Un ringraziamento speciale va alla mia migliore amica Ila, per avermi ascoltata e supportata, con infinita pazienza, durante i miei momenti di difficoltà e stanchezza. Sei stata un sostegno fondamentale per me durante tutti questi anni e, anche se lontane, è come se fossi sempre stata accanto a me.
Ringrazio la mia famiglia, in particolare zia Leda e zia Barbara. Ho sentito forte il vostro affetto durante questo percorso e vi ringrazio per avermi sostenuta con il cuore.
Il ringraziamento più importante di tutti va a mamma e papà. Senza il vostro immenso supporto tutto questo non sarebbe stato possibile. Grazie per avermi permesso di seguire i miei sogni e desideri, per non avermi mai tarpato le ali e per avermi sostenuta e incoraggiata in tutte le mie scelte, anche in quelle che so vi hanno fatto soffrire. Vi sarò sempre grata per aver creduto in me, in ogni momento.
Infine, ringrazio me stessa, per la mia perseveranza, forza di volontà e coraggio di mettermi in gioco.

Table of Contents

List of Tables VII

List of Figures IX

Motion Analysis Lab 2

1 Introduction 7
- 1.1 Background . . . 8
  - 1.1.1 Aging and Motor Decline . . . 8
  - 1.1.2 The Role of Tai Chi in Rehabilitation . . . 10
  - 1.1.3 Tele-Rehabilitation and Remote Monitoring . . . 11
  - 1.1.4 Keypoint-Based Human Motion Analysis . . . 13
  - 1.1.5 Parametric Human Body Models and SMPL . . . 16

2 Materials and Methods 18
- 2.1 Study Design and Dataset . . . 19
  - 2.1.1 Participants of the Study . . . 19
  - 2.1.2 Tai Chi Protocol and Exercises . . . 20
  - 2.1.3 Recording Setup . . . 26
  - 2.1.4 Scoring System . . . 26
  - 2.1.5 Videos Excluded from the Analysis . . . 27
- 2.2 Biomechanical Modeling Pipeline . . . 28
  - 2.2.1 SMPLest-X-based Pose Estimation . . . 28
  - 2.2.2 Keypoint and Pose Parameter Extraction . . . 29
  - 2.2.3 Video Segmentation - GUI . . . 29
  - 2.2.4 Keypoint Normalization and Filtering . . . 31
- 2.3 Motor Behavior Analysis . . . 32
  - 2.3.1 Repetitions Analysis . . . 37
- 2.4 Modeling Framework . . . 40
  - 2.4.1 Feature Extraction . . . 40
  - 2.4.2 Feature Selection . . . 41

2.4.3 Dimensionality Reduction - UMAP . . . 41
2.4.4 Modeling Algorithms . . . 43
2.4.5 Cross Validation . . . 48
2.4.6 Performance Metrics . . . 49

3 Results 51
3.1 Ex 1 . . . 51
3.1.1 MB 1 . . . 52
3.1.2 MB 2 . . . 54
3.1.3 MB 3 . . . 57
3.1.4 MB 4 . . . 58
3.1.5 MB 5 . . . 60
3.1.6 MB 6 . . . 62
3.1.7 Proficiency Assessment . . . 65
3.2 Ex 6 . . . 67
3.2.1 MB 1 . . . 67
3.2.2 MB 2 . . . 69
3.2.3 MB 3 . . . 71
3.2.4 MB 4 . . . 73
3.2.5 MB 5 . . . 75
3.2.6 MB 6 . . . 77
3.2.7 MB 7 . . . 79
3.2.8 Proficiency Assessment . . . 83

4 Discussion 85
4.1 Interpretation of Biomechanical Findings . . . 85
4.2 Impact on Tele-Tai Chi Rehabilitation . . . 87
4.3 Study Limitations . . . 88
4.4 Future Developments . . . 89

5 Conclusions 91

A Appendix 93
A.1 Ex 1 . . . 94
A.1.1 MB 1 . . . 94
A.1.2 MB 2 . . . 95
A.1.3 MB 3 . . . 96
A.1.4 MB 4 . . . 97
A.1.5 MB 5 . . . 98
A.1.6 MB 6 . . . 99
A.2 Ex 6 . . . 100
A.2.1 MB 1 . . . 100

A.2.2 MB 2 101
A.2.3 MB 3 102
A.2.4 MB 4 103
A.2.5 MB 5 104
A.2.6 MB 6 105
A.2.7 MB 7 106

Bibliography 108

List of Tables

2.1 Inclusion and exclusion criteria. 20
2.2 Final Dataset. 28
2.3 Total number of Motor Behaviors. 32
2.4 Frequency of Motor Behavior occurrence in the study population, Raising the Power. 34
2.5 Frequency of Motor Behavior occurrence in the study population, Golden Rooster Stands on One Leg. 37
2.6 UMAP parameters used in this analysis. 43
2.7 Hyperparameter search space for the Logistic Regression model. . . 44
2.8 Hyperparameter search space and range of values used in the RandomizedSearchCV process for the Random Forest Classifier. 47
2.9 Hyperparameter search space and range of values used in the RandomizedSearchCV process for the Random Forest Regressor. 47

3.1 Selected features for Motor Behavior 1, Exercise 1: Raising the Power. 52
3.2 Classification results reported as macro-averaged metrics, for Motor Behavior 1, Exercise 1. 53
3.3 Selected features for Motor Behavior 2, Exercise 1: Raising the Power. 55
3.4 Classification results reported as macro-averaged metrics, for Motor Behavior 2, Exercise 1. 56
3.5 Selected features for Motor Behavior 3, Exercise 1: Raising the Power. 57
3.6 Classification results reported as macro-averaged metrics, for Motor Behavior 3, Exercise 1. 58
3.7 Selected features for Motor Behavior 4, Exercise 1: Raising the Power. 59
3.8 Classification results reported as macro-averaged metrics, for Motor Behavior 4, Exercise 1. 60
3.9 Selected features for Motor Behavior 5, Exercise 1: Raising the Power. 60
3.10 Classification results reported as macro-averaged metrics, for Motor Behavior 5, Exercise 1. 61
3.11 Selected features for Motor Behavior 6, Exercise 1: Raising the Power. 62

3.12 Classification results reported as macro-averaged metrics, for Motor Behavior 6, Exercise 1. 63
3.13 Selected features for Proficiency evaluation, Exercise 1: Raising the Power. 65
3.14 Regression results for Proficiency evaluation, Exercise 1. 67
3.15 Selected features for Motor Behavior 1, Exercise 6: Golden Rooster Stands on One Leg. 68
3.16 Classification results reported as macro-averaged metrics, for Motor Behavior 1, Exercise 6. 69
3.17 Selected features for Motor Behavior 2, Exercise 6: Golden Rooster Stands on One Leg. 70
3.18 Classification results reported as macro-averaged metrics, for Motor Behavior 2, Exercise 6. 71
3.19 Selected features for Motor Behavior 3, Exercise 6: Golden Rooster Stands on One Leg. 72
3.20 Classification results reported as macro-averaged metrics, for Motor Behavior 3, Exercise 6. 73
3.21 Selected features for Motor Behavior 4, Exercise 6: Golden Rooster Stands on One Leg. 74
3.22 Classification results reported as macro-averaged metrics, for Motor Behavior 4, Exercise 6. 75
3.23 Selected features for Motor Behavior 5, Exercise 6: Golden Rooster Stands on One Leg. 76
3.24 Classification results reported as macro-averaged metrics, for Motor Behavior 5, Exercise 6. 77
3.25 Selected features for Motor Behavior 6, Exercise 6: Golden Rooster Stands on One Leg. 78
3.26 Classification results reported as macro-averaged metrics, for Motor Behavior 6, Exercise 6. 79
3.27 Selected features for Motor Behavior 7, Exercise 6: Golden Rooster Stands on One Leg. 80
3.28 Classification results reported as macro-averaged metrics, for Motor Behavior 7, Exercise 6. 81
3.29 Selected features for Proficiency evaluation, Exercise 6: Golden Rooster Stands on One Leg. 83
3.30 Regression results for Proficiency evaluation, Exercise 6. 84

List of Figures

1 ExoRoboWalker. 3
2 Frame Runner. 4
3 Experimental setup. 5

1.1 BlazePose keypoints skeleton reconstruction. 14
1.2 Movenet keypoints skeleton reconstruction. 15
1.3 OpenPose skeleton reconstruction. Image adapted from [31]. 15
1.4 From left to right: RGB image, estimated 2D body joints, skeletal representation, SMPL body model, and SMPL-X model. SMPL-X extends SMPL by including hands and face representations. Image adapted from [34]. 17

2.1 Overview of the data analysis pipeline. 19
2.2 Raising the Power exercise. 21
2.3 Push and Withdraw exercise. 22
2.4 Wave Hands like Clouds exercise. 23
2.5 Grasp the Sparrow's Tail exercise. 24
2.6 Brush Knee and Twist Step exercise. 25
2.7 Golden Rooster Stands on One Leg exercise. 26
2.8 Distribution of proficiency scores across subjects for Exercise 1 (a) and Exercise 6 (b). 27
2.9 Example of a video excluded from the analysis. 28
2.10 GUI for manual video segmentation. 30
2.11 Example of the starting and ending position for Exercise 1. 30
2.12 Example of the starting and ending position for Exercise 6. 31
2.13 Dataset summary for Exercise 1, Raising the Power. 34
2.14 Dataset summary for Exercise 6, Golden Rooster Stands on One Leg. 37
2.15 Distribution of an example feature for Exercise 1, Motor Behavior 1. 38
2.16 Distribution of an example feature for Exercise 1, Motor Behavior 4. 39
2.17 Intra-subject variability ratio for Exercise 1, Motor Behavior 1. . . . 39
2.18 Intra-subject variability ratio for Exercise 1, Motor Behavior 4. . . . 40
2.19 Nested Leave-One-Group-Out (LOGO) cross-validation workflow. . 49

3.1 UMAP projection, Motor Behavior 1, Exercise 1: Raising the Power. 53
3.2 Confusion Matrices for Motor Behavior 1: Random Forest (a) and Logistic Regression (b), Exercise 1: Raising the Power. 53
3.3 UMAP projection, Motor Behavior 2, Exercise 1: Raising the Power. 56
3.4 Confusion Matrices for Motor Behavior 2: Random Forest (a) and Logistic Regression (b), Exercise 1: Raising the Power. 56
3.5 UMAP projection, Motor Behavior 3, Exercise 1: Raising the Power. 57
3.6 Confusion Matrices for Motor Behavior 3: Random Forest (a) and Logistic Regression (b), Exercise 1: Raising the Power. 58
3.7 UMAP projection, Motor Behavior 4, Exercise 1: Raising the Power. 59
3.8 Confusion Matrices for Motor Behavior 4: Random Forest (a) and Logistic Regression (b), Exercise 1: Raising the Power. 59
3.9 UMAP projection, Motor Behavior 5, Exercise 1: Raising the Power. 61
3.10 Confusion Matrices for Motor Behavior 5: Random Forest (a) and Logistic Regression (b), Exercise 1: Raising the Power. 61
3.11 UMAP projection, Motor Behavior 6, Exercise 1: Raising the Power. 63
3.12 Confusion Matrices for Motor Behavior 6: Random Forest (a) and Logistic Regression (b), Exercise 1: Raising the Power. 63
3.13 Summary of Motor Behavior analysis results for Exercise 1 with Random Forest. 64
3.14 Summary of Motor Behavior analysis results for Exercise 1 with Logistic Regression. 64
3.15 Nested mRMR approach: UMAP projection (a) and regression performance (b), Exercise 1: Raising the Power. 66
3.16 Manual feature selection approach: UMAP projection (a) and regression performance (b), Exercise 1: Raising the Power. 66
3.17 Prediction error per subject, manual feature selection, Exercise 1: Raising the Power. 67
3.18 UMAP projection, Motor Behavior 1, Exercise 6: Golden Rooster Stands on One Leg. 68
3.19 Confusion Matrices for Motor Behavior 1: Random Forest (a) and Logistic Regression (b), Exercise 6: Golden Rooster Stands on One Leg. 69
3.20 UMAP projection, Motor Behavior 2, Exercise 6: Golden Rooster Stands on One Leg. 70
3.21 Confusion Matrices for Motor Behavior 2: Random Forest (a) and Logistic Regression (b), Exercise 6: Golden Rooster Stands on One Leg. 71
3.22 UMAP projection, Motor Behavior 3, Exercise 6: Golden Rooster Stands on One Leg. 72

3.23 Confusion Matrices for Motor Behavior 3: Random Forest (a) and Logistic Regression (b), Exercise 6: Golden Rooster Stands on One Leg. . . . 73
3.24 UMAP projection, Motor Behavior 4, Exercise 6: Golden Rooster Stands on One Leg. . . . 74
3.25 Confusion Matrices for Motor Behavior 4: Random Forest (a) and Logistic Regression (b), Exercise 6: Golden Rooster Stands on One Leg. . . . 75
3.26 UMAP projection, Motor Behavior 5, Exercise 6: Golden Rooster Stands on One Leg. . . . 76
3.27 Confusion Matrices for Motor Behavior 5: Random Forest (a) and Logistic Regression (b), Exercise 6: Golden Rooster Stands on One Leg. . . . 77
3.28 UMAP projection, Motor Behavior 6, Exercise 6: Golden Rooster Stands on One Leg. . . . 78
3.29 Confusion Matrices for Motor Behavior 6: Random Forest (a) and Logistic Regression (b), Exercise 6: Golden Rooster Stands on One Leg. . . . 79
3.30 UMAP projection, Motor Behavior 7, Exercise 6: Golden Rooster Stands on One Leg. . . . 80
3.31 Confusion Matrices for Motor Behavior 7: Random Forest (a) and Logistic Regression (b), Exercise 6: Golden Rooster Stands on One Leg. . . . 81
3.32 Summary of Motor Behavior analysis results for Exercise 6 with Random Forest. . . . 82
3.33 Summary of Motor Behavior analysis results for Exercise 6 with Logistic Regression. . . . 82
3.34 Proficiency assessment, Exercise 6: UMAP projection (a) and regression performance (b), Golden Rooster Stands on One Leg. . . . 84
3.35 Prediction error per subject, Proficiency assessment, Exercise 6: Golden Rooster Stands on One Leg. . . . 84

A.1 Feature ranking — Exercise 1, Motor Behavior 1. . . . 94
A.2 Top 3 features by importance — Exercise 1, Motor Behavior 1. . . . 94
A.3 Feature ranking — Exercise 1, Motor Behavior 2. . . . 95
A.4 Top 3 features by importance — Exercise 1, Motor Behavior 2. . . . 95
A.5 Feature ranking — Exercise 1, Motor Behavior 3. . . . 96
A.6 Top 3 features by importance — Exercise 1, Motor Behavior 3. . . . 96
A.7 Feature ranking — Exercise 1, Motor Behavior 4. . . . 97
A.8 Top 3 features by importance — Exercise 1, Motor Behavior 4. . . . 97
A.9 Feature ranking — Exercise 1, Motor Behavior 5. . . . 98

A.10 Top 3 features by importance — Exercise 1, Motor Behavior 5. . . . 98
A.11 Feature ranking — Exercise 1, Motor Behavior 6. 99
A.12 Top 3 features by importance — Exercise 1, Motor Behavior 6. . . . 99
A.13 Feature ranking — Exercise 6, Motor Behavior 1. 100
A.14 Top 3 features by importance — Exercise 6, Motor Behavior 1. . . . 100
A.15 Feature ranking — Exercise 6, Motor Behavior 2. 101
A.16 Top 3 features by importance — Exercise 6, Motor Behavior 2. . . . 101
A.17 Feature ranking — Exercise 6, Motor Behavior 3. 102
A.18 Top 3 features by importance — Exercise 6, Motor Behavior 3. . . . 102
A.19 Feature ranking — Exercise 6, Motor Behavior 4. 103
A.20 Top 3 features by importance — Exercise 6, Motor Behavior 4. . . . 103
A.21 Feature ranking — Exercise 6, Motor Behavior 5. 104
A.22 Top 3 features by importance — Exercise 6, Motor Behavior 5. . . . 104
A.23 Feature ranking — Exercise 6, Motor Behavior 6. 105
A.24 Top 3 features by importance — Exercise 6, Motor Behavior 6. . . . 105
A.25 Feature ranking — Exercise 6, Motor Behavior 7. 106
A.26 Top 3 features by importance — Exercise 6, Motor Behavior 7. . . . 106

Motion Analysis Lab

During my period as a Research Scholar at the Motion Analysis Laboratory led by Professor Paolo Bonato, I had the opportunity not only to carry out my own Master's Thesis research project but also to observe, participate in, and explore several other projects developed by fellow students and visiting PhD researchers. The continuous exchange of ideas with the rest of the team and the chance to learn something new every day were among the most enriching aspects of this experience. Even everyday informal conversations often led to interesting insights and new ideas.
Some of the projects I explored include the following:

Exoskeleton Study

The ExoRoboWalker, shown in Figure 1, is a lower-limb overground exoskeleton with six degrees of freedom, designed to support and assist gait. Its articulated joints and adjustable lateral bars allow the device to accommodate various body sizes, while a pelvic belt with shoulder braces provides essential trunk support and ensures user stability. Each joint integrates a suite of sensors to capture kinetic and kinematic gait data:

- **Potentiometers:** measure joint angular position.
- **Strain gauges:** track the interaction forces between the user and the device.
- **Pressure sensors:** embedded in the insoles to detect the phases of the gait cycle.

The device operates in a "transparent mode," following the user's natural gait pattern to minimize interference and prevent unnecessary increases in metabolic cost. In this study, I participated as a healthy subject for data acquisition, to characterize the device's interaction with the user and validate the control framework. The aim of the study was to investigate the neuromuscular mechanisms underlying motor adaptation in response to specific perturbations delivered during walking. To this end, a dedicated perturbation paradigm was implemented, using the exoskeleton

as a testing platform. The experimental sessions involved a multimodal assessment, simultaneously recording electroencephalography (EEG) and electromyography (EMG) signals. While the EMG recordings provided insight into spinal reflex responses to perturbations, the EEG data allowed the higher-level neural processes involved in motor adaptation to be examined, such as internal model formation, error processing, and the balance between feedforward and feedback control mechanisms.

This study has significant clinical implications. By gaining a deeper understanding of how the nervous system perceives and responds to gait perturbations, it becomes possible to better predict how individuals might respond to robot-assisted rehabilitation. Ultimately, these insights will help develop personalized, optimized rehabilitation protocols that adapt to the specific impairments and recovery potential of each patient.

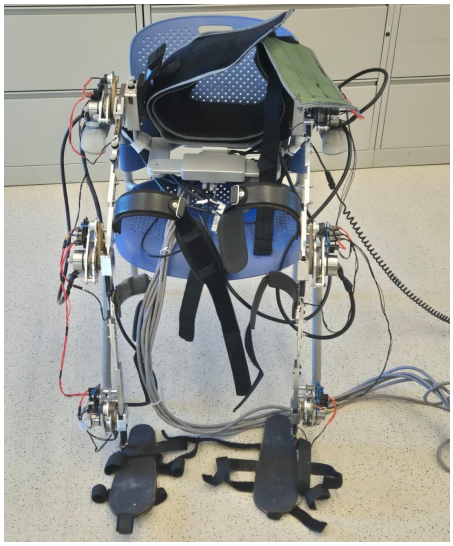


Figure 1: ExoRoboWalker.

Frame Runner Study

The Frame Runner, shown in Figure 2, is an adaptive mobility device designed to support walking and running in individuals with neurological motor impairments. It enables dynamic, self-propelled movement in an engaging and inclusive environment and may improve physical activity participation, cardiovascular fitness, muscle strength, and overall quality of life, particularly in pediatric populations where motivation and adherence to repetitive therapy are often challenging. However, its biomechanical effects on lower-limb loading and movement patterns remain poorly understood.

In this study, I participated as a healthy subject in a laboratory-based dry run. This study evaluates the safety and clinical suitability of the Frame Runner by analyzing gait and running biomechanics in both healthy participants and individuals with neurological motor impairments. The aims of the project are to characterize lower-limb kinematics, kinetics, and joint loading, as well as to assess how these parameters change during use in both simulated and real pathological conditions.
To enable the estimation of lower-body biomechanics in out-of-laboratory settings, a musculoskeletal model was developed in OpenSim using data collected from preliminary laboratory trials involving healthy participants and a physical therapist simulating cerebral palsy gait patterns. The model was validated against gold-standard 3D motion capture and force plate data, showing good agreement in estimating joint kinematics and dynamics.
To extend biomechanical assessment outside the laboratory, markerless human pose estimation methods are being explored as an alternative to wearable sensors. OpenPose and SMPLest-X were initially tested using single-view RGB recordings from frontal and lateral perspectives. Due to evident inaccuracies in body alignment and motion reconstruction, OpenPose was excluded from further analysis, while current work focuses on SMPLest-X, which showed more stable and accurate reconstructions and is undergoing further validation.



Figure 2: Frame Runner.

Accelera Study

One of the projects I had the opportunity to explore involved stochastic resonance (SR) stimulation, a nonlinear phenomenon in which subthreshold noise can improve weak signal detection. In this study, I participated as a subject in a dry run in order to directly experience the intervention being investigated.
The study evaluated the immediate effects of vibrotactile SR applied with 2 piezoelectric tactors in stroke survivors, using a single-session crossover design. The experimental setup is shown in Figure 3. Participants completed tasks under both stimulation and no-stimulation conditions. Outcomes included measures of upper-limb motor performance (Wolf Motor Function Test (WMFT) subset and

gross arm movements), as well as physiological and kinematic measures including surface EMG and accelerometer-based assessment of movement smoothness.
Preliminary results showed that SR stimulation improved distal movements in a subset of patients, with improvements observed in WMFT subtasks, including reduced completion times in the manipulation phase and, in one stroke survivor, successful execution of a task that was not possible without stimulation. Effects on gross motor tasks and movement quality were more variable.
These results indicate that vibrotactile SR stimulation can acutely improve distal motor performance, particularly in tasks that rely on sensory feedback, supporting its potential as a non-invasive wearable approach for upper-limb rehabilitation after stroke.

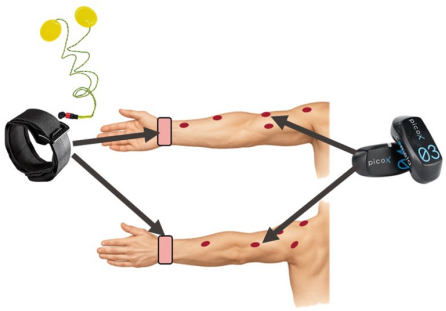


Figure 3: Experimental setup.

Chapter 1

Introduction

Population aging is associated with progressive physiological decline, including reduced balance control, decreased muscular strength, and increased risk of falls. Falls are a leading cause of injury in older adults, affecting roughly one third of individuals aged 65 years and over [1], and may result in fractures, head trauma, loss of independence, and extensive healthcare resource utilization [2]. Beyond physical injury, falls also have significant psychological and social consequences, including fear of falling, reduced self-confidence, and limitations in social interactions [3].
For all these reasons, fall prevention is considered a major priority in geriatric care and public health.
Evidence shows that interventions targeting balance, strength, and postural control can mitigate fall risk, highlighting the importance of early assessment and preventive strategies for maintaining mobility and safety in aging individuals.
Tai Chi, a traditional Chinese martial art used as a health-promoting exercise, incorporates slow, controlled movements with weight shifting, trunk rotation, and semi-squat postures, which are known to enhance balance performance and lower-limb strength. Due to its low-impact and slow-motion characteristics, Tai Chi is particularly suitable for older adults and can be easily integrated into daily routines. Numerous studies have shown that Tai Chi is an effective exercise for preventing falls and enhancing balance in older adults, both in healthy individuals and those at higher risk of falling [2, 4]. Its benefits become greater with increased practice duration and frequency.
This Thesis focuses on the development of Machine Learning algorithms to automatically assess the performance of older adults practicing Tai Chi from video recordings, collected primarily in home environments and, to a lesser extent, in the laboratory setting. Body joint mechanical information is extracted using a parametric Human Pose Estimation approach, enabling quantitative evaluation of Tai Chi proficiency. The Thesis is structured as follows.
This First Chapter begins with a brief overview of the physical and psychological

effects of aging on the older adult population. It then presents two sections focusing on the rehabilitative benefits of Tai Chi practice and its potential applications in tele-rehabilitation. Finally, the chapter provides a concise introduction to Human Pose Estimation models, including both keypoint detection-based and parametric approaches, setting the basis for automated assessment of movement proficiency.
In the Second Chapter, the Materials and Methods applied in this study are described. The Chapter is organized into four main sections:

The first consists of an overview of the study design and dataset, including participant details, the Tai Chi exercise protocol, the setup for video recordings in both home and laboratory settings, and the scoring system used to quantify performance levels.

The second is dedicated to the description of the biomechanical modeling pipeline, which involves SMPLest-X-based pose estimation, extraction of keypoints and pose parameters, video segmentation, normalization and filtering of keypoints.

The third presents the analysis of the different motor behaviors, focusing on the identification of deviations from ideal Tai Chi movements.

The fourth and last one describes the modeling framework, covering feature extraction and selection, dimensionality reduction, model training, validation, and evaluation metrics used to assess the models' performance.

The Third Chapter presents the results obtained in this study, organized according to each Tai Chi exercise analyzed.
The Fourth Chapter discusses these findings in detail, highlighting their relevance and addressing the limitations of the study while suggesting directions for future research.
Finally, the Fifth Chapter concludes the Thesis by summarizing the main outcomes and highlighting the key takeaway messages from the work.

1.1 Background

1.1.1 Aging and Motor Decline

With advancing age, sensorimotor control and motor function gradually decline. Motor performance difficulties related to aging are caused by complex interactions between the central and peripheral nervous systems as well as the neuromuscular system [5]. Common motor impairments associated with aging include reduced muscle strength, slower gait speed, increased movement variability, and poorer

balance control [6]. These deficits have important functional effects on daily activities. In particular, impairments in gait and balance are clinically significant, as falls represent a major cause of injury and health-related problems among older adults.

Population-based studies show that approximately 10% of older adults experience at least two falls per year. Following a fall, about one quarter of older adults limit their daily activities for at least one day or require medical attention. More severe injuries, such as sprains, strains, joint dislocations, fractures, and concussions, occur in roughly 10% of fall events [7].

Approximately 20–30% of falls result in moderate to severe injuries, which can limit mobility and reduce quality of life [8]. Nearly one-third of older adults report a fear of falling, even without having experienced a previous fall. This fear often causes limitations in mobility, reduced physical activity, depression, social isolation, and loneliness [9].

In terms of fall biomechanics, when the body is stationary, the center of mass and the center of pressure are vertically aligned over the support surface. During gait initiation, lifting one foot reduces the base of support and creates a displacement between the center of mass and the center of pressure, challenging balance and potentially leading to a fall if stability is not properly controlled [10]. Skeletal muscle coordination plays a major role in controlling body sway and gait parameters. With aging, reductions in muscle mass, contractility, and joint function, together with degenerative spinal changes, can alter the position of the center of mass and increase the risk of falls [11].

Several factors have been identified as contributing to falls in older adults. Among the most important are impairments in balance and gait, polypharmacy, and a history of previous falls. Additional factors such as visual impairments, cognitive decline, and environmental hazards can further increase the risk of falling. These factors often interact with each other, meaning that falls rarely result from a single cause. Instead, the risk of falling typically increases as the number of contributing factors rises [12].

For this reason, fall prevention has become a major public health priority. Among the available strategies, exercise programs have been shown to effectively reduce both the rate of falls and the likelihood of falling in older adults. Well-designed physical activity programs can therefore play a key role in preventing falls [13].

In this context, Tai Chi is particularly beneficial, as it combines slow, controlled movements with balance, coordination, and postural control training, making it effective for reducing fall risk.

1.1.2 The Role of Tai Chi in Rehabilitation

Physical activities with a focus on improving balance and muscle strength are one of the most effective strategies for preventing falls in older adults. In addition to reducing fall risk, physical activity offers multiple broader health benefits and can be incorporated into a healthy lifestyle, rather than targeting a single health problem [14].
Tai Chi is a traditional Chinese form of exercise that has been practiced for centuries to support health and longevity and has gained increasing popularity in Western countries in recent decades.
It consists of slow, continuous movements coordinated with controlled breathing and mindful awareness, supporting both physical and cognitive functions. The practice is typically performed in a semi-squat posture, with movements executed at a very slow pace. Although the overall intensity of Tai Chi exercises is generally considered low, the prolonged maintenance of postures and the controlled movement speed can increase muscular demand, particularly in the lower limbs. Furthermore, the height of the posture and the duration of movements can be adjusted, allowing the intensity of the practice to be adapted to different levels of physical ability.
Due to its slow, low-impact, and controlled nature, Tai Chi is widely regarded as a safe and gentle form of physical activity that incorporates elements of balance, strength, and coordination [15, 16].
Evidence from recent systematic reviews and meta-analyses highlights its broad range of health benefits across physical, cognitive, and psychological areas [15]:

- **Physical function:** Moderate to strong evidence indicates that regular Tai Chi practice can significantly improve balance, mobility, muscular strength, and proprioception, while also promoting cardiorespiratory health.
- **Cognitive function:** Positive effects have been reported on global cognition and executive function, suggesting that Tai Chi may help maintain cognitive health in older adults.
- **Sleep and psychological well-being:** Improvements in sleep quality have been observed, along with reductions in anxiety, depression, and stress, showing benefits for mental health and overall quality of life.
- **Risk reduction:** Tai Chi has been associated with a lower risk of falls and improvements in stroke-related risk factors, highlighting its potential role in preventive medicine.
- **Other potential benefits:** New evidence suggests additional benefits for body composition, bone mineral density, and general physical function, further supporting its overall health benefits.

Among the different benefits associated with Tai Chi practice, improvements in balance are particularly relevant. Balance is a fundamental component of mobility and independence and is closely associated with fall risk in older adults.
Several studies have examined the effects of Tai Chi practice on balance control. For instance, a 12-week Tai Chi training program was shown to significantly improve balance in older adults compared with a control group.
In the same study, no falls were reported among participants practicing Tai Chi during the intervention period, while falls were observed in the control group.
Interestingly, no significant changes in lower-limb strength were observed, suggesting that improvements in balance may occur even in the absence of noticeable increases in muscle strength [17].
Evidence from randomized controlled trials also supports the positive effects of Tai Chi on balance. For instance, a 24-week Tai Chi training program was found to significantly improve postural stability compared with resistance training and stretching in individuals with Parkinson's disease. Participants practicing Tai Chi also showed improvements in mobility and gait performance and experienced fewer falls compared with the group following a stretching-based exercise program [18].
Overall, current evidence highlights the potential of Tai Chi as an effective practice for improving balance control and reducing the risk of falls, supporting its growing use in fall prevention strategies for older adults.

1.1.3 Tele-Rehabilitation and Remote Monitoring

In recent years, telerehabilitation has emerged as a promising approach to provide rehabilitation and exercise interventions remotely through digital technologies.
It is considered a branch of telehealth and allows rehabilitation activities to be performed while the physical therapist and the patient are in different locations. One of its main objectives is to increase accessibility to rehabilitation practices and improve continuity of care, particularly for populations or individuals who live far from rehabilitation centers or hospitals [19].
Telerehabilitation may also reduce barriers related to transportation, mobility limitations, and time constraints, allowing patients to participate in rehabilitation programs directly from their own homes. These advantages are particularly relevant for older adults, who may face difficulties in attending in-person exercise sessions regularly.
Telerehabilitation has generally been well accepted by both patients and healthcare professionals, with several studies reporting high levels of satisfaction and feasibility across different clinical settings. In many cases, patients appreciate the possibility of performing rehabilitation exercises at home while maintaining regular contact with healthcare providers through digital platforms [20].
However, despite these positive findings, important challenges remain. In particular,

accurately evaluating patients' movement execution and overall motion quality during telerehabilitation sessions can be more challenging in the absence of direct physical supervision by clinicians.
For this reason, the development of reliable technologies to remotely monitor patients and objectively assess their performance is becoming increasingly important to ensure the effectiveness and safety of telerehabilitation programs.
Musculoskeletal rehabilitation often includes performing specific task-oriented exercises designed to restore joint mobility, muscle strength, and neuromuscular coordination. However, when these exercises are performed without supervision at home, there is a risk of incorrect execution, which may reduce therapeutic effectiveness and potentially lead to injury or delayed recovery. In addition, the lack of real-time feedback on posture, movement accuracy, and exercise quality may interfere with the achievement of optimal rehabilitation outcomes [21].
These challenges are particularly relevant for exercise-based practices such as Tai Chi, which rely on slow and controlled movements as well as correct body alignment. Ensuring proper execution is especially critical when Tai Chi is delivered remotely through telerehabilitation programs.
In recent years, several studies have explored the feasibility of delivering Tai Chi programs remotely through digital platforms, an approach referred to as "Tele-Tai Chi". For example, previous studies have developed Tai Chi programs delivered through videoconferencing systems, allowing participants to follow exercise sessions from home while interacting with instructors in real time. These studies reported good adherence and demonstrated that remotely delivered Tai Chi practice is feasible for older adults [22].
More broadly, telewellness programs delivered through videoconferencing technologies can improve access to structured exercise protocols for older adults with mobility limitations and may also provide opportunities for social interaction among participants, potentially enhancing motivation and adherence to the rehabilitation programs [23].
Despite these advantages, delivering exercise-based interventions remotely still presents several challenges related to the objective evaluation of movement execution. To address these limitations, various technological approaches have been proposed to support remote monitoring and evaluation of physical exercises. Existing approaches include wearable sensors, motion capture systems, and vision-based techniques. Wearable devices, such as inertial measurement units (IMUs), are electronic sensors incorporating accelerometers and gyroscopes, which measure the acceleration and angular velocity of body segments. IMUs provide objective information on joint angles, limb orientation, and movement dynamics [24, 25]. However, they require subjects to wear multiple sensors, which could limit their practicality in home-based rehabilitation scenarios.
In contrast, markerless motion capture systems (camera-based approaches) have

gained increasing attention, as they allow the analysis of human movements without the need for wearable devices. Recent advances in pose estimation techniques allow the automatic extraction of body keypoints from video recordings, offering a promising solution for monitoring exercise performance in telerehabilitation contexts [26]. In the following sections, some representative examples of these approaches, focusing on keypoint-based and mesh-based pose estimation methods, will be briefly presented.

1.1.4 Keypoint-Based Human Motion Analysis

Quantitative analysis of human movement is central to clinical assessment, rehabilitation, and sports science, where accurate estimation of body motion is often required. For this reason, several approaches have been developed, including marker-based motion capture systems and markerless methods based on visual pose estimation. Marker-based motion capture systems are generally considered the gold standard for three-dimensional kinematic analysis, as they provide highly accurate and detailed measurements of joint motion.
However, these systems require specialized and complex equipment, controlled laboratory environments and time-consuming setup procedures, which limit their use in more practical scenarios and real-world applications [27].

MediaPipe

Among the markerless methods, MediaPipe Pose (BlazePose) is a lightweight deep learning solution for real-time human pose estimation. It is part of the MediaPipe framework, developed by Google, which provides tools to build pipelines for processing visual data such as images and videos.
BlazePose estimates 2D (and optionally 3D) coordinates of 33 body keypoints from RGB images. This enables fast and non-invasive motion analysis, even on mobile devices. As illustrated in Figure 1.1, this representation corresponds to a structured skeletal model of the human body. The pipeline typically combines a person detector with a pose tracking model, allowing continuous estimation of joint positions across frames while keeping a high computational efficiency.
Due to these characteristics, it is widely used in applications related to fitness monitoring and rehabilitation [28, 29].

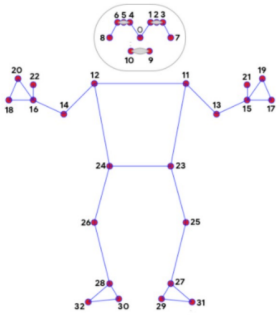


Figure 1.1: BlazePose keypoints skeleton reconstruction.

MoveNet

MoveNet is a real-time human pose estimation model developed by Google, designed for fast and accurate detection of 2D body keypoints [30]. It is based on deep convolutional neural networks and trained on large datasets, including COCO and internal Google datasets derived from fitness, yoga, and dance activities. This allows robust performance across different poses, environments, and visual conditions.
As illustrated in Figure 1.2, the model detects 17 keypoints corresponding to major anatomical joints, such as shoulders, elbows, wrists, hips, knees, and ankles. Two main versions are available: Lightning, optimized for speed, and Thunder, optimized for higher accuracy. Both single-person and multi-person configurations are supported, although the single-person version is often preferred for real-time applications such as biomechanics and rehabilitation.
The outputs of this model are normalized keypoint coordinates and confidence scores. These are obtained through a multi-stage pipeline that includes localization of the person, keypoint regression, and refinement based on heatmaps. Temporal filtering is then applied to reduce jitter and ensure smoother motion, even during fast movements.
Overall, MoveNet is well suited for real-time motion analysis in applications such as rehabilitation, sports, and human movement monitoring.

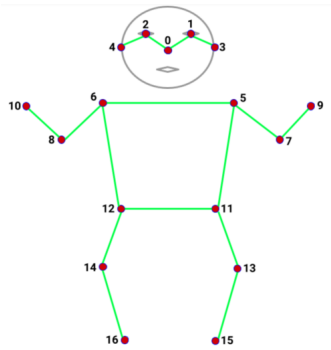


Figure 1.2: Movenet keypoints skeleton reconstruction.

OpenPose

OpenPose is a widely used markerless method for human pose estimation that can detect body keypoints in real time from images and videos. It is able to estimate the 2D pose of multiple people at the same time, making it suitable for scenes with several interacting subjects, as shown in Figure 1.3. The method follows a bottom-up approach: body parts are first detected independently and then grouped to form individual poses. This is done using Part Affinity Fields (PAFs), which describe the spatial relationships between joints. This strategy allows the system to maintain good accuracy while achieving real-time performance, even when many people are present in the scene.
In addition to detecting body joints, OpenPose can also estimate keypoints of the hands, face, and feet, providing a rich set of information for human motion analysis [31].



Figure 1.3: OpenPose skeleton reconstruction. Image adapted from [31].

Despite their effectiveness in many different applications, keypoint-based human pose estimation methods present limitations in representing the full geometry

and physical structure of the human body. These shortcomings have led to the development of more expressive body modeling approaches.

1.1.5 Parametric Human Body Models and SMPL

Keypoint-based human pose estimation methods have significantly progressed in recent years. However, they still present several limitations. First, occlusions and complex body configurations can reduce the accuracy of joint localization. Moreover, when estimating 3D pose from 2D images, the problem is rather ambiguous, due to the lack of depth information. Another relevant challenge is the limited availability of accurate 3D annotated data, which makes training and evaluation more difficult [32].
To address the limitations presented by keypoint-based approaches of human pose estimation, parametric human body models have been introduced. These models represent the human body in a more complete manner by including both pose and body shape, typically in the form of a 3D mesh. This allows a more realistic and consistent description of human motion.
Among the existing parametric human body models for human pose estimation, SMPL, STAR, and GHUM are commonly used to represent the human body in a low-dimensional parametric space.

SMPL and SMPL-X

SMPL (Skinned Multi-Person Linear model) is a parametric statistical model for human body representation that describes both the pose and shape of the body through a compact, low-dimensional set of parameters. Its key advantage is that, despite this small set of parameters, it still produces a detailed 3D mesh with realistic geometry: a triangulated surface that approximates the shape of the human body in three dimensions.
The main idea behind SMPL is to combine a standard skeletal structure with learned deformations that correct and refine the shape of the body during various movements. This allows the model to be more realistic than traditional skinning methods, which often produce unnatural deformations of joints [33].
Thanks to its simplicity and flexibility, SMPL can be easily used in different computer vision and graphics applications, and it has become one of the most widely used models for 3D human body representation.
An extension of SMPL is **SMPL-X (SMPL eXpressive)**, which enriches the representation of the first SMPL model by including hands and face representation in addition to the full body. As shown in Figure 1.4, this extension allows the modeling of more expressive human motions, such as hand gestures and facial expressions, which are not captured by the previously mentioned standard

SMPL model. Despite this increased level of detail, SMPL-X maintains the same advantages of the original model, such as a compact parameter representation, differentiability, and compatibility with standard graphics pipelines.
For all these reasons, SMPL-X has become widely used in full-body human motion capture and reconstruction tasks [34].
In this work, SMPL-X parameters are estimated directly from video using SMPLest-X [35], as described in detail in the Materials and Methods chapter.

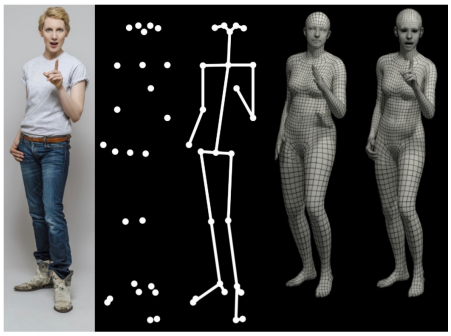


Figure 1.4: From left to right: RGB image, estimated 2D body joints, skeletal representation, SMPL body model, and SMPL-X model. SMPL-X extends SMPL by including hands and face representations. Image adapted from [34].

Chapter 2

Materials and Methods

This Chapter describes the materials and methods used to obtain the results presented in this thesis. It includes both the data collection process and the subsequent data analysis pipeline.
First, the study design and dataset are presented, including details about the participants, the Tai Chi exercise protocol, the recording setup, the evaluation scores, and the videos excluded from the analysis.
The chapter then presents the data processing and analysis pipeline, which includes extraction of joint keypoints, video segmentation, and keypoint normalization and filtering. Relevant movement features are then extracted, followed by feature selection and dimensionality reduction.
Finally, classification and regression models are used to evaluate movement performance.
A schematic overview of the entire pipeline is provided in Figure 2.1.

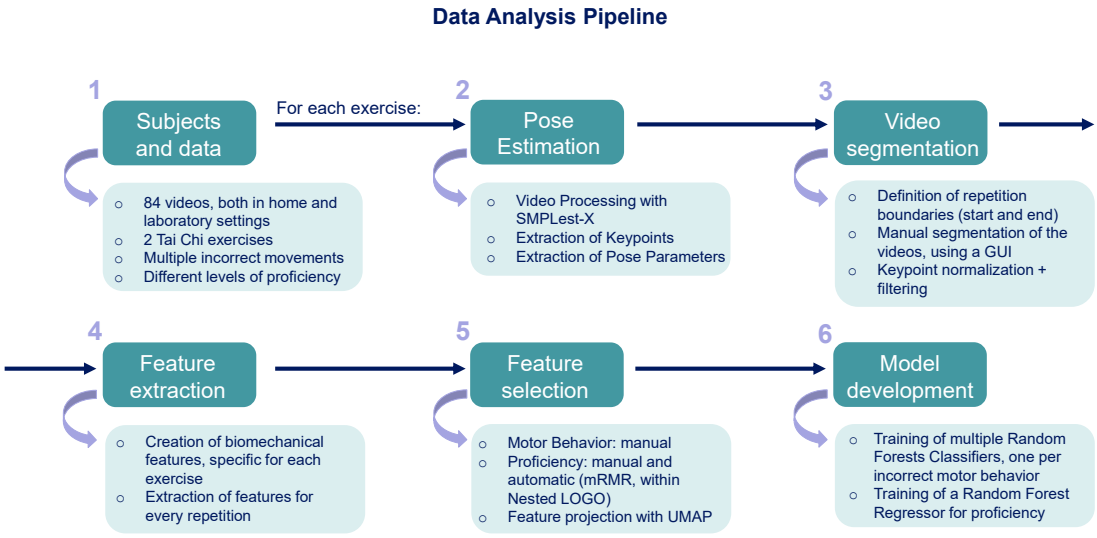


Figure 2.1: Overview of the data analysis pipeline.

2.1 Study Design and Dataset

2.1.1 Participants of the Study

Thirty older adults living in the community were recruited to take part in this study. Participants were Tai Chi–naïve and took part in a single-arm longitudinal study investigating the learning of Tai Chi exercises in a home-based setting. Data collection was carried out both remotely and in person, with in-person sessions taking place at the Motion Analysis Laboratory (MAL) of Spaulding Rehabilitation Hospital. Eligible participants to take part in this study were required to meet the inclusion criteria and not present any of the exclusion conditions listed in Table 2.1. Participants were recruited through various ways, including clinician referrals, institutional registries, community-based recruitment, and informational flyers. All participants provided informed consent before taking part in the study.

Table 2.1: Inclusion and exclusion criteria.

Inclusion Criteria	Exclusion Criteria
Age 60–85 years	Chronic neuromuscular conditions (e.g., Parkinson's disease, multiple sclerosis, stroke)
Tai Chi–naïve	Recent hospitalization (past 6 months)
Ability to walk 15 min without assistive device	Active cancer
Access to computer, tablet or smartphone	Musculoskeletal conditions requiring chronic pain medication
Working email address	Diagnosed dementia or cognitive impairment (MoCA-Blind < 18)

2.1.2 Tai Chi Protocol and Exercises

The intervention protocol consisted of a 12-week simplified Tai Chi training program provided through the Tele-TC platform. This protocol was designed for older adults and focused on a small set of essential Tai Chi movements that can be performed safely and repetitively in a smooth and continuous way.
Participants were asked to practice the training exercises at home for 45–60 minutes per session, at least three times per week. In addition to practicing on their own using the tablet application, participants attended two or more live training sessions via Zoom with a certified Tai Chi instructor, each lasting up to 45 minutes.
The Tele-TC program was based on a simplified version of Cheng Man-Ch'ing's Yang-style short form and included six main Tai Chi movements. The training also included an additional set of traditional warm-up exercises designed to loosen the body, help with relaxation, increase breathing awareness, and encourage mindful movement.
Chairs were included in the protocol for some seated warm-up exercises and could also be used for stability or rest if needed, during the execution of more difficult exercises. The intervention focused on slow, controlled weight shifts, coordinated upper and lower limb movements, and postural alignment.
Training data including practice frequency, duration of the sessions, adherence to the protocol, and performance measures were automatically recorded through the Tele-TC platform.
The Tai Chi protocol included the following core movements:

1. Raising the Power

This exercise involves coordinated movements of the upper and lower body and helps improve joint mobility and balance. It mainly engages the ankles, knees, hips, and the torso, while promoting coordinated arm and leg movements.
As shown in Fig. 2.2, participants begin in a neutral standing posture with the feet about shoulder-width or hip-width apart. From this position, the knees are gently flexed while the arms are slowly raised in front of the body. The wrists are kept relaxed and slightly flexed, with the hands placed about shoulder-width apart and the fingers pointing downward in a natural manner. When the wrists reach approximately shoulder height, the palms gradually open and the fingers extend while the shoulders and elbows remain relaxed.
The elbows then sink downward, allowing the forearms and hands to move down smoothly along the front of the body. During this downward movement, the legs slowly extend again but without locking the knees. The movement is repeated in a continuous and controlled way, coordinating the knee flexion with the upward motion of the arms and the leg extension with the downward motion.
Throughout this exercise, the spine remains straight and elongated, the shoulders relaxed, and the movement is performed smoothly with a constant rhythm. The body moves smoothly as one single unit, with both feet remaining in stable contact with the ground [36].



Figure 2.2: Raising the Power exercise.

2. Push and Withdraw

This exercise introduces a forward and backward weight shift between the two legs and aims to improve dynamic balance, lower-limb strength, and movement coordination. It combines lower-body weight transfer with coordinated arm and trunk movements.
As shown in Fig. 2.3, participants begin from a standing position with the feet approximately shoulder-width apart. From this position, the weight is shifted on

one leg while the opposite foot is turned slightly outward. A step forward is then performed with the other foot to reach a bow stance, maintaining a comfortable distance between the two feet. In this position, the front knee is gently flexed while the back leg supports part of the body weight, with the trunk facing forward.
The arms are raised in front of the body to about shoulder level, with the elbows slightly flexed and the hands open in a relaxed pushing position. From here, the movement continues with a smooth backward weight shift onto the back leg, together with a circular downward motion of the arms. The weight is then transferred forward again onto the front leg while the arms rise again and return to the initial pushing position.
The movement is performed continuously and in a controlled way, coordinating the forward and backward weight transfer with the circular motion of the upper limbs. Throughout the exercise, a straight posture is maintained, with relaxed shoulders and stable contact of the feet with the ground. The same sequence of movements can be repeated on the opposite side, as well [36].

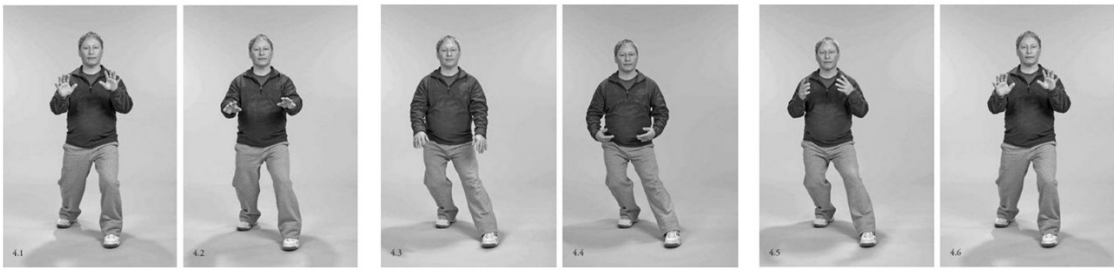


Figure 2.3: Push and Withdraw exercise.

3. Wave Hands like Clouds

This exercise focuses on lateral weight shifting performed together with trunk rotation and coordinated movements of the arms, with the aim of improving balance, posture control, and coordination of the entire body. The movement mainly involves the hips, trunk, shoulders, and upper limbs, while requiring controlled transfer of body weight from one leg to the other.
As shown in Fig. 2.4, participants begin in a standing posture with the feet parallel and approximately shoulder-width apart. The movement starts by shifting the body weight toward one leg while the trunk and waist rotate slightly toward the same side, keeping both feet in contact with the ground. The weight is then transferred to the opposite leg, together with a slight rotation of the trunk in the opposite direction.
Circular arm movements are then added to the pattern of the lower body. One

arm moves across the front of the body along a smooth curved path, approximately between waist and chest height, while the other arm follows the movement in coordination. The lateral transfer of body weight is synchronized with the motion of the upper limbs, creating a continuous sequence involving the legs, trunk, and arms.
Throughout the exercise, the movement is performed in a slow and controlled way, maintaining an upright posture, relaxed shoulders, and continuous and stable contact of the feet with the floor. The sequence is performed on both sides with smooth and non-stop transitions between phases [36].



Figure 2.4: Wave Hands like Clouds exercise.

4. Grasp the Sparrow's Tail

This exercise is characterized by forward and backward weight shifting with trunk rotation and coordinated movements of both arms, with the goal of improving balance, joint mobility, and total body coordination. The movement mainly involves the hips, trunk, shoulders, and arms, and highlights the role of the torso in guiding the sequence.
As shown in Fig. 2.5, participants begin in a bow stance, with one foot placed forward and the other slightly turned outward. From this position, the body weight is first shifted toward the front leg while the trunk is oriented forward. The trunk then rotates slightly toward the same side, after which the weight is gradually moved backward onto the opposite leg. During this shift, the torso rotates in the opposite direction, and the movement is then reversed to return to the starting bow stance.
Coordinated arm movements are performed together with the pattern of the legs. One arm is positioned in front of the chest in a pushing gesture, while the other is flexed in a rounded position. As the trunk rotates and the body weight shifts backward, both arms move downward in a wide circular pattern and then return upward as the weight is shifted forward again. The coordinated action of the arms, trunk, and lower limbs produces a continuous and flowing sequence of movements.

The exercise is performed slowly and gently, maintaining a correct posture, relaxed shoulders, and stable and continuous contact of both feet with the ground. The body moves as a whole unit, with transitions between forward and backward shifts executed in a controlled and fluid manner [36].



Figure 2.5: Grasp the Sparrow's Tail exercise.

5. Brush Knee and Twist Step

This exercise consists of slow forward steps and coordinated arm movements, aiming to improve balance, leg strength, and integration of upper and lower body movements. The movement involves the hips, legs, arms and the trunk, while emphasizing control and stability during the steps.
As shown in Fig. 2.6, participants begin the execution with one foot forward, toes rotated slightly outward, and the weight mostly on that leg. The other foot stays close to the forward foot for balance, with little weight on the corresponding leg. The arms are positioned in front of the body, holding a small imaginary sphere gently, with one hand placed on top of the other.
A step forward into a bow stance with the back foot is then performed, putting most of the weight onto the front leg. As the body weight moves forward, hips and torso are turned toward the direction of the step. The arm on the side of the stepping leg moves down in front of the thigh, while the other arm rises to shoulder height in a soft push position.
After completing the step, participants balance briefly on the front leg, bringing the back knee close to the front knee, then returning the back foot to its starting position. The weight is shifted back to the rear leg while the arms move following a circular pattern back to the starting "sphere" position.
This sequence of movements can be performed on the opposite side as well, stepping with the other leg forward and repeating the same coordinated weight shift and arm movements.
The exercise is performed slowly and steadily, keeping the spine straight and the shoulders relaxed. The feet stay firmly planted, while the body moves as a

coordinated unit. Forward steps are coordinated with circular arm movements, maintaining balance and control throughout the exercise [36].



Figure 2.6: Brush Knee and Twist Step exercise.

6. Golden Rooster Stands on One Leg

This exercise focuses on single-leg balance, coordinated arm movements, and controlled weight transfer, to improve stability, leg strength, and whole-body coordination. It mainly involves the hips, knees, and ankles while focusing on correct alignment of the trunk.
As shown in Fig. 2.7 [37], participants begin with the feet close together and the arms relaxed along the body. The left leg then steps forward, followed by the right leg, until both feet are close together again. Then, the right leg and right hand are raised at the same time, while the left hand descends with the palm facing backward. Subjects stand on the right leg for a moment, before lowering the previously raised knee.
During the lowering phase of the right leg, the right hand moves in a circular pattern toward the right thigh, while keeping the elbow stable. As both hands are lowered, the body weight slowly moves backward. The exercise ends by moving the left leg backward so that the feet return to the starting position. The same sequence is then performed on the opposite side.
Throughout the exercise, the spine remains upright, the shoulders relaxed, and the feet stay firmly on the ground whenever both are in contact with the floor. The movements are performed slowly and smoothly, with coordinated action of the arms, legs, and trunk to maintain correct balance and control of the limbs [38].



Figure 2.7: Golden Rooster Stands on One Leg exercise.

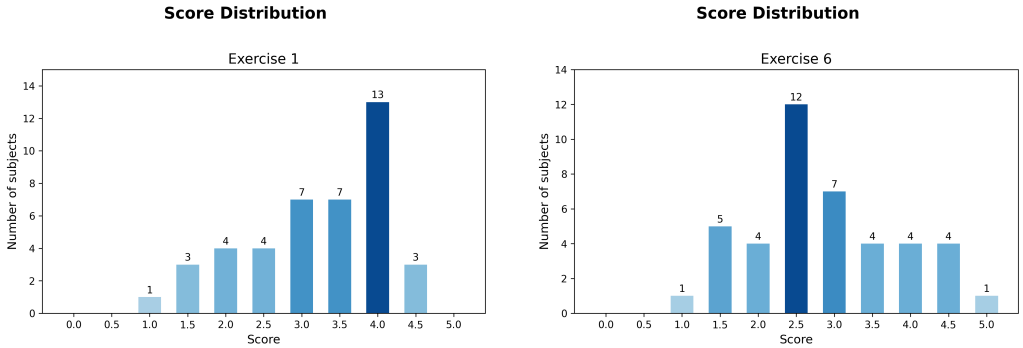
2.1.3 Recording Setup

Participants' movement executions were recorded using wearable sensors during all assessment visits. Each participant was equipped with five wearable inertial sensors, placed on the trunk and on each limb using elastic bands.
During the visits, participants followed the instructions provided by the Tele-TC platform while their movements were recorded. Sensor data were used to quantify Tai Chi performance metrics including movement quality, adherence to the protocol, and safety-related measures.
All sessions were also video recorded via Zoom to provide a reference for offline analysis and for evaluating Tai Chi proficiency.
In addition to the wearable sensors, participants wore a GeneActiv wrist-worn activity monitor for approximately one week following the baseline visit and again during the week before the final visit, in order to measure daily physical activity levels. In this thesis, only the video recordings were used for the analysis; the wearable sensor data were collected as part of the broader study but are not analyzed here.

2.1.4 Scoring System

A Tai Chi expert evaluated in detail the videos recorded at baseline and at the end of the protocol (week 12). Each video was reviewed carefully, with attention to the motor behaviors shown by each subject. During these review sessions, any deviations from the correct execution or any incorrect movements were noted, resulting in a global evaluation of motor performance quality.
The multiple repetitions (nine) of each subject were evaluated collectively, rather than individually, allowing for an overall evaluation of the participants' execution throughout the set of repetitions.
Following this review, each subject received a single overall score for their complete set of repetitions. Scores ranged from 0, indicating an extremely poor performance, to 5, representing a perfect execution, using 0.5-point increments to allow for a

finer distinction between performance levels. In addition to this overall proficiency score, the specific deviations noted during the review were used to define the binary motor-behavior labels (presence/absence) described in Section 2.3. The proficiency score serves as the target for the regression task, while the binary labels serve as the targets for the motor-behavior classification. The resulting distribution of proficiency scores across subjects is shown in Figure 2.8 for both exercises.



(a) Exercise 1 (b) Exercise 6

Figure 2.8: Distribution of proficiency scores across subjects for Exercise 1 (a) and Exercise 6 (b).

2.1.5 Videos Excluded from the Analysis

Before conducting the analysis, a data quality control was performed to ensure the reliability of the results. Some subjects were excluded from the study based on predefined criteria. The primary reasons were related to improper camera framing, which resulted in the subject being partially or entirely out of the field of view. Following a visual inspection of the meshes reconstructed using the SMPL model, as well as the time-varying 2D projection of keypoints, additional videos were excluded due to poor quality and poor lighting conditions, which did not allow for sufficient mesh reconstruction accuracy. Similarly, a small number of videos were discarded due to significant lag, which prevented a reliable biomechanical analysis. In Figure 2.9, there is an example of a video excluded from the analysis, due to the subject being positioned out of the field of view of the camera. The final number of videos included in the analysis is summarized in Table 2.2.



Figure 2.9: Example of a video excluded from the analysis.

Table 2.2: Final Dataset.

Exercise	Original number of videos	Final number of videos, after exclusion	Final number of repetitions
Raising the Power	55	42	378
Golden Rooster Stands on One Leg	55	42	378

2.2 Biomechanical Modeling Pipeline

2.2.1 SMPLest-X-based Pose Estimation

The proposed pipeline processes a video sequence frame-by-frame, where each video corresponds to a full execution of a specific exercise, including multiple repetitions. Human detection is performed in each frame to localize the subject of interest. The detected region is then used to crop and normalize the input image before feeding it into SMPLest-X [35] for 3D human body reconstruction. The model outputs a 3D mesh, joint positions, and SMPL-X pose parameters. These results are projected back onto the original image space for visualization, allowing a qualitative assessment of the reconstructed motion and the model's ability to consistently follow the subject across the entire exercise sequence. In addition, the estimated 2D/3D keypoints and pose parameters are stored for subsequent quantitative analysis. Each video is processed at a fixed frame rate of 30 fps.

2.2.2 Keypoint and Pose Parameter Extraction

The SMPL model is controlled by pose parameters (θ), which define the 3D rotations of the body joints in a kinematic tree (in axis-angle format). The shape parameters (β) determine a 3D body template, from which a predefined set of joint locations is obtained via a linear regressor applied to the template mesh vertices; then these joints are articulated along the kinematic tree according to the pose parameters, yielding both the final joint positions and, through linear blend skinning, the posed 3D body mesh. These joint positions are used as 3D keypoints (origin in the pelvis) in our pose estimation framework. In this work, both the 3D keypoints representing body joints and the pose parameters have been used.

2.2.3 Video Segmentation - GUI

To enable manual segmentation of full patient exercise videos into movement repetitions, a graphical user interface (GUI) was developed in Python to ensure a fast, intuitive, and reproducible annotation workflow. The interface allows the user to select a video file through a file dialog and visualize it using OpenCV. The video is loaded frame-by-frame, and the current frame index is continuously displayed on the screen together with user instructions for navigation and annotation.
As shown in Figure 2.10, the tool allows the user to navigate through long videos using both small and large jumps, making it easier and faster to review the recording. Frame-by-frame navigation is enabled using the keys “a” (previous frame) and “d” (next frame), while larger jumps are supported using “z” and “f” for ± 10 frames, and “e” for a forward jump of 100 frames. This navigation strategy allows the user to quickly explore the video while maintaining the possibility of precise frame-level control. The key “q” is used to exit the visualization.
During video inspection, the beginning and end of an entire exercise execution are manually annotated using the key “s”, which stores the current frame index and its corresponding timestamp (in seconds, computed using the video frame rate). An “undo” functionality is implemented through the key “u”, allowing the removal of the most recently saved timestamp in case of annotation errors. All selected timestamps are saved in an organized format, including both the frame number and the corresponding time in seconds.
After completion of the annotation process, all timestamps are exported and saved into an Excel (.xlsx) file in the same directory as the input video. The exported timestamps were later used as temporal boundaries to segment the full motion recordings into individual repetitions. In particular, they are used as reference markers for dividing the extracted keypoint trajectories into repetition-specific segments, enabling further biomechanical analysis of each movement cycle independently.



Figure 2.10: GUI for manual video segmentation.

Definition of repetition boundaries

To ensure consistent segmentation across subjects, the start and end of each repetition were defined based on specific postural criteria.
For Exercise 1, the beginning of a repetition was identified as the moment when the subject was in a relaxed upright position, with the arms resting alongside the body and the knees extended but not actively engaged, just before executing the arm elevation movement. The end of the repetition, and consequently the start of the following one, was defined as the instant when the arms returned to a relaxed position alongside the body and the knees were again extended, corresponding to the recovery of the same relaxed upright starting posture. An example of the starting and ending position for Exercise 1 is shown in Figure 2.11.



Figure 2.11: Example of the starting and ending position for Exercise 1.

For Exercise 6, the beginning of a repetition was identified as the moment when the subject had already assumed the initial stance, with one foot positioned forward and

the other backward, the body weight evenly distributed, and the arms held away from the body, before any lifting of the arm or the corresponding leg had started. The end of the repetition was defined as the instant following the lifting phase, when the limbs returned to a neutral transition position between two consecutive movements. This moment corresponds to the transition phase in which the subject completes one repetition and prepares to initiate the next. An example of the starting and ending position for Exercise 6 is shown in Figure 2.12.



Figure 2.12: Example of the starting and ending position for Exercise 6.

2.2.4 Keypoint Normalization and Filtering

For each subject and repetition, 3D keypoint coordinates were normalized using a single global scaling factor to account for inter-subject variability in body size. The start of the first repetition was identified from timestamp data and converted into the corresponding frame index, using the frame rate of 30 fps. From this point, the first five frames, assumed to represent a quasi-static posture, were selected. Within this short temporal window, the Euclidean distance between the pelvis and neck joints was computed for each frame, and the mean value was used as the normalization factor. This factor was then applied to all joints and all frames by dividing the corresponding 3D coordinates (x, y, z). As a result, the full sequence of keypoints representing the motion was rescaled using an anatomical reference length extracted from the initial posture.

The normalized keypoints and the pose parameters were then temporally filtered to reduce high-frequency noise while preserving the underlying motion dynamics. Specifically, each component of the keypoint coordinates (x, y, z) and of the pose parameters was processed independently using a fourth-order low-pass Butterworth filter with a cutoff frequency of 10 Hz. The filtering was applied along the time dimension for each joint. The sampling frequency was set equal to the SMPL frame rate (30 fps). This procedure resulted in a smoothed representation of both the keypoint and pose parameter time series.

2.3 Motor Behavior Analysis

Correct execution of Tai Chi exercises plays a key role in achieving the desired therapeutic outcomes. However, patients often perform incorrect movement strategies during training. For this reason, the analysis focuses on the identification and characterization of the most common deviations from correct motor execution, with the objective of providing feedback strategies that can improve motor performance. For each video, deviations from the correct execution of the movements were annotated. These movement patterns were investigated in order to identify the most common execution errors and provide meaningful feedback to the subjects, enabling them to recognize and correct mistakes during future training sessions. A great range of movement deviations was observed at first; however, only the most frequent ones were considered for the subsequent analysis. Each deviation was labeled as "presence" if the subject exhibited the incorrect movement pattern, and "absence" otherwise. The total number of Motor Behaviors considered for each exercise is presented in Table 2.3.

Table 2.3: Total number of Motor Behaviors.

Exercise	**Number of Motor Behaviors analyzed**
Raising the Power	6
Golden Rooster Stands on One Leg	7

For Exercise 1, Raising the Power, six movement deviations were analyzed. Their occurrence within the study population is presented in Table 2.4:

- **Motor Behavior 1: Wrist Height** – at peak elevation, the vertical position of the wrists should not exceed shoulder height. A lower elevation is considered acceptable, particularly in patients with some sort of motor impairment. Subjects for whom the majority of repetitions exhibited a maximum wrist height above shoulder level were labeled as "presence".
- **Motor Behavior 2: Synchronization Knees-Wrists** – during movement execution, knee flexion should begin simultaneously with the upward motion of the wrists. Peak wrist elevation should correspond in time with the maximum knee flexion. Subsequently, the knees should extend while the wrists lower, returning to the initial position. Subjects who exhibited a clear lack of synchronization, such as initiating knee flexion before wrist elevation, or vice versa, were labeled as "presence".

- **Motor Behavior 3: Feet Width** – throughout the whole execution of the exercise, the feet should remain in contact with the ground and be positioned at a width at least equal to that of the hips, and ideally comparable to shoulder width. Subjects who exhibited an insufficient stance width were labeled as "presence".

- **Motor Behavior 4: Wrists-Shoulders Parallelism** – during the upward movement, the wrists should maintain a width consistent with that of the shoulders, keeping a parallel trajectory. At peak elevation, the wrists should reach approximately shoulder height, and the same alignment should be maintained during the downward phase of the movement. Subjects who exhibited an excessive wrist width relative to shoulder width were labeled as "presence".

- **Motor Behavior 5: Feet Parallelism** – throughout the execution of the exercise, the feet should remain firmly in contact with the ground and maintain a stable position. They should be kept oriented parallel to each other, without external rotation, in order to ensure proper lower-limb alignment and balance. This alignment should be preserved across all phases of the movement. Subjects who exhibited a noticeable external rotation of the feet during the execution were labeled as "presence".

- **Motor Behavior 6: Arm Three-Dimensionality** – during the execution of the movement, the upper limbs should follow a smooth and continuous trajectory in the space in front of the body, rather than being restricted to a single plane. From an initial neutral standing posture, the arms are elevated with a coordinated motion involving shoulders and elbows, maintaining a natural degree of flexion at the wrists and a comfortable spacing between the hands. As the arms reach shoulder level, the movement should preserve a rounded and forward-oriented configuration, rather than collapsing into a vertical or overly rigid pattern. During the following downward phase, the arms should descend along a similarly controlled and spatially consistent path, maintaining continuity of motion. Subjects who exhibited a planar or reduced three-dimensional arm trajectory were labeled as "presence".

Other movement deviations were identified, although they involved only a small portion of the study population and were not included in the main analysis. These included an incorrect motion of the fingers during the ascending, peak, and descending phases, often characterized by excessive stiffness, whereas the fingers should follow a more fluid and natural movement in coordination with the wrists. A limited number of subjects exhibited knee valgus during flexion, while one subject lifted the heels during execution, despite the requirement to maintain full foot

contact with the ground.
Some participants showed noticeable stiffness in the neck and shoulders, affecting the overall fluidity of the movement. In a few cases, the wrists were positioned posterior to the hips at the beginning or end of the lifting phase. Additionally, some subjects did not raise the wrists symmetrically, showing vertical displacement asymmetries between the two sides.
Most subjects, if not all, tended to fully extend the knees at the end of the movement, locking their joints. However, proper execution requires the knees to remain slightly flexed throughout, avoiding complete extension. This aspect proved particularly challenging, and only a minority of participants were able to maintain the correct slight flexion. Finally, a few subjects did not perform any noticeable knee flexion, keeping the legs extended for the entire duration of the movement.

Table 2.4: Frequency of Motor Behavior occurrence in the study population, Raising the Power.

MB 1	MB 2	MB 3	MB 4	MB 5	MB 6
17 subjects	20 subjects	11 subjects	9 subjects	26 subjects	20 subjects

Figure 2.13 provides a comprehensive summary of the dataset for Exercise 1, Raising the Power, illustrating each participant's motor behavior profile sorted by their respective scores in ascending order.

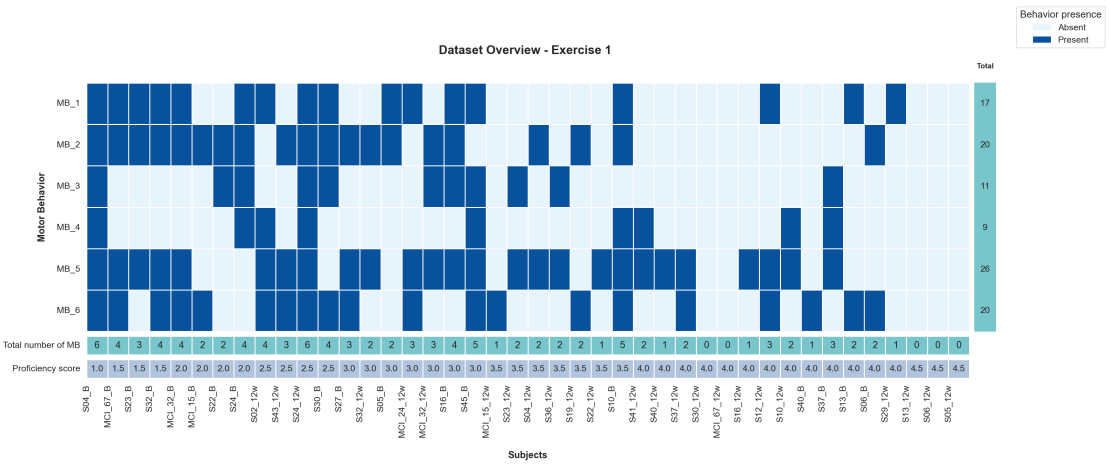


Figure 2.13: Dataset summary for Exercise 1, Raising the Power.

For Exercise 6, Golden Rooster Stands on One Leg, seven movement deviations were analyzed. Their occurrence within the study population is also reported in

Table 2.5:

- **Motor Behavior 1: Wrist Height** – at peak elevation, the vertical position of the raised wrist should reach approximately shoulder height and should not exceed it. This position ensures proper upper-limb control and coordination with the lower-limb movement during the single-leg stance. Subjects for whom the majority of repetitions exhibited a wrist height above shoulder level were labeled as "presence".

- **Motor Behavior 2: Limb Side Coordination** – during movement execution, the raised arm and the lifted leg should belong to the same side of the body (i.e., right arm with right leg, or left arm with left leg). This coordinated pattern is essential to ensure correct execution of the exercise and proper balance control. Subjects who exhibited inconsistent side coordination, such as raising the arm and the contralateral leg, were labeled as "presence".

- **Motor Behavior 3: Initial Foot Stance Width** – prior to knee elevation, the feet should be positioned at an adequate distance, with at least one visible step between them. This wider stance is required to ensure stability during the subsequent single-leg movement. Subjects who did not adopt a sufficiently wide foot stance at the beginning of the exercise were labeled as "presence".

- **Motor Behavior 4: Wrist–Elbow Alignment of Raised Arm** – during arm elevation, the wrist of the raised arm should remain vertically aligned above the elbow, maintaining a perpendicular alignment that ensures correct upper-limb posture and control. Subjects who exhibited misalignment between the wrist and elbow during the movement were labeled as "presence".

- **Motor Behavior 5: Ankle–Knee Alignment of Raised Leg** – during the execution of the movement, the ankle of the raised leg should remain vertically aligned above the knee, maintaining a perpendicular alignment to ensure proper lower-limb control during single-leg stance. Subjects who exhibited misalignment between the ankle and knee of the raised leg were labeled as "presence".

- **Motor Behavior 6: Hip–Knee Height Alignment of Raised Leg** – during single-leg stance, the knee of the raised leg should be elevated to approximately the same height as the hip, ensuring adequate range of motion. Subjects who raised the knee insufficiently, resulting in a knee position clearly below hip level, were labeled as "presence".

- **Motor Behavior 7: Elbow–Knee Alignment of Raised Side** – during the execution of the movement, the elbow of the raised arm should remain

aligned with the knee of the lifted leg, both oriented in the same spatial direction. This alignment ensures coordinated whole-body control and correct contralateral motor pattern execution. Subjects in which the elbow and knee of the raised side were not aligned, were labeled as "presence".

Other movement deviations were identified during the execution of Golden Rooster Stands on One Leg; however, they were not included in the main analysis due to their limited observability or reduced relevance for the final scoring.
In particular, the position of the non-active contralateral arm was considered. Specifically, the non-active arm should remain positioned laterally to the body, in front of the hip, and should not move behind the trunk during the execution of the movement. However, this movement characteristic was often difficult to assess due to partial occlusion of the contralateral arm by the trunk of the subject and the resulting limitation in accurately estimating its depth position from a monocular camera view. Moreover, it was frequently observed that some subjects tended to move the arm posteriorly only during repetitions performed on one side, and not when engaging the other side of the body. This subtle deviation from the correct movement execution was observed in the majority of videos analyzed. For all the reasons described above, this motor behavior was not included in the final analysis.
Another very subtle movement component concerned foot dorsiflexion of the raised leg. According to expert evaluation, this aspect was considered a fine detail with minimal influence on the final score. The foot should not be actively held in dorsiflexion, but rather remain relaxed or plantar-flexed (pointing downward).
However, SMPL does not reliably capture fine foot articulation, in particular toe and foot segment motion, making this feature difficult to assess consistently. For this reason, and given its marginal relevance for scoring, it was not included in the analysis.
Balance during the single-leg stance was also considered as a candidate motor behavior. However, the degree of balance impairment showed relevant variability both across subjects and repetitions of the same subject, making it difficult to assign a consistent binary presence/absence label. Moreover, balance instability appeared to influence the execution of several other motor behaviors analyzed for this exercise, acting as a confounding factor rather than an independent, well-defined deviation. For all these reasons, balance was not considered as a separate motor behavior in the final analysis.

Table 2.5: Frequency of Motor Behavior occurrence in the study population, Golden Rooster Stands on One Leg.

MB 1	MB 2	MB 3	MB 4	MB 5	MB 6	MB 7
14 subjects	5 subjects	7 subjects	12 subjects	9 subjects	17 subjects	19 subjects

Figure 2.14 provides a comprehensive summary of the dataset, Exercise 6, Golden Rooster Stands on One Leg, illustrating each participant's motor behavior profile sorted by their respective scores in ascending order.

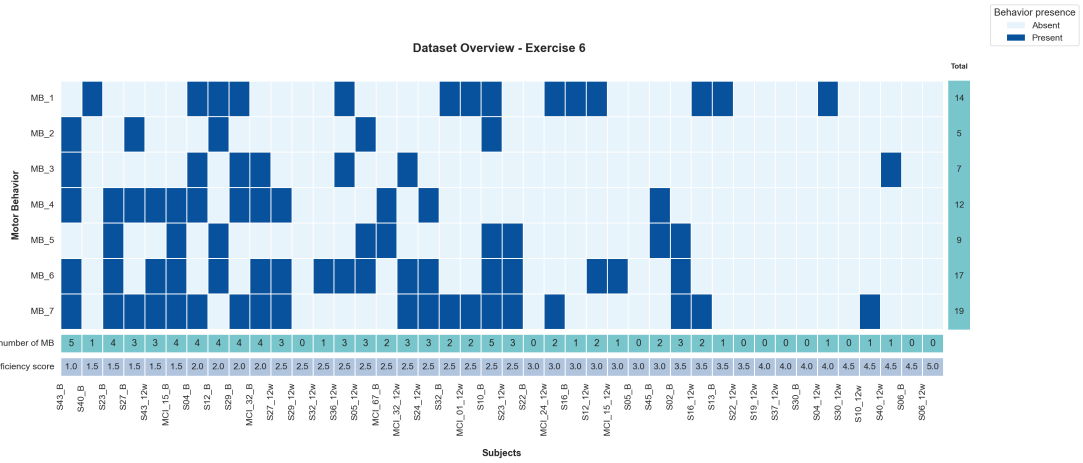


Figure 2.14: Dataset summary for Exercise 6, Golden Rooster Stands on One Leg.

2.3.1 Repetitions Analysis

Before the classification phase, an initial data analysis was performed to determine the optimal strategy for handling the multi-repetition structure of the dataset. Each subject performed nine repetitions per exercise; however, both the presence of incorrect motor behaviors and the overall proficiency score (continuous 0-5) were assigned to the entire execution rather than to individual repetitions.
Consequently, we needed to determine whether aggregating the repetitions, through pooling techniques, was appropriate, particularly in cases of high intra-subject variability. To investigate this, for each motor behavior, the features of individual repetitions were plotted to assess intra-subject variance across the nine repetitions. In these plots, individual repetitions are marked by circular markers, while the aggregated values, obtained through mean pooling, are represented with square

markers. Subjects are color-coded according to the presence or absence of the motor behavior error.
As illustrated in the examples in Figures 2.15 and 2.16 (representing two features for Exercise 1, Motor Behaviors 1 and 4), the repetitions are generally consistent with one another. Qualitatively, it can be observed that the intra-subject variance is minimal compared to the distinct separation between the two classes.
To quantitatively support this observation, an intra-subject variability ratio was computed for each feature, defined as:

$$\text{ratio} = \frac{\sigma_{\text{within}}}{\Delta_{\text{between}}} = \frac{\sqrt{\frac{1}{N}\sum_{i=1}^{N}\sigma_i^2}}{|\mu_{\text{YES}} - \mu_{\text{NO}}|} \tag{2.1}$$

where $\sigma_{\text{within}} = \sqrt{\frac{1}{N}\sum_{i=1}^{N}\sigma_i^2}$ is the pooled within-subject standard deviation across the nine repetitions, $\Delta_{\text{between}} = |\mu_{\text{YES}} - \mu_{\text{NO}}|$ is the absolute difference between the mean values of the two classes, N is the number of subjects, and σ_i^2 is the variance of the nine repetitions of subject i.
All features were standardized prior to computation using z-score normalization. A ratio below 1 indicates that intra-subject variability is smaller than the between-class separation.
As shown in Figures 2.17 and 2.18, the majority of features exhibit a ratio well below 1, confirming quantitatively that repetitions are consistent within subjects. Given these observations, we opted against aggregation. Instead, the label assigned to the overall performance was applied to each individual repetition, increasing the sample size available for training.

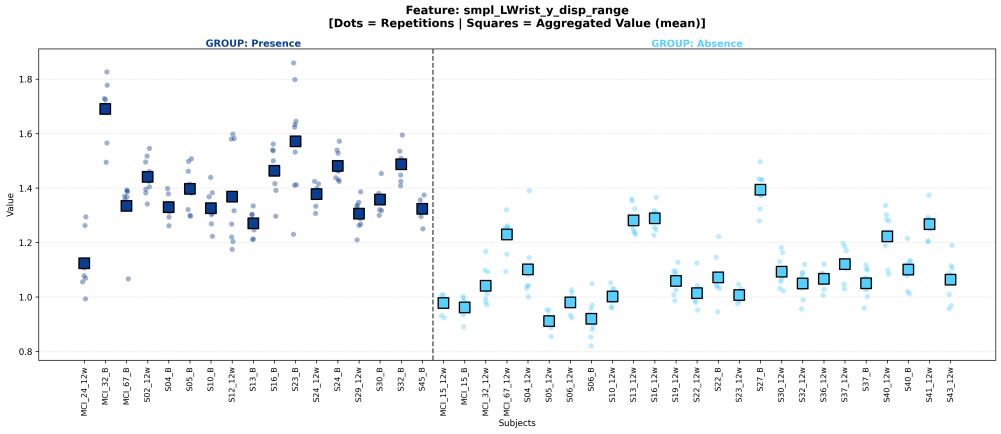


Figure 2.15: Distribution of an example feature for Exercise 1, Motor Behavior 1.

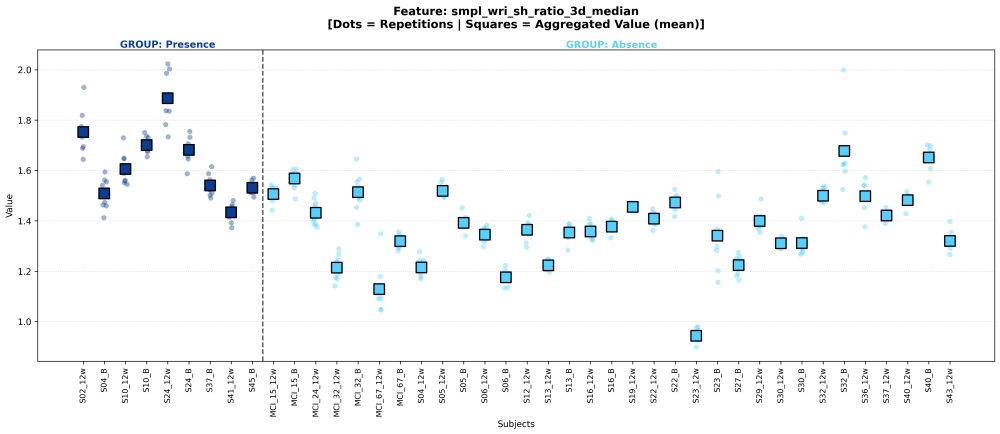


Figure 2.16: Distribution of an example feature for Exercise 1, Motor Behavior 4.

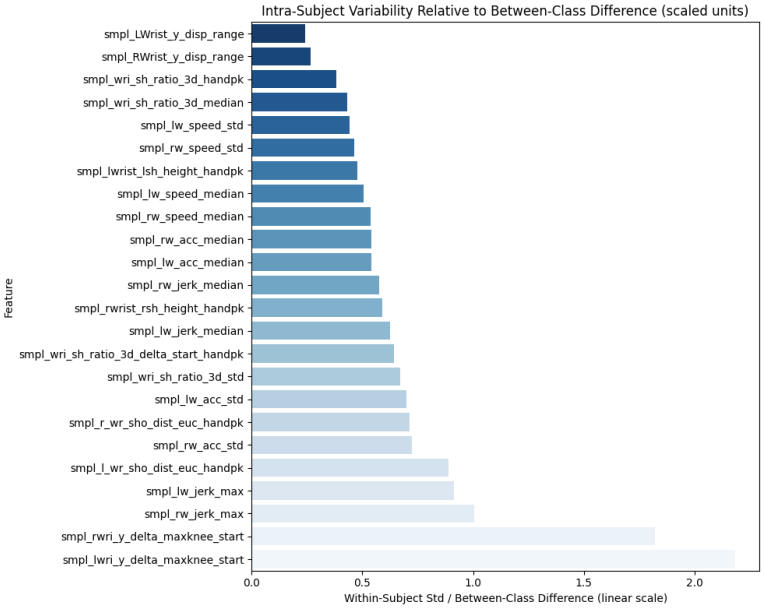


Figure 2.17: Intra-subject variability ratio for Exercise 1, Motor Behavior 1.

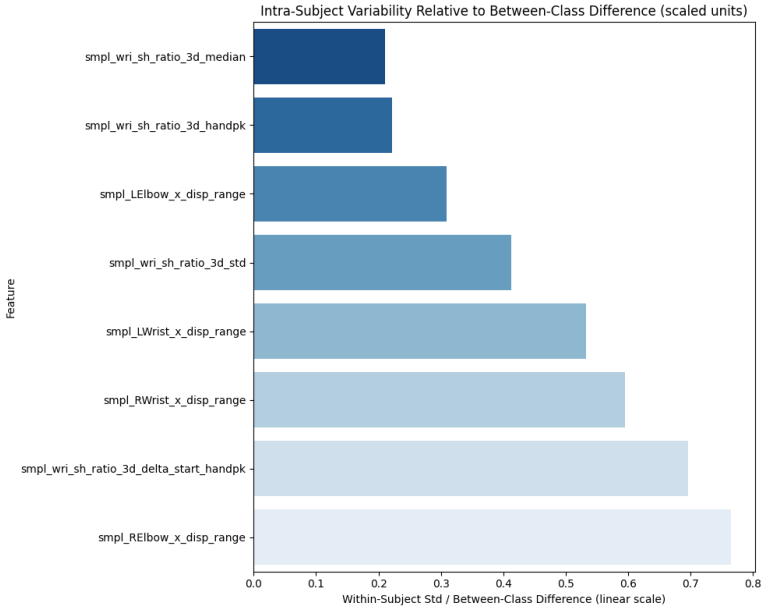


Figure 2.18: Intra-subject variability ratio for Exercise 1, Motor Behavior 4.

2.4 Modeling Framework

2.4.1 Feature Extraction

To capture the essential biomechanical characteristics of the analyzed Tai Chi exercises, a set of motion-specific features was created to describe the distinct motor behaviors observed.
We computed static features by identifying key anatomical configurations at specific, movement-relevant instants, such as the peak elevation of the wrist. Then we extracted some dynamic features, which characterize the movement as it evolves over time. These are derived from the trajectory of joints over defined intervals or across the full duration of a repetition, using statistical descriptors such as median, standard deviation, and range, as well as scalar quantities like correlation coefficients and spatial distances.
Outliers were identified per feature using a Tukey-style rule applied to the 5th–95th percentile range, with values beyond 1.5× this range from either bound flagged as outliers. Outlier repetitions were replaced with the subject's within-trial mean.
For Exercise 6, the term "active" refers to the raised limb (arm or leg) involved in the single-leg stance, as opposed to the supporting side.
Given that each Tai Chi exercise presents a unique biomechanical structure, the feature sets are exercise-specific, ensuring that the extracted metrics accurately

capture the physiological aspects of the movement under analysis.
A comprehensive list of the features extracted and utilized for each motor behavior is provided in Appendix A.

2.4.2 Feature Selection

For the feature selection for the binary motor behavior analysis, we relied on a manual selection process based on prior domain knowledge. We specifically focused on features that best captured the characteristics of the different motor behaviors considered in this study. For more complex movements, we naturally selected a larger number of features to ensure the model could capture their nuances. This approach allowed us to prioritize variables with clear clinical or biomechanical meaning, ensuring that the model remains interpretable while focusing on the most informative aspects of the data.
For the proficiency assessment task, two feature selection approaches were tested. The first was an **automated nested feature selection** using the mRMR (Minimum Redundancy Maximum Relevance) algorithm [39]. At each cross-validation fold, mRMR was run on the aggregated training data to identify the top-K features (K $\in \{4, 5, 6, 7, 8\}$), balancing relevance to the proficiency score with low redundancy among selected variables. The best K was determined by selecting the configuration that achieved the lowest validation MAE.
The second approach was a **manual selection** based on prior domain knowledge, selecting variables with clear biomechanical meaning relevant to the exercise being analyzed. Features were picked to capture the movement qualities a Tai Chi expert would focus on when assigning a proficiency score. In addition to the biomechanical features, binary ground-truth labels representing the presence or absence of specific incorrect motor behaviors were included as additional features, providing the model with explicit information about movement quality and level of performance.

2.4.3 Dimensionality Reduction - UMAP

Dimensionality reduction techniques are essential for managing high-dimensional datasets without sacrificing structural information. Mapping data into a lower-dimensional space not only improves visualization and reduces noise, but also facilitates subsequent feature extraction, improving the performance of tasks such as classification and clustering [40].
Traditional linear approaches, such as Principal Component Analysis (PCA), remain largely used due to their simplicity and computational efficiency. Their main drawback, however, is the inability to capture the nonlinear dynamics typical of real-world datasets.
To address this limitation, nonlinear frameworks, like t-SNE, UMAP, Isomap, and

Locally Linear Embedding (LLE) are preferred when the goal is to preserve the data's intrinsic geometry [41].

In this context, **Uniform Manifold Approximation and Projection (UMAP)** is a recent technique that stands out for its unique capacity to preserve both local and global relationships. It is a nonlinear dimensionality reduction method based on manifold learning and topological data analysis, which constructs a graph representation of the high-dimensional data and then computes a low-dimensional embedding that respects the original neighborhood structure.

UMAP assumes that the data lie on a low-dimensional manifold embedded within a higher-dimensional space, where local connectivity can be modeled through neighborhood relationships. Based on all these assumptions, it builds a fuzzy topological representation and optimizes the embedding to mirror the data structure as accurately as possible [41].

The UMAP algorithm depends on a set of hyperparameters that influence the structure of the low-dimensional embedding. The hyperparameters are [41]:

- **n_neighbors** – number of neighboring samples used to approximate the local structure of the data. Smaller values focus on local relationships, while larger values capture broader global patterns. Default value: 15.

- **n_components** – defines the number of dimensions in the reduced space. Default value: 2.

- **min_dist** – minimum distance between points in the low-dimensional space. Lower values lead to tighter, more compact clusters, while higher values preserve a more uniform distribution of points. Default value: 0.1.

- **metric** – specifies the distance metric used to compute relationships within the ambient space of the input data.

In this study, UMAP is employed exclusively as an exploratory tool for qualitative visualization. Our goal is not to perform dimensionality reduction for further processing, but to visually inspect the capability of the selected features to discriminate between the different classes in a low-dimensional space ($n_components = 3$). This allows for an intuitive assessment of how clearly our subjects are grouped before moving on to the subsequent analysis. The UMAP parameters used in this analysis are presented in Table 2.6.

Table 2.6: UMAP parameters used in this analysis.

Parameter	Value
n_neighbors	100 (Motor Behavior), 80 (Proficiency)
n_components	3
min_dist	0.6
metric	euclidean

2.4.4 Modeling Algorithms

Logistic Regression

For the binary Motor Behavior classification task, we first tested a **Logistic Regression** (LR) model as a simple, interpretable baseline before moving to more complex methods.
LR is a classification algorithm that belongs to the family of Generalized Linear Models (GLM). While standard linear regression models are designed to predict continuous values, LR adapts this idea to binary classification by mapping a linear combination of input features into a probability range of $[0, 1]$ [42, 43].
The main mechanism of LR is the sigmoid function (also known as the logistic function), which maps real-valued numbers into the $(0, 1)$ interval, giving the characteristic S-shaped curve used for probabilistic output. Mathematically, for a set of input features X, the model calculates the probability p that an observation belongs to a specific class as follows:

$$p(X) = \frac{1}{1 + e^{-(\beta_0 + \beta_1 X_1 + \cdots + \beta_d X_d)}} \tag{2.2}$$

Each observation is treated as a Bernoulli trial, and the model estimates the outcome probability from the linear relationship between the predictors and the log-odds of the dependent variable.
Classification then follows a standard threshold of 0.50: an instance is assigned to the positive class when its predicted probability exceeds this value.
Hyperparameter tuning was performed using **RandomizedSearchCV** (10 iterations per fold), optimizing for Macro F1-score within the inner LOGO loop. Two hyperparameters were tuned: the regularization strength `C` and the penalty type, using a `liblinear` solver. The search ranges are summarized in Table 2.7. As for the Random Forest classifier, the `class_weight='balanced'` option was used to account for class imbalance in the motor behavior labels.

Table 2.7: Hyperparameter search space for the Logistic Regression model.

Hyperparameter	Values
C (regularization)	0.001, 0.01, 0.1, 1
penalty	L1, L2

Random Forest Classifier

Random Forest is an ensemble learning method introduced by Breiman [44] in 2001, based on the combination of multiple decision trees to improve prediction accuracy and robustness. Instead of relying on one single model, this method builds a collection of trees, each trained on a slightly different subset of the data, and aggregates their outputs to produce the final prediction.
The method relies on two main sources of randomness. First, each tree is trained on a bootstrap sample of the original dataset, obtained by sampling with replacement (a technique known as Bagging [45]). This process introduces variability among the trees and helps reduce overfitting. Second, at each split in the tree construction, only a random subset of the input features is considered. The best split is then selected among this subset, rather than among all the features available.
Each decision tree is typically grown to its maximum depth without pruning, following the standard CART (Classification and Regression Trees [46]) methodology. Single trees can have high variance and tend to overfit the training dataset, but the aggregation of multiple trees significantly improves generalization performance. In classification tasks, the final prediction is usually obtained through majority voting, while in regression problems it is computed as the average of the outputs of all the trees.
One of the key advantages of Random Forest is its ability to handle high-dimensional data and complex, non-linear relationships without requiring strong assumptions on the data distribution. Additionally, it is relatively resistant to overfitting due to the combination of multiple uncorrelated trees. This method also provides an internal estimate of the generalization error through out-of-bag samples, as well as measures of variable importance, which can be useful for feature selection and model interpretation [47]. Random Forests also work well even when only a few features are actually informative: their performance depends mainly on the number of relevant features, rather than on the total number of features in the input [48]. Thanks to all these properties, Random Forests are versatile and robust, able to handle mixed data types, high-dimensional inputs, and complex interactions among their variables. They are also computationally efficient and can be parallelized, making them widely used in both regression and classification tasks [44, 47].
Random Forest models depend on a set of hyperparameters that can be tuned to

optimize performance and control overfitting. The main hyperparameters include:

- **n_estimators** – number of trees in the forest. Default value: 100. Increasing this generally improves stability and accuracy but increases computational costs.
- **max_features** – number of features to consider at each split. Default values: `sqrt` for classification, `None` for regression. Options include `sqrt`, `log2`, or `None`. Choosing fewer features increases tree diversity and reduces overfitting.
- **max_depth** – maximum depth of each tree. Default value: `None` (trees grow until all leaves are pure or contain less than `min_samples_split` samples). A non-deep tree may underfit, while a very deep tree may overfit.
- **min_samples_split** – minimum number of samples required to split an internal node. Default: 2. Larger values prevent overfitting as they require more samples to create the split.
- **min_samples_leaf** – minimum number of samples required to create a leaf. Default: 1. Controls leaf size and smooths predictions.
- **max_leaf_nodes** – maximum number of leaf nodes. Default value: `None` (no limit). Helps control the complexity of the tree.
- **bootstrap** – whether to use bootstrap sampling (bagging) to train each tree. Default: `True`.
- **max_samples** – proportion of the dataset sampled for each tree during bootstrapping. Default: `None` (all samples are used).

To account for class imbalance in the motor behavior labels, the balanced class-weight option was used, which automatically adjusts the weight of each class inversely to its frequency.

Random Forest Regressor

The same Random Forest algorithm can be used for regression problems, in which case it is referred to as Random Forest Regressor. In this study, it was used to predict the Proficiency score as a continuous variable. The mechanism is the same as before: each tree is trained on a bootstrap sample of the data (Bagging [45]), with a random subset of features considered at each split. The only difference is in how the trees' predictions are combined: instead of majority voting, the final output is the average of all the trees' predictions, giving a more stable estimate than a single regression tree.

The Random Forest Regressor maintains the same advantages as the standard Random Forest: it resists overfitting, can model non-linear relationships, and can handle high-dimensional data [48].
Since the output is continuous, performance was evaluated using the Mean Absolute Error (MAE) and the Root Mean Squared Error (RMSE).
Since the proficiency scores were unevenly distributed across subjects, with the extremes of the scale poorly represented, a weighting scheme was applied during training to prevent the model from focusing only on the most frequent score values. Each training sample was assigned a weight inversely proportional to the squared frequency of its score, so that subjects with rare scores contributed more strongly to the loss. This is the regression counterpart of the balanced class weighting used in classification, helping the model give due importance to under-represented performance levels. The Random Forest Regressor uses the same hyperparameters as the classifier described above, but with different search ranges, as summarized in Table 2.9.

Hyperparameter tuning is typically performed using **GridSearchCV** or **RandomizedSearchCV**. GridSearchCV tests every possible parameter combination, whereas RandomizedSearchCV samples combinations randomly and is faster, especially for large search spaces.
For the Random Forest Classifier, only a subset of hyperparameters (n_estimators, max_features, max_depth, min_samples_leaf) was tuned by performing a randomized search. Specifically, the search was conducted over 10 iterations per fold to balance model performance with computational efficiency. This number was selected to explore the parameter space effectively while minimizing the risk of overfitting. The parameter ranges were restricted to moderate tree depths and leaf sizes; this strategy was chosen to favor model simplicity, ensuring that the Random Forest remains robust and avoids simply memorizing the training data. This approach is crucial not only to improve generalization on the limited sample size of 42 subjects, but also to maintain clinical interpretability, allowing the model to act as a transparent decision-support tool rather than an uninterpretable system. This tuning process was integrated within a Nested Leave-One-Group-Out Cross-Validation framework to ensure a reliable estimate of the model's performance [49]. Table 2.8 summarizes the hyperparameters selected for the tuning process, along with their corresponding search ranges.

Table 2.8: Hyperparameter search space and range of values used in the RandomizedSearchCV process for the Random Forest Classifier.

Hyperparameter	**Value**
n_estimators	30, 40, 50, 60
max_depth	2, 3, 4, 5
min_samples_leaf	3, 5, 7
max_features	sqrt, log2

Hyperparameter tuning for the Random Forest Regressor followed the same approach as for the classifier, using **RandomizedSearchCV**. Five hyperparameters (n_estimators, max_features, max_depth, min_samples_split, min_samples_leaf) were tuned, while bootstrap sampling was kept enabled and max_samples fixed at 0.7 for every tree. The search ran for 20 iterations per fold, balancing performance with computational cost.
As for the classifier, the parameter ranges were kept narrow, favoring trees with limited depth and larger leaf/split sizes to keep the model simple and avoid overfitting on the training data. With only 42 subjects in the sample, this also helps preserve the model's clinical interpretability. Tuning was carried out within the same Nested Leave-One-Group-Out Cross-Validation framework described above. Table 2.9 lists the hyperparameters tuned and their search ranges.

Table 2.9: Hyperparameter search space and range of values used in the RandomizedSearchCV process for the Random Forest Regressor.

Hyperparameter	**Value**
n_estimators	100, 300
max_depth	3, 4, 5
min_samples_split	5, 7
min_samples_leaf	3, 5
max_features	sqrt, 0.5

In this study, Random Forest is used to classify presence or absence of incorrect motor behaviors and, in its regression form, to predict the continuous proficiency score, both from features extracted by the proposed biomechanical modeling pipeline. Its ability to handle correlated features and capture non-linear patterns makes it well

suited to movement data.

2.4.5 Cross Validation

Leave-One-Group-Out

To evaluate the model's ability to generalize, a **Leave-One-Group-Out (LOGO)** cross-validation strategy was adopted [50], where each subject constitutes a group. Given a dataset with N subjects, the procedure consists of N iterations. At each iteration, data from one subject are used as the test set, while the model is trained on the data from the remaining $N - 1$ subjects.

This method ensures that no data from the same subject are present at the same time in both the train and the test sets, in order to prevent subject-specific bias. As a result, LOGO allows a subject-independent evaluation, where each held-out recording is never seen during training, reflecting the model's ability to generalize to unseen executions, which is essential in applications that involve human movement data [51].

In our dataset, each subject had multiple repetitions of the same exercise. During LOGO, all repetitions from a subject were kept together and exclusively assigned to either the training set or the test set in each iteration. This ensured that no repetition from the same subject appeared simultaneously in both sets, thus maintaining an unbiased evaluation with respect to individual subjects.

Hyperparameter optimization was performed within a nested LOGO loop applied to the training fold of each outer iteration. At each inner step, one of the $N - 1$ training subjects is held out as a validation set, and the model is fit on the remaining subjects. The hyperparameter configuration that achieved the best average validation performance across all inner folds was then used to retrain the model on the full training fold before evaluation on the outer test subject.
At test time, the model produces a prediction for each of the held-out subject's repetitions. These repetition-level predictions are then aggregated into a single subject-level prediction, by majority voting for the classifier and by taking the median for the regressor, which is compared against the subject's reference label to compute the performance metrics.
A complete and detailed schematic representation of the nested Leave-One-Group-Out (LOGO) cross-validation workflow is provided in Figure 2.19.

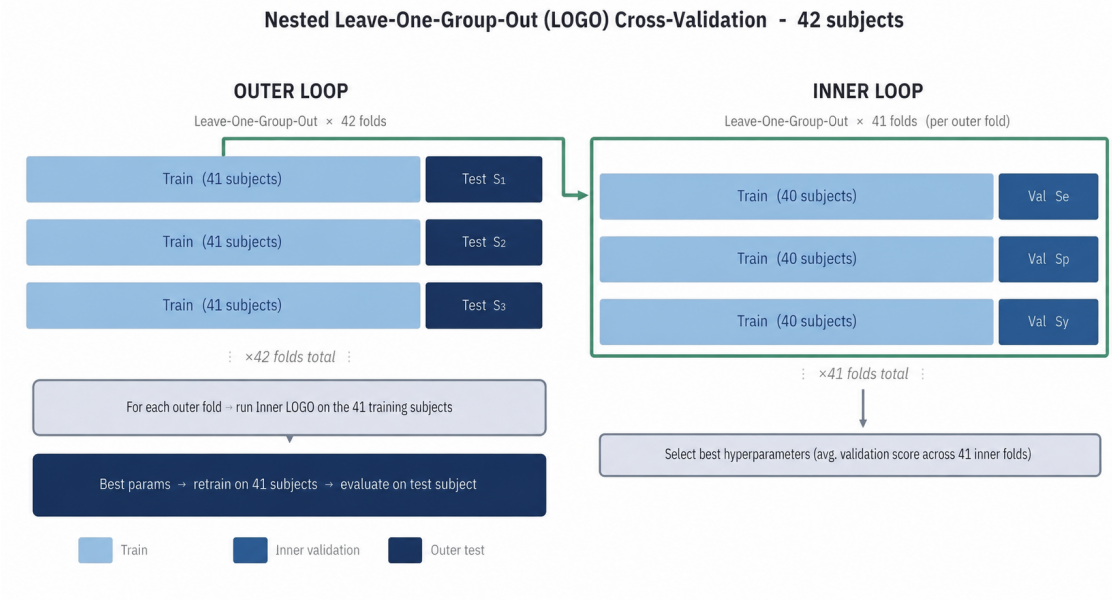


Figure 2.19: Nested Leave-One-Group-Out (LOGO) cross-validation workflow.

2.4.6 Performance Metrics

To assess the performance of the binary models for the detection of incorrect motor behavior, particularly in the presence of class imbalance, the following metrics were adopted: Macro Precision, Macro Recall, and the Macro F1-score. All metrics were computed using macro-averaging, which assigns equal weight to each class regardless of its prevalence.

Precision measures the accuracy of positive predictions, representing the proportion of correctly identified positive instances out of all instances predicted as positive. **Recall** (or sensitivity) quantifies the model's ability to identify all relevant instances by calculating the proportion of correctly identified positive instances out of all actual positive cases:

$$\text{Precision} = \frac{TP}{TP + FP}, \quad \text{Recall} = \frac{TP}{TP + FN}$$

with TP = true positives, FP = false positives, and FN = false negatives.

All three metrics are first computed independently for each class, and then averaged across classes with equal weight (macro-averaging):

$$\text{Metric}_{\text{macro}} = \frac{1}{C}\sum_{c=1}^{C}\text{Metric}_c$$

where C is the number of classes.
The **Macro F1-score** is the unweighted average of the F1-scores calculated for each class individually. It gives equal importance to all the classes, regardless of their prevalence, which is useful when the dataset contains underrepresented classes [52].
The F1-score for a single class is the harmonic mean of Precision and Recall:

$$\text{F1} = 2 \cdot \frac{\text{Precision} \cdot \text{Recall}}{\text{Precision} + \text{Recall}}$$

Using these three metrics together is crucial for a robust evaluation in imbalanced scenarios. While Macro Precision and Macro Recall provide insight into the specific types of errors made by the model across all classes equally, the Macro F1-score summarizes the model's ability to correctly classify each class, independent of class size.
For the Random Forest Regressor, performance was assessed using the **Mean Absolute Error** (MAE) and the **Root Mean Squared Error** (RMSE), since the predicted Proficiency score is continuous rather than categorical. MAE measures the average magnitude of the prediction errors, while RMSE penalizes larger errors more heavily by squaring them before averaging:

$$\text{MAE} = \frac{1}{n}\sum_{i=1}^{n}|y_i - \hat{y}_i|, \quad \text{RMSE} = \sqrt{\frac{1}{n}\sum_{i=1}^{n}(y_i - \hat{y}_i)^2}$$

where y_i is the true Proficiency score, $\hat{y}_i$ the predicted value, and n the number of observations.
In addition, the **Coefficient of Determination** (R^2) was computed to quantify the proportion of variance in the true Proficiency scores explained by the model's predictions:

$$R^2 = 1 - \frac{\sum_{i=1}^{n}(y_i - \hat{y}_i)^2}{\sum_{i=1}^{n}(y_i - \bar{y})^2}$$

where $\bar{y}$ is the mean of the true Proficiency scores. Values closer to 1 indicate a better fit.
Used together, these metrics give a fuller picture of the regression error, with RMSE more sensitive to large deviations.

Chapter 3

Results

This Chapter presents the results of the classification and regression models obtained in this analysis. For each exercise, the results include:

- The features selected for each motor behavior and for the final proficiency assessment.
- The UMAP projection of the selected features for each motor behavior and for the final proficiency assessment.
- The Macro F1-scores, Precision, and Recall for each motor behavior.
- The MAE, RMSE, and R^2 for the final proficiency assessment.
- The Confusion Matrices of the classifiers for each motor behavior.
- The scatter plot of true versus predicted Proficiency scores, with a fitted regression line, for the final proficiency assessment.
- The bar plot of the per-subject prediction error for the final proficiency assessment, highlighting subjects for which the model over- or under-estimates the score.

3.1 Ex 1

This Section presents the results obtained in the analysis of Exercise 1, Raising the Power.

3.1.1 Motor Behavior 1

This Subsection presents the results obtained in the analysis of the Motor Behavior 1, Wrist Height, for Exercise 1, Raising the Power.
The features selected for this motor behavior are listed in Table 3.1, while the underlying data structure is visualized through UMAP projections in Figure 3.1. Furthermore, Table 3.2 presents the final evaluation metrics, specifically the Macro F1-score, Precision, and Recall calculated at the subject level. The classification performance is further detailed in the confusion matrices shown in Figure 3.2.

Table 3.1: Selected features for Motor Behavior 1, Exercise 1: Raising the Power.

Selected Features	**Values considered**
Wrist vertical displacement (Y-axis), L & R	range
Wrist 3D speed, L & R	median, std
Wrist 3D acceleration magnitude, L & R	median, std
Wrist 3D jerk magnitude, L & R	median, max
Inter-wrist / inter-shoulder 3D distance ratio	median, std, handpk, Δ(start $\rightarrow$ handpk)
Wrist to Shoulder vertical distance, L & R	handpk
Wrist to Shoulder 3D Euclidean distance, L & R	handpk
Δ Wrist Y from rep start to max knee bend, L & R	$\Delta(\max_{\text{knee}} - \text{start})$

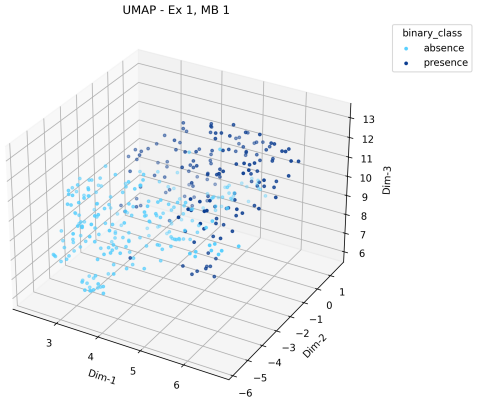
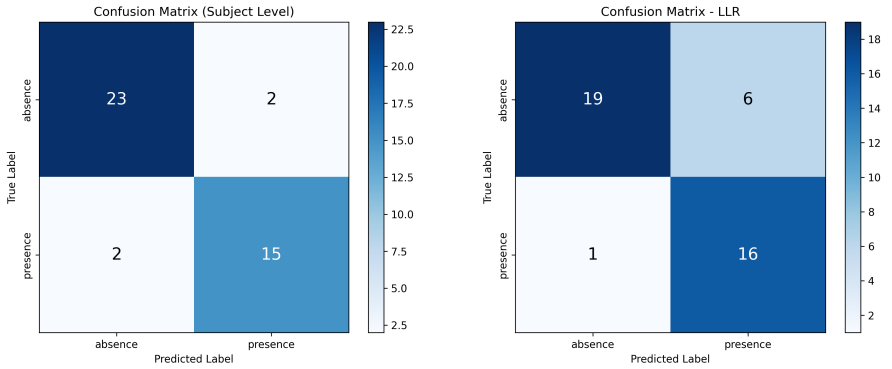


Figure 3.1: UMAP projection, Motor Behavior 1, Exercise 1: Raising the Power.



(a) Random Forest (b) Logistic Regression

Figure 3.2: Confusion Matrices for Motor Behavior 1: Random Forest (a) and Logistic Regression (b), Exercise 1: Raising the Power.

Table 3.2: Classification results reported as macro-averaged metrics, for Motor Behavior 1, Exercise 1.

Metric	**Random Forest**	**Logistic Regression**
Macro F1-score	0.90	0.83
Macro Precision	0.90	0.84
Macro Recall	0.90	0.85

3.1.2 Motor Behavior 2

This Subsection presents the results obtained in the analysis of the Motor Behavior 2, Synchronization Knees-Wrists, for Exercise 1, Raising the Power.
The features selected for this motor behavior are listed in Table 3.3, while the underlying data structure is visualized through UMAP projections in Figure 3.3. Furthermore, Table 3.4 presents the final evaluation metrics, specifically the Macro F1-score, Precision, and Recall calculated at the subject level. The classification performance is further detailed in the confusion matrices shown in Figure 3.4.

Table 3.3: Selected features for Motor Behavior 2, Exercise 1: Raising the Power.

Selected Features	**Values considered**
Wrist vertical displacement (Y-axis), L & R	range
Knee vertical displacement (Y-axis), L & R	range
Asymmetry between time-to-max-bend and time-from-max-to-end	single value
Asymmetry between wrist raise and lower duration	single value
Time from rep start to max knee bend	single value
Time from max knee bend to rep end	single value
Absolute time difference between wrist and knee peaks	single value
Cross-correlation Wrist–Knee trajectory, L & R	min
Zero-lag Pearson correlation between knee and wrist	single value
Optimal time lag for max wrist–knee cross-correlation	single value
Normalized time difference between wrist and knee peaks	single value
Pearson correlation of wrist and knee velocity profiles	single value
MSE between normalized wrist and knee trajectories	single value
Wrist vertical position at max knee bend, L & R	single value
Δ Wrist Y from rep start to max knee bend, L & R	$\Delta(\max_{\text{knee}} - \text{start})$
Knee flexion angle, L & R	start, end, $\Delta(\text{start} \rightarrow \text{end})$, $\Delta(\text{start} \rightarrow \text{max})$
Wrist 3D speed, L & R	median, std
Wrist 3D acceleration magnitude, L & R	median, std
Wrist 3D jerk magnitude, L & R	median, max

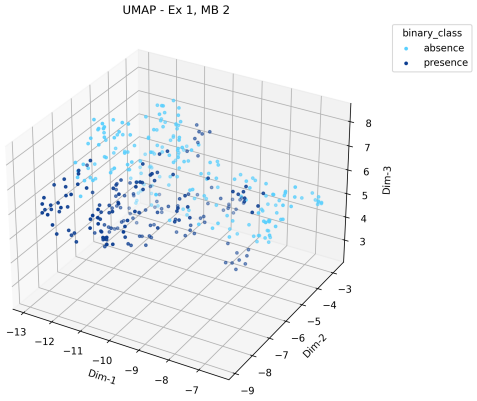
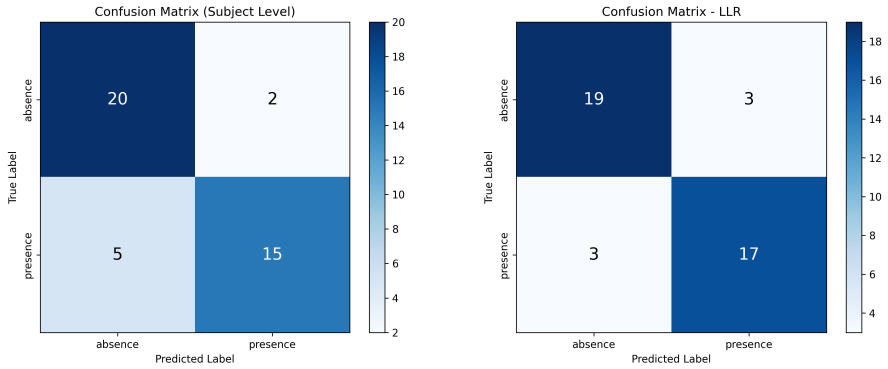


Figure 3.3: UMAP projection, Motor Behavior 2, Exercise 1: Raising the Power.



(a) Random Forest (b) Logistic Regression

Figure 3.4: Confusion Matrices for Motor Behavior 2: Random Forest (a) and Logistic Regression (b), Exercise 1: Raising the Power.

Table 3.4: Classification results reported as macro-averaged metrics, for Motor Behavior 2, Exercise 1.

Metric	**Random Forest**	**Logistic Regression**
Macro F1-score	0.83	0.86
Macro Precision	0.84	0.86
Macro Recall	0.83	0.86

3.1.3 Motor Behavior 3

This Subsection presents the results obtained in the analysis of the Motor Behavior 3, Feet Width, for Exercise 1, Raising the Power.
The features selected for this motor behavior are listed in Table 3.5, while the underlying data structure is visualized through UMAP projections in Figure 3.5. Furthermore, Table 3.6 presents the final evaluation metrics, specifically the Macro F1-score, Precision, and Recall calculated at the subject level. The classification performance is further detailed in the confusion matrices shown in Figure 3.6.

Table 3.5: Selected features for Motor Behavior 3, Exercise 1: Raising the Power.

Selected Features	**Values considered**
Ankle-to-ankle / hip-to-hip 3D distance ratio	median, Δ(start $\rightarrow$ end)
Ankle-to-ankle / shoulder-to-shoulder 3D distance ratio	median, Δ(start $\rightarrow$ end)
3D Euclidean distance between ankles	median
Knee-to-knee Euclidean distance	median, std, Δ(start $-$ $\max_{\text{knee}}$)

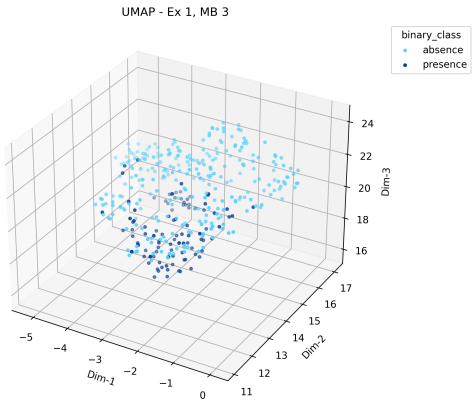
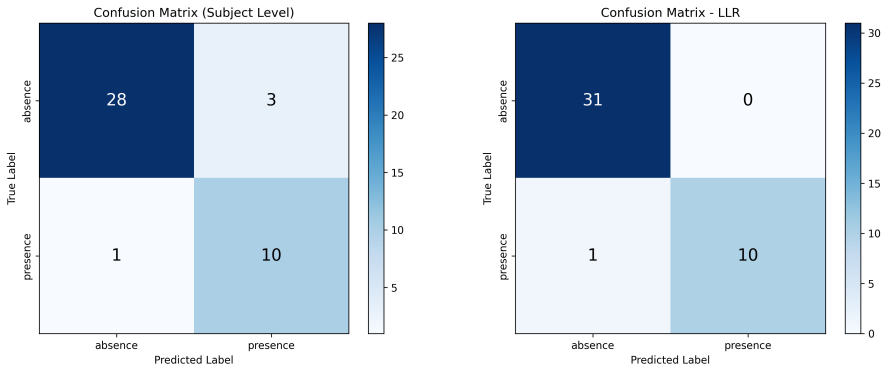


Figure 3.5: UMAP projection, Motor Behavior 3, Exercise 1: Raising the Power.



(a) Random Forest (b) Logistic Regression

Figure 3.6: Confusion Matrices for Motor Behavior 3: Random Forest (a) and Logistic Regression (b), Exercise 1: Raising the Power.

Table 3.6: Classification results reported as macro-averaged metrics, for Motor Behavior 3, Exercise 1.

Metric	**Random Forest**	**Logistic Regression**
Macro F1-score	0.88	0.97
Macro Precision	0.87	0.98
Macro Recall	0.91	0.95

3.1.4 Motor Behavior 4

This Subsection presents the results obtained in the analysis of the Motor Behavior 4, Wrists-Shoulders Parallelism, for Exercise 1, Raising the Power.
The features selected for this motor behavior are listed in Table 3.7, while the underlying data structure is visualized through UMAP projections in Figure 3.7. Furthermore, Table 3.8 presents the final evaluation metrics, specifically the Macro F1-score, Precision, and Recall calculated at the subject level. The classification performance is further detailed in the confusion matrices shown in Figure 3.8.

Table 3.7: Selected features for Motor Behavior 4, Exercise 1: Raising the Power.

Selected Features	**Values considered**
Elbow lateral displacement (X-axis), L & R	range
Wrist lateral displacement (X-axis), L & R	range
Inter-wrist / inter-shoulder 3D distance ratio	median, std, handpk, Δ(start $\rightarrow$ handpk)

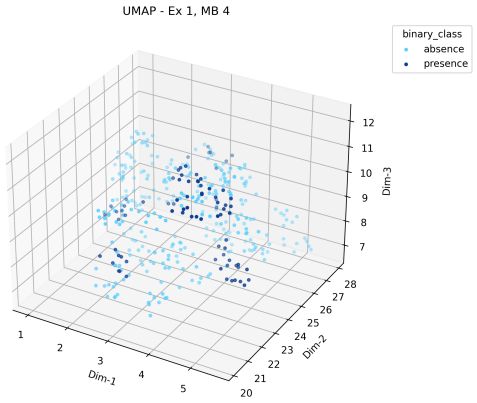
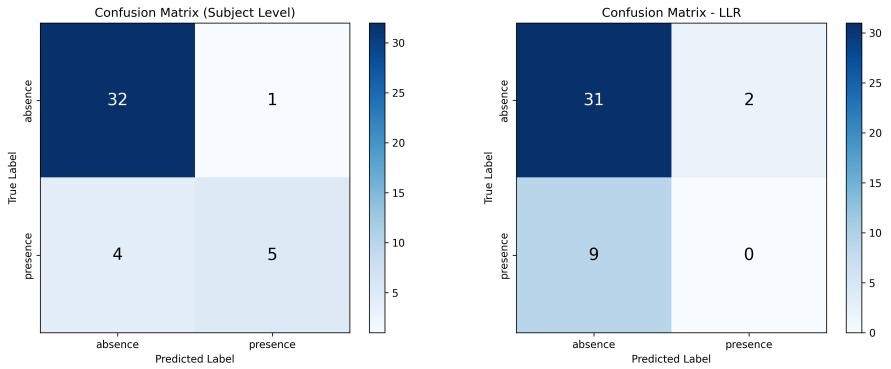


Figure 3.7: UMAP projection, Motor Behavior 4, Exercise 1: Raising the Power.



(a) Random Forest (b) Logistic Regression

Figure 3.8: Confusion Matrices for Motor Behavior 4: Random Forest (a) and Logistic Regression (b), Exercise 1: Raising the Power.

Table 3.8: Classification results reported as macro-averaged metrics, for Motor Behavior 4, Exercise 1.

Metric	Random Forest	Logistic Regression
Macro F1-score	0.80	0.42
Macro Precision	0.86	0.39
Macro Recall	0.76	0.47

3.1.5 Motor Behavior 5

This Subsection presents the results obtained in the analysis of the Motor Behavior 5, Feet Parallelism, for Exercise 1, Raising the Power.
The features selected for this motor behavior are listed in Table 3.9, while the underlying data structure is visualized through UMAP projections in Figure 3.9. Furthermore, Table 3.10 presents the final evaluation metrics, specifically the Macro F1-score, Precision, and Recall calculated at the subject level. The classification performance is further detailed in the confusion matrices shown in Figure 3.10.

Table 3.9: Selected features for Motor Behavior 5, Exercise 1: Raising the Power.

Selected Features	Values considered
Angle between L and R foot axes (toe–heel vectors)	median, max, min
Angle between L and R foot axes (small toe–ankle vectors)	median, max, min
Big-toe-to-big-toe / heel-to-heel distance ratio	median, max, min
Big-toe-to-big-toe / ankle-to-ankle distance ratio	median, max, min
Small-toe-to-small-toe / ankle-to-ankle distance ratio	median, max, min

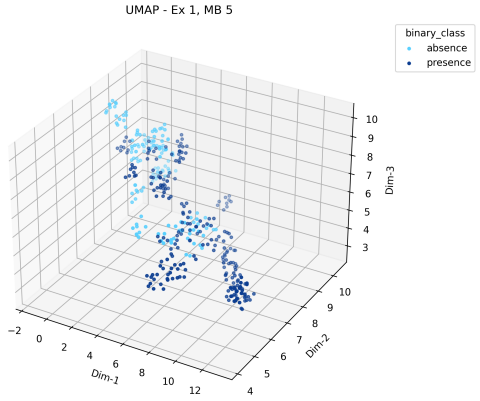
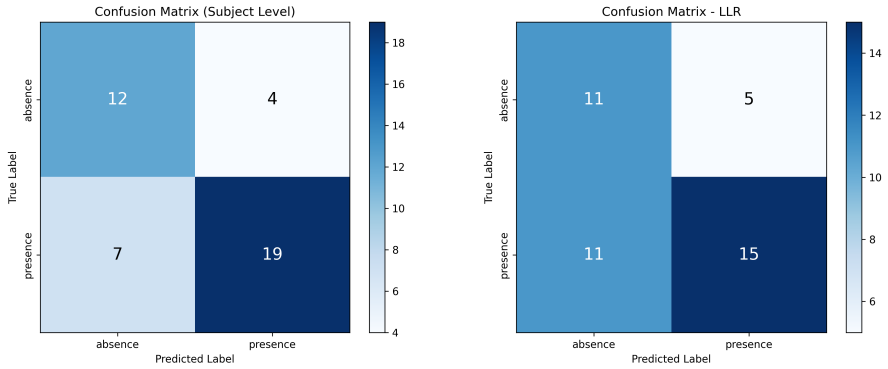


Figure 3.9: UMAP projection, Motor Behavior 5, Exercise 1: Raising the Power.



(a) Random Forest (b) Logistic Regression

Figure 3.10: Confusion Matrices for Motor Behavior 5: Random Forest (a) and Logistic Regression (b), Exercise 1: Raising the Power.

Table 3.10: Classification results reported as macro-averaged metrics, for Motor Behavior 5, Exercise 1.

Metric	**Random Forest**	**Logistic Regression**
Macro F1-score	0.73	0.62
Macro Precision	0.73	0.63
Macro Recall	0.74	0.63

3.1.6 Motor Behavior 6

This Subsection presents the results obtained in the analysis of the Motor Behavior 6, Arm Three-Dimensionality, for Exercise 1, Raising the Power.
The features selected for this motor behavior are listed in Table 3.11, while the underlying data structure is visualized through UMAP projections in Figure 3.11. Furthermore, Table 3.12 presents the final evaluation metrics, specifically the Macro F1-score, Precision, and Recall calculated at the subject level. The classification performance is further detailed in the confusion matrices shown in Figure 3.12.

Table 3.11: Selected features for Motor Behavior 6, Exercise 1: Raising the Power.

Selected Features	**Values considered**
Elbow lateral displacement (X-axis), L & R	range
Elbow vertical displacement (Y-axis), L & R	range
Elbow sagittal displacement (Z-axis), L & R	range
Wrist lateral displacement (X-axis), L & R	range
Wrist vertical displacement (Y-axis), L & R	range
Wrist sagittal displacement (Z-axis), L & R	range
Wrist to Shoulder 3D Euclidean distance, L & R	median, min, max, std, range
Elbow flexion angle, L & R	median, min, range, std, min_afterhandpk, Δ(start $\rightarrow$ max)
Min Z-depth difference Wrist vs. Shoulder, L & R	min
Total 3D path length of wrist, L & R	cumulative sum

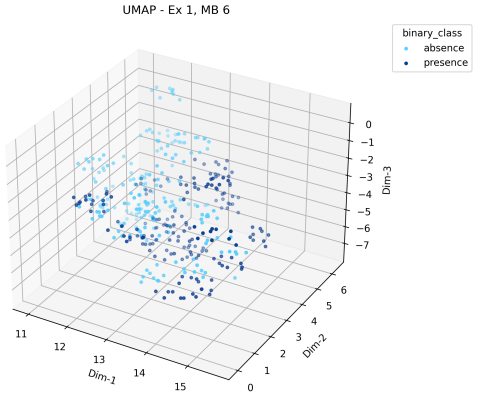
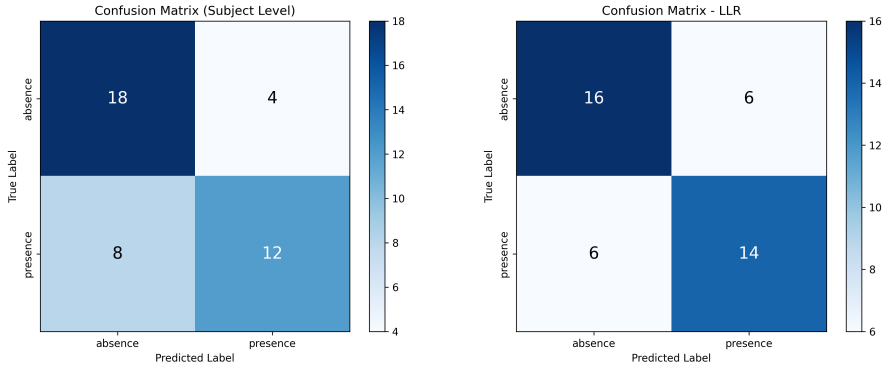


Figure 3.11: UMAP projection, Motor Behavior 6, Exercise 1: Raising the Power.



(a) Random Forest (b) Logistic Regression

Figure 3.12: Confusion Matrices for Motor Behavior 6: Random Forest (a) and Logistic Regression (b), Exercise 1: Raising the Power.

Table 3.12: Classification results reported as macro-averaged metrics, for Motor Behavior 6, Exercise 1.

Metric	Random Forest	Logistic Regression
Macro F1-score	0.71	0.71
Macro Precision	0.72	0.71
Macro Recall	0.71	0.71

Figures 3.13 and 3.14 provide a comprehensive summary of the evaluated motor behaviors, reporting the Macro F1-score achieved for each.

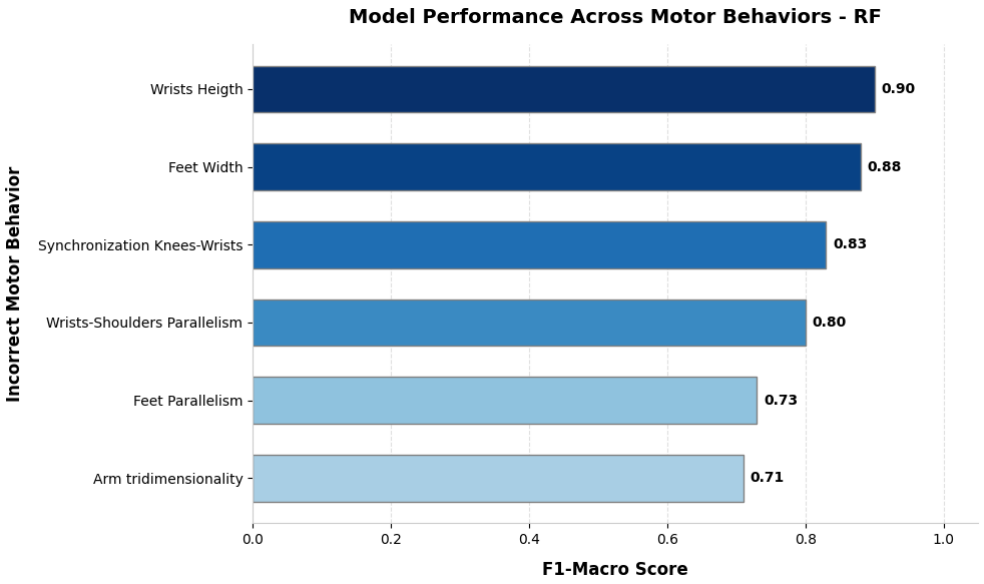


Figure 3.13: Summary of Motor Behavior analysis results for Exercise 1 with Random Forest.

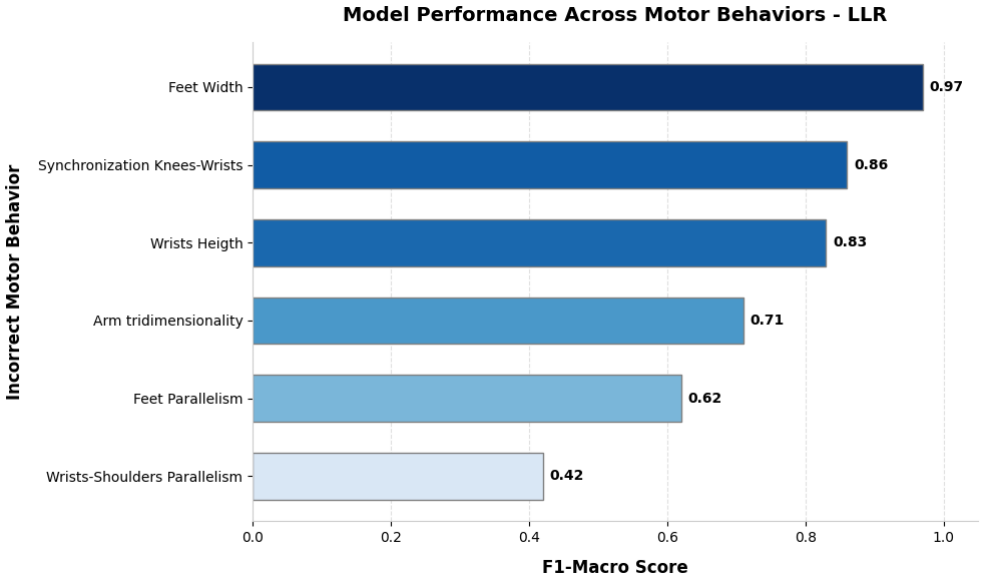


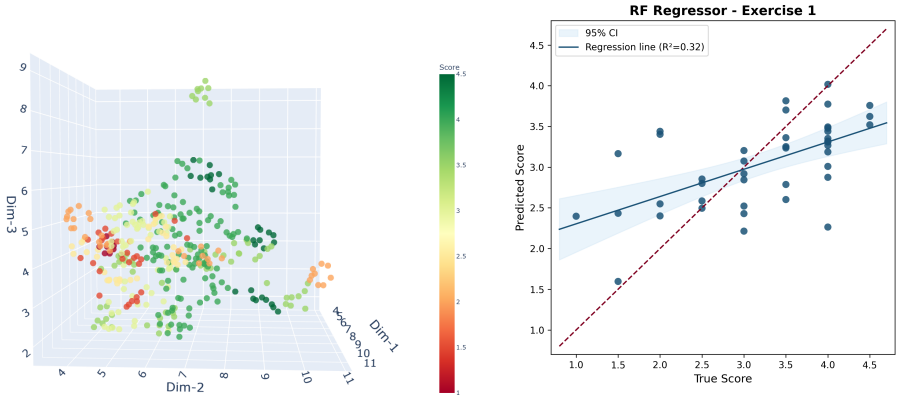
Figure 3.14: Summary of Motor Behavior analysis results for Exercise 1 with Logistic Regression.

3.1.7 Proficiency Assessment

This Subsection presents the results obtained for the proficiency analysis for Exercise 1, Raising the Power. Two feature selection approaches were evaluated, as described in Subsection 2.4.2: an automated nested mRMR selection and a manual selection based on domain knowledge.
The features selected for the manual approach are listed in Table 3.13. The underlying data structure is visualized through UMAP projections in Figures 3.15(a) and 3.16(a) for the nested mRMR and manual approach respectively. For the nested mRMR approach, the UMAP projection was computed using the feature set selected at fold 37, chosen as the fold whose MAE was closest to the mean across folds, and therefore representative of the average regression performance.
The regression performance is further illustrated in Figures 3.15(b) and 3.16(b), which show the true versus predicted Proficiency scores together with a fitted regression line. The prediction error per subject for the manual approach is shown in Figure 3.17, highlighting the subjects for which the model over- or under-estimates the proficiency score. Furthermore, Table 3.14 presents the final evaluation metrics, specifically the MAE, RMSE, and R^2 calculated at the subject level, for both approaches.

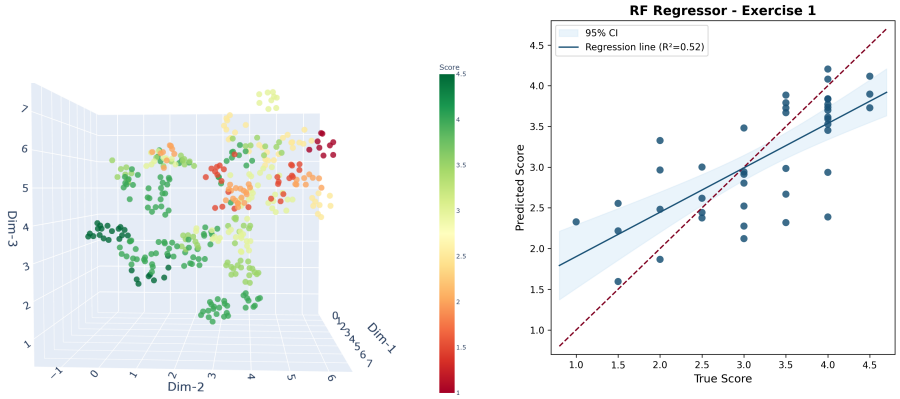
Table 3.13: Selected features for Proficiency evaluation, Exercise 1: Raising the Power.

Selected Features	**Description**
Zero-lag Pearson correlation between knee and wrist	single value
Normalized time difference between wrist and knee peaks	single value
MSE between normalized wrist and knee trajectories	single value
Right wrist vertical displacement (Y-axis)	range
Left wrist vertical displacement (Y-axis)	range
L-wrist to L-shoulder vertical distance	handpk
Presence or Absence of all 6 MB	single binary value



(a) UMAP (b) True vs Predicted

Figure 3.15: Nested mRMR approach: UMAP projection (a) and regression performance (b), Exercise 1: Raising the Power.



(a) UMAP (b) True vs Predicted

Figure 3.16: Manual feature selection approach: UMAP projection (a) and regression performance (b), Exercise 1: Raising the Power.

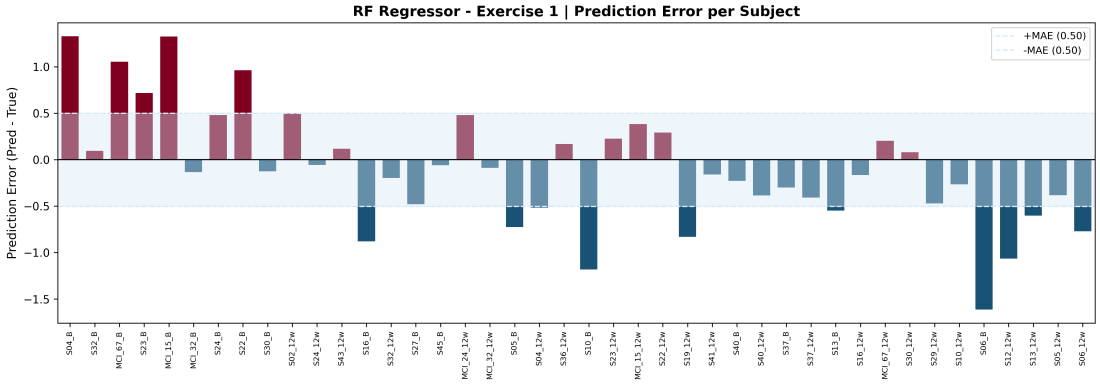


Figure 3.17: Prediction error per subject, manual feature selection, Exercise 1: Raising the Power.

Table 3.14: Regression results for Proficiency evaluation, Exercise 1.

Metric	Nested mRMR	Manual FS
MAE	0.610	0.501
RMSE	0.759	0.638
R^2	0.320	0.520

3.2 Ex 6

This Section presents the results obtained in the analysis of Exercise 6, Golden Rooster Stands on One Leg.

3.2.1 Motor Behavior 1

This Subsection presents the results obtained in the analysis of the Motor Behavior 1, Wrist Height, for Exercise 6, Golden Rooster Stands on One Leg.
The features selected for this motor behavior are listed in Table 3.15, while the underlying data structure is visualized through UMAP projections in Figure 3.18. Furthermore, Table 3.16 presents the final evaluation metrics, specifically the Macro F1-score, Precision, and Recall calculated at the subject level. The classification performance is further detailed in the confusion matrices shown in Figure 3.19.

Table 3.15: Selected features for Motor Behavior 1, Exercise 6: Golden Rooster Stands on One Leg.

Selected Features	**Values considered**
Wrist vertical displacement (Y-axis), L & R	range
Elbow vertical displacement (Y-axis), L & R	range
Active wrist Y position	kneepk, Δ(start $\rightarrow$ kneepk), Δ(kneepk $\rightarrow$ end)
Active elbow Y position	kneepk, Δ(start $\rightarrow$ kneepk), Δ(kneepk $\rightarrow$ end)
Active wrist vertical and euclidean displacement	range
Active elbow vertical and euclidean displacement	range
Active wrist to shoulder vertical distance	kneepk
Active wrist to shoulder 3D euclidean distance	kneepk

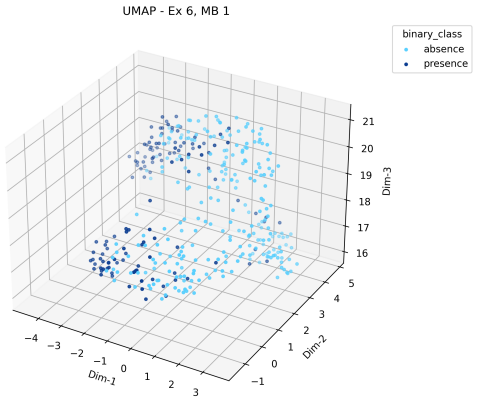
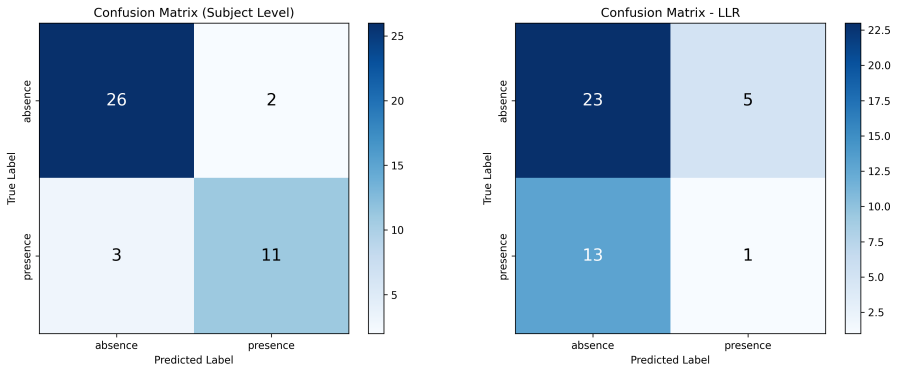


Figure 3.18: UMAP projection, Motor Behavior 1, Exercise 6: Golden Rooster Stands on One Leg.



(a) Random Forest (b) Logistic Regression

Figure 3.19: Confusion Matrices for Motor Behavior 1: Random Forest (a) and Logistic Regression (b), Exercise 6: Golden Rooster Stands on One Leg.

Table 3.16: Classification results reported as macro-averaged metrics, for Motor Behavior 1, Exercise 6.

Metric	Random Forest	Logistic Regression
Macro F1-score	0.86	0.41
Macro Precision	0.87	0.40
Macro Recall	0.86	0.45

3.2.2 Motor Behavior 2

This Subsection presents the results obtained in the analysis of the Motor Behavior 2, Limb Side Coordination, for Exercise 6, Golden Rooster Stands on One Leg. The features selected for this motor behavior are listed in Table 3.17, while the underlying data structure is visualized through UMAP projections in Figure 3.20. Furthermore, Table 3.18 presents the final evaluation metrics, specifically the Macro F1-score, Precision, and Recall calculated at the subject level. The classification performance is further detailed in the confusion matrices shown in Figure 3.21.

Table 3.17: Selected features for Motor Behavior 2, Exercise 6: Golden Rooster Stands on One Leg.

Selected Features	Values considered
Duration knee remains elevated	single value
Duration wrist remains elevated	single value
Duration knee and wrist remain elevated simultaneously	single value
Overlap time / knee hold time	single value
Zero-lag Pearson correlation between knee and wrist trajectories	single value
Optimal time lag for max wrist–knee cross-correlation	single value
Normalized time difference between wrist and knee peaks	single value
Pearson correlation of wrist and knee velocity profiles	single value
MSE between normalized wrist and knee trajectories	single value

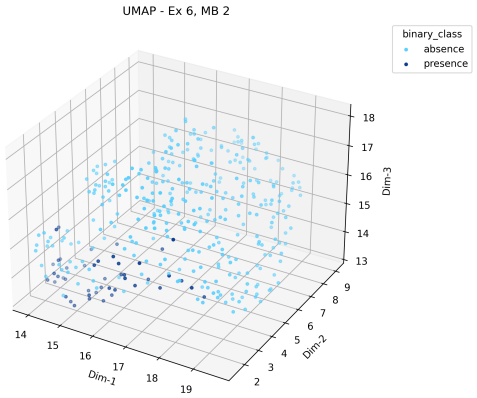
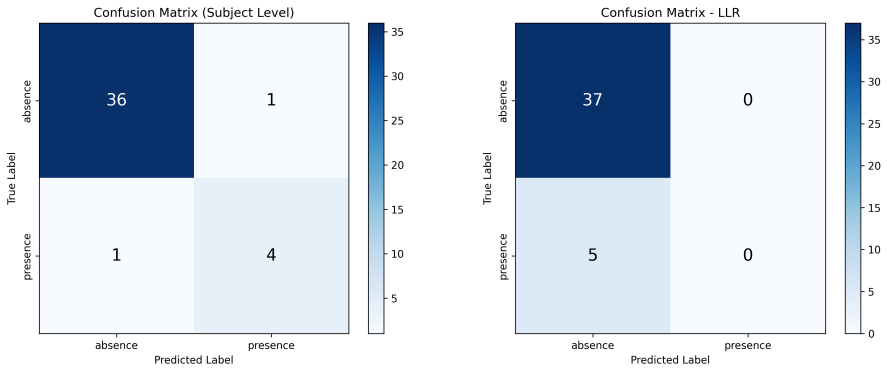


Figure 3.20: UMAP projection, Motor Behavior 2, Exercise 6: Golden Rooster Stands on One Leg.



(a) Random Forest (b) Logistic Regression

Figure 3.21: Confusion Matrices for Motor Behavior 2: Random Forest (a) and Logistic Regression (b), Exercise 6: Golden Rooster Stands on One Leg.

Table 3.18: Classification results reported as macro-averaged metrics, for Motor Behavior 2, Exercise 6.

Metric	**Random Forest**	**Logistic Regression**
Macro F1-score	0.89	0.47
Macro Precision	0.89	0.44
Macro Recall	0.89	0.50

3.2.3 Motor Behavior 3

This Subsection presents the results obtained in the analysis of the Motor Behavior 3, Initial Foot Stance Width, for Exercise 6, Golden Rooster Stands on One Leg. The features selected for this motor behavior are listed in Table 3.19, while the underlying data structure is visualized through UMAP projections in Figure 3.22. Furthermore, Table 3.20 presents the final evaluation metrics, specifically the Macro F1-score, Precision, and Recall calculated at the subject level. The classification performance is further detailed in the confusion matrices shown in Figure 3.23.

Table 3.19: Selected features for Motor Behavior 3, Exercise 6: Golden Rooster Stands on One Leg.

Selected Features	**Values considered**
3D Euclidean distance between ankles (start phase)	frame_0, start_median
Ankle-to-ankle / hip-to-hip 3D distance ratio (start phase)	frame_0, start_median
Ankle-to-ankle / shoulder-to-shoulder 3D distance ratio (start phase)	frame_0, start_median
Knee-to-knee Euclidean distance (start phase)	frame_0, start_median, Δ_{rel}(start $\rightarrow$ maxbend)

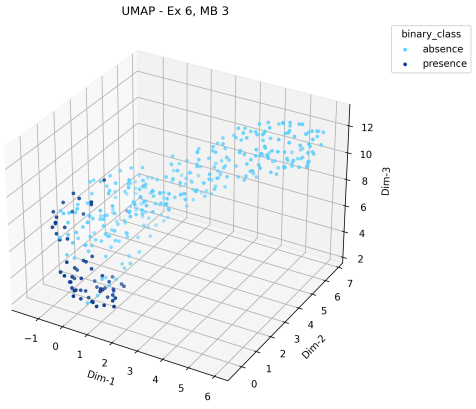
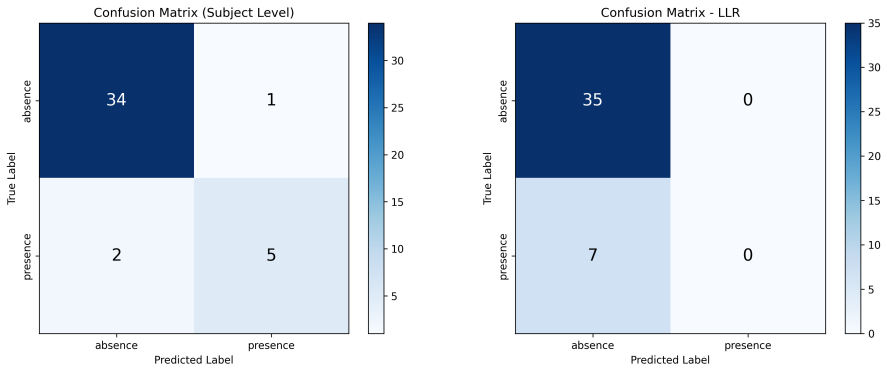


Figure 3.22: UMAP projection, Motor Behavior 3, Exercise 6: Golden Rooster Stands on One Leg.



(a) Random Forest (b) Logistic Regression

Figure 3.23: Confusion Matrices for Motor Behavior 3: Random Forest (a) and Logistic Regression (b), Exercise 6: Golden Rooster Stands on One Leg.

Table 3.20: Classification results reported as macro-averaged metrics, for Motor Behavior 3, Exercise 6.

Metric	Random Forest	Logistic Regression
Macro F1-score	0.86	0.45
Macro Precision	0.89	0.42
Macro Recall	0.84	0.50

3.2.4 Motor Behavior 4

This Subsection presents the results obtained in the analysis of the Motor Behavior 4, Wrist-Elbow Alignment of Raised Arm, for Exercise 6, Golden Rooster Stands on One Leg.
The features selected for this motor behavior are listed in Table 3.21, while the underlying data structure is visualized through UMAP projections in Figure 3.24. Furthermore, Table 3.22 presents the final evaluation metrics, specifically the Macro F1-score, Precision, and Recall calculated at the subject level. The classification performance is further detailed in the confusion matrices shown in Figure 3.25.

Table 3.21: Selected features for Motor Behavior 4, Exercise 6: Golden Rooster Stands on One Leg.

Selected Features	**Values considered**
Active elbow flexion angle	kneepk, Δ(start $\rightarrow$ kneepk), Δ(kneepk $\rightarrow$ end)
Active wrist vertical position	kneepk, Δ(start $\rightarrow$ kneepk), Δ(kneepk $\rightarrow$ end)
Active wrist displacement	x_range, y_range, euc_range
Wrist displacement, L & R	x_range, y_range
Active wrist to shoulder vertical distance	kneepk
Active wrist to shoulder lateral distance	kneepk
Active wrist to shoulder 3D Euclidean distance	kneepk
Active elbow–wrist distance	x_kneepk, z_kneepk
Angle between forearm vector and vertical axis	kneepk

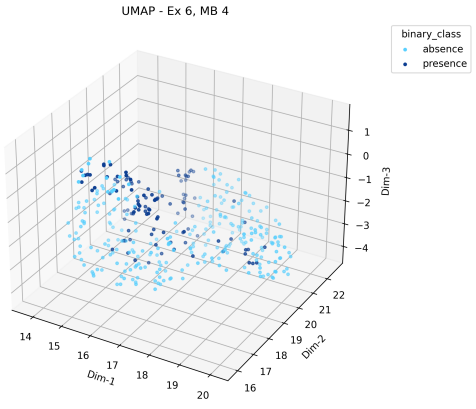
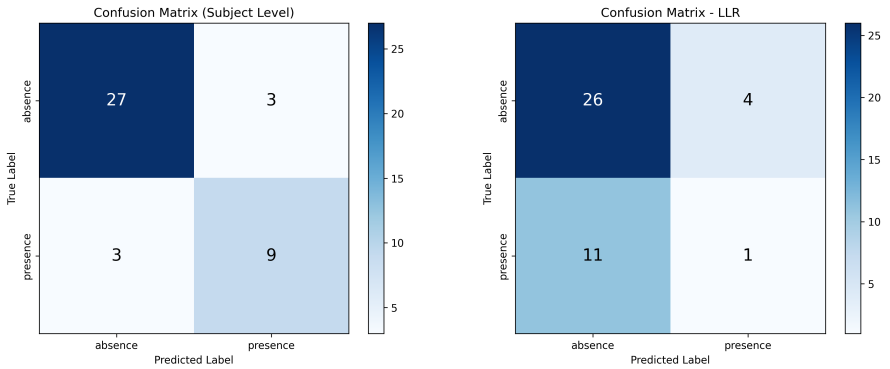


Figure 3.24: UMAP projection, Motor Behavior 4, Exercise 6: Golden Rooster Stands on One Leg.



(a) Random Forest (b) Logistic Regression

Figure 3.25: Confusion Matrices for Motor Behavior 4: Random Forest (a) and Logistic Regression (b), Exercise 6: Golden Rooster Stands on One Leg.

Table 3.22: Classification results reported as macro-averaged metrics, for Motor Behavior 4, Exercise 6.

Metric	Random Forest	Logistic Regression
Macro F1-score	0.83	0.45
Macro Precision	0.83	0.45
Macro Recall	0.83	0.48

3.2.5 Motor Behavior 5

This Subsection presents the results obtained in the analysis of the Motor Behavior 5, Ankle-Knee Alignment of Raised Leg, for Exercise 6, Golden Rooster Stands on One Leg.
The features selected for this motor behavior are listed in Table 3.23, while the underlying data structure is visualized through UMAP projections in Figure 3.26. Furthermore, Table 3.24 presents the final evaluation metrics, specifically the Macro F1-score, Precision, and Recall calculated at the subject level. The classification performance is further detailed in the confusion matrices shown in Figure 3.27.

Table 3.23: Selected features for Motor Behavior 5, Exercise 6: Golden Rooster Stands on One Leg.

Selected Features	**Values considered**
Active knee displacement	y_range, z_range, euc_range
Active ankle displacement	y_range, z_range, euc_range
Knee displacement, L & R	y_range, z_range
Ankle sagittal displacement (Z-axis), L & R	range
Heel sagittal displacement (Z-axis), L & R	range
Active ankle–knee distance at knee peak	x_kneepk, z_kneepk
Angle between lower leg vector (knee→ankle) and vertical axis	kneepk

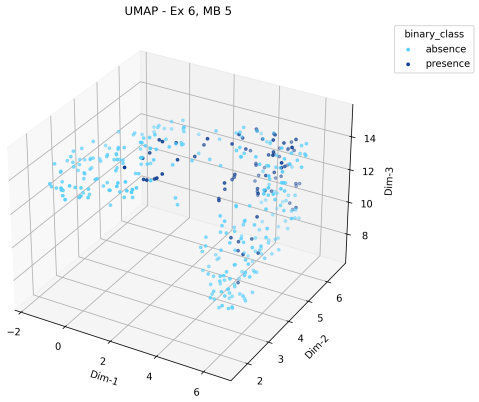
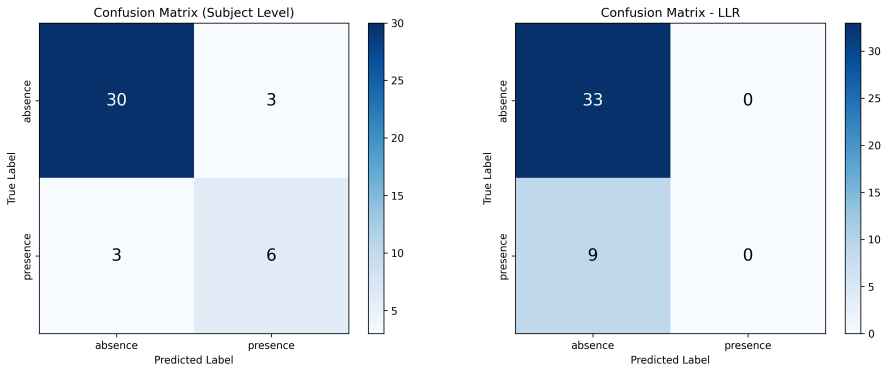


Figure 3.26: UMAP projection, Motor Behavior 5, Exercise 6: Golden Rooster Stands on One Leg.



(a) Random Forest (b) Logistic Regression

Figure 3.27: Confusion Matrices for Motor Behavior 5: Random Forest (a) and Logistic Regression (b), Exercise 6: Golden Rooster Stands on One Leg.

Table 3.24: Classification results reported as macro-averaged metrics, for Motor Behavior 5, Exercise 6.

Metric	Random Forest	Logistic Regression
Macro F1-score	0.79	0.44
Macro Precision	0.79	0.39
Macro Recall	0.79	0.50

3.2.6 Motor Behavior 6

This Subsection presents the results obtained in the analysis of the Motor Behavior 6, Hip-Knee Height Alignment of Raised Leg, for Exercise 6, Golden Rooster Stands on One Leg.
The features selected for this motor behavior are listed in Table 3.25, while the underlying data structure is visualized through UMAP projections in Figure 3.28. Furthermore, Table 3.26 presents the final evaluation metrics, specifically the Macro F1-score, Precision, and Recall calculated at the subject level. The classification performance is further detailed in the confusion matrices shown in Figure 3.29.

Table 3.25: Selected features for Motor Behavior 6, Exercise 6: Golden Rooster Stands on One Leg.

Selected Features	**Values considered**
Active knee vertical position	kneepk, Δ(start $\rightarrow$ kneepk), Δ(kneepk $\rightarrow$ end)
Active ankle vertical position	kneepk, Δ(start $\rightarrow$ kneepk), Δ(kneepk $\rightarrow$ end)
Active knee flexion angle	kneepk, max, Δ(start $\rightarrow$ kneepk), Δ(kneepk $\rightarrow$ end)
Active knee displacement	y_range, euc_range
Active ankle displacement	y_range, euc_range
Knee vertical displacement (Y-axis), L & R	range
Ankle vertical displacement (Y-axis), L & R	range
Active hip–knee vertical distance	kneepk

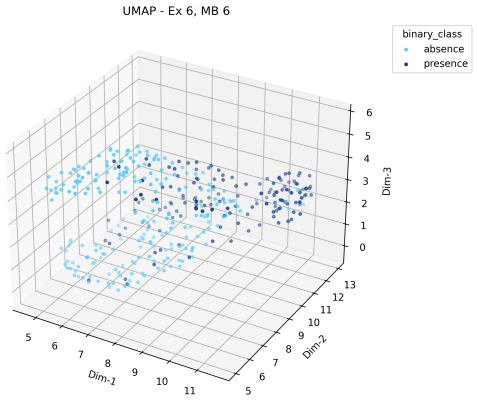
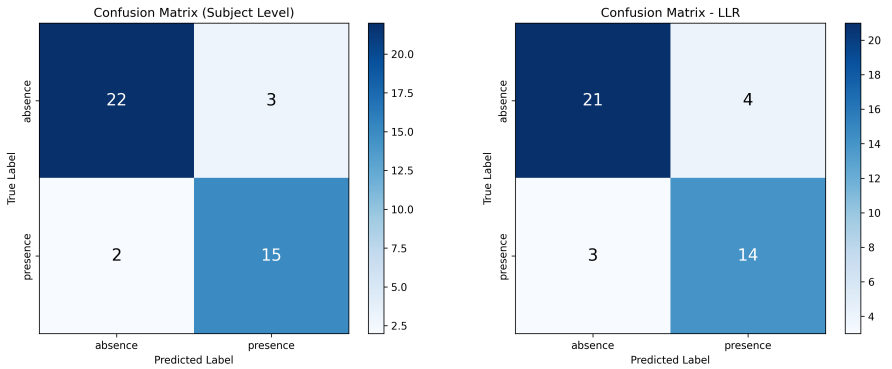


Figure 3.28: UMAP projection, Motor Behavior 6, Exercise 6: Golden Rooster Stands on One Leg.



(a) Random Forest (b) Logistic Regression

Figure 3.29: Confusion Matrices for Motor Behavior 6: Random Forest (a) and Logistic Regression (b), Exercise 6: Golden Rooster Stands on One Leg.

Table 3.26: Classification results reported as macro-averaged metrics, for Motor Behavior 6, Exercise 6.

Metric	Random Forest	Logistic Regression
Macro F1-score	0.88	0.83
Macro Precision	0.88	0.83
Macro Recall	0.88	0.83

3.2.7 Motor Behavior 7

This Subsection presents the results obtained in the analysis of the Motor Behavior 7, Elbow-Knee Alignment of Raised Side, for Exercise 6, Golden Rooster Stands on One Leg.
The features selected for this motor behavior are listed in Table 3.27, while the underlying data structure is visualized through UMAP projections in Figure 3.30. Furthermore, Table 3.28 presents the final evaluation metrics, specifically the Macro F1-score, Precision, and Recall calculated at the subject level. The classification performance is further detailed in the confusion matrices shown in Figure 3.31.

Table 3.27: Selected features for Motor Behavior 7, Exercise 6: Golden Rooster Stands on One Leg.

Selected Features	**Values considered**
Active elbow flexion angle	kneepk, Δ(start $\rightarrow$ kneepk), Δ(kneepk $\rightarrow$ end)
Active elbow vertical position	kneepk
Active knee vertical position	kneepk
Active elbow–knee distance	euc_kneepk, x_kneepk, y_kneepk
Active elbow displacement	x_range, y_range, euc_range
Active knee displacement	x_range, y_range, euc_range
Elbow displacement, L & R	x_range, y_range
Knee displacement, L & R	x_range, y_range

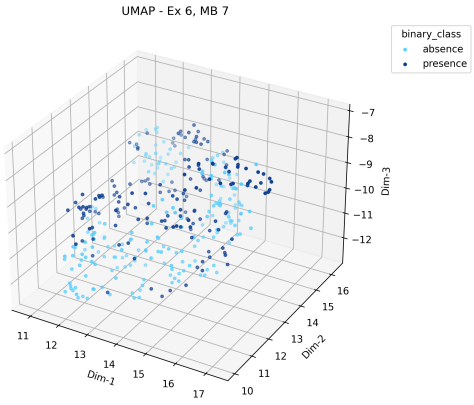
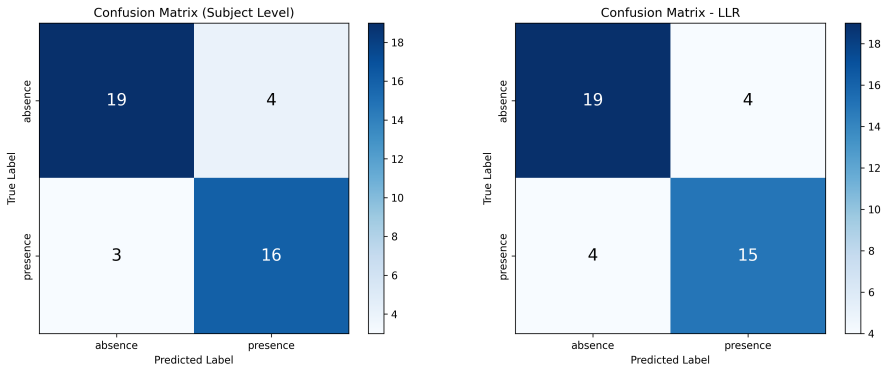


Figure 3.30: UMAP projection, Motor Behavior 7, Exercise 6: Golden Rooster Stands on One Leg.



(a) Random Forest (b) Logistic Regression

Figure 3.31: Confusion Matrices for Motor Behavior 7: Random Forest (a) and Logistic Regression (b), Exercise 6: Golden Rooster Stands on One Leg.

Table 3.28: Classification results reported as macro-averaged metrics, for Motor Behavior 7, Exercise 6.

Metric	**Random Forest**	**Logistic Regression**
Macro F1-score	0.83	0.81
Macro Precision	0.83	0.81
Macro Recall	0.83	0.81

Figures 3.32 and 3.33 provide a comprehensive summary of the evaluated motor behaviors, reporting the Macro F1-score achieved for each.

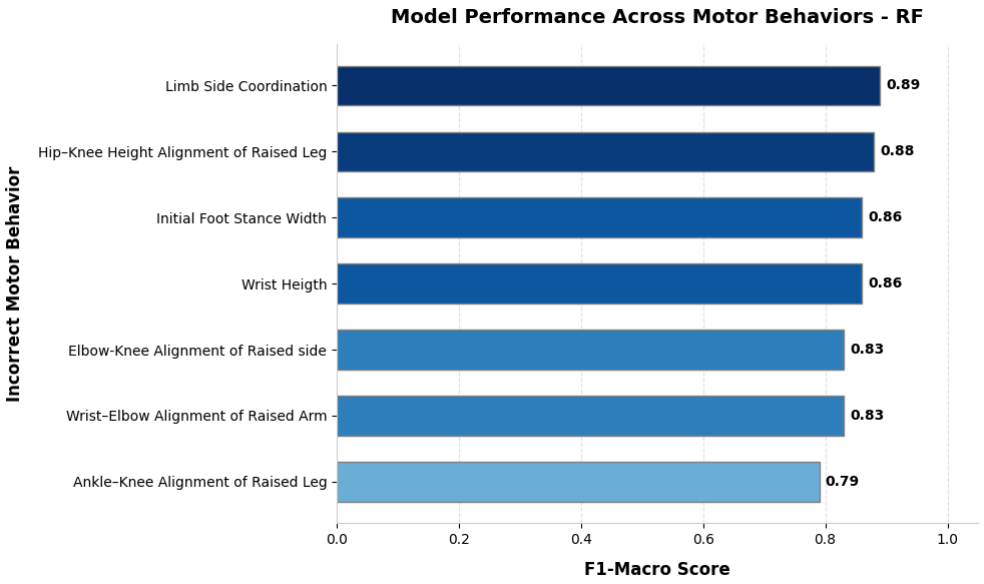


Figure 3.32: Summary of Motor Behavior analysis results for Exercise 6 with Random Forest.

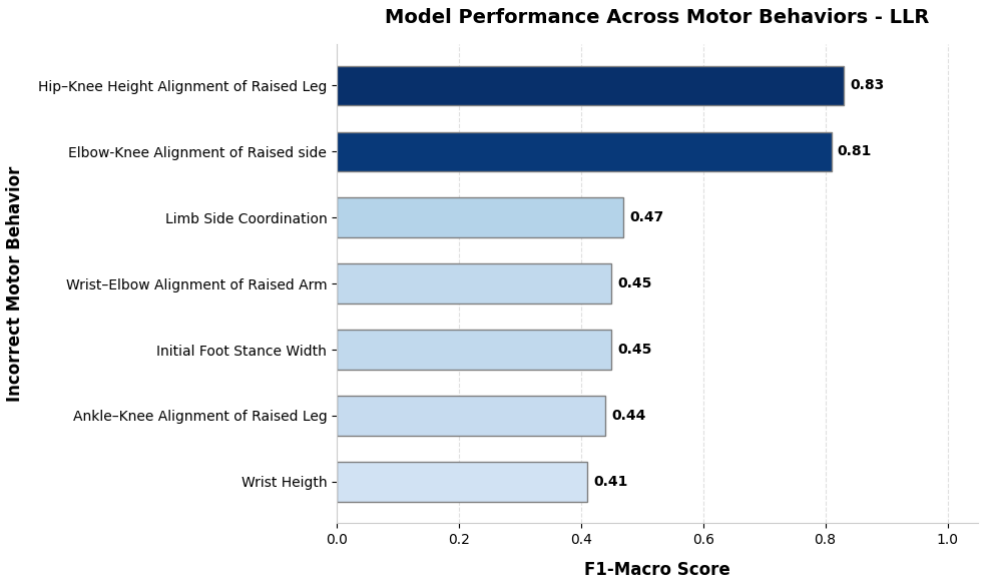


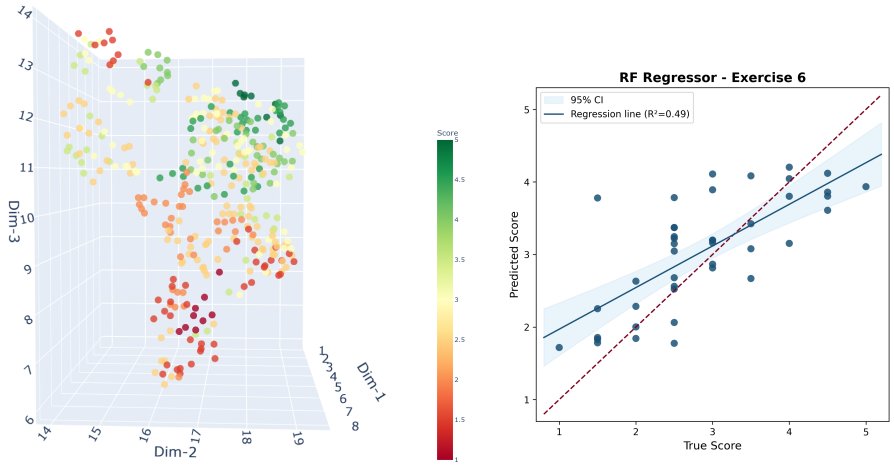
Figure 3.33: Summary of Motor Behavior analysis results for Exercise 6 with Logistic Regression.

3.2.8 Proficiency Assessment

This Subsection presents the results obtained for the proficiency analysis for Exercise 6, Golden Rooster Stands on One Leg. For Exercise 6, only the manual feature selection approach was applied, as described in Subsection 2.4.2.
The features selected for the proficiency assessment are listed in Table 3.29, while the underlying data structure is visualized through UMAP projections in Figure 3.34(a). Furthermore, Table 3.30 presents the final evaluation metrics, specifically the MAE, RMSE, and R^2 calculated at the subject level. The regression performance is further illustrated in Figure 3.34(b), which shows the true versus predicted Proficiency scores together with a fitted regression line. The prediction error per subject is shown in Figure 3.35, highlighting the subjects for which the model over- or under-estimates the proficiency score.

Table 3.29: Selected features for Proficiency evaluation, Exercise 6: Golden Rooster Stands on One Leg.

Selected Features	**Values considered**
Active wrist to shoulder vertical distance	kneepk
Overlap time / knee hold time	single value
Right heel sagittal displacement (Z-axis)	range
Active elbow–knee distance	x_kneepk
Active elbow flexion angle	Δ(start $\rightarrow$ kneepk)
Angle between forearm vector and vertical axis	kneepk
Presence or Absence of all 7 MB	single binary value



(a) UMAP projection (b) True vs Predicted

Figure 3.34: Proficiency assessment, Exercise 6: UMAP projection (a) and regression performance (b), Golden Rooster Stands on One Leg.

Table 3.30: Regression results for Proficiency evaluation, Exercise 6.

Metric	**Results**
MAE	0.540
RMSE	0.690
R^2	0.490

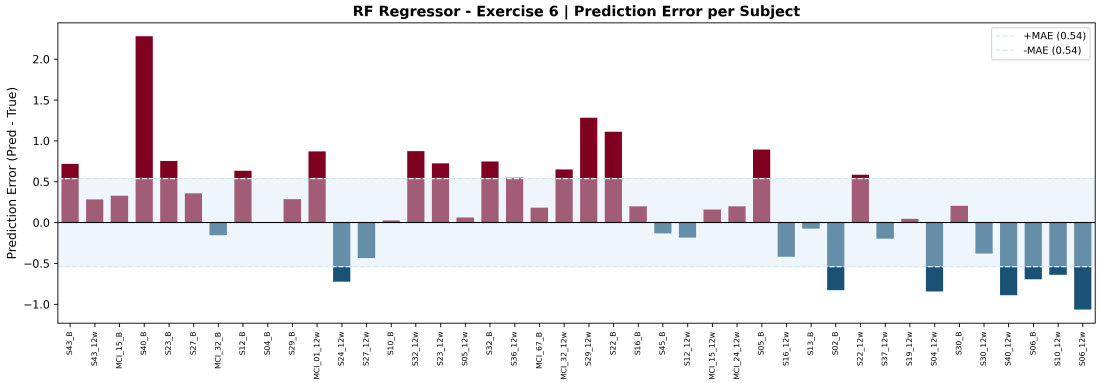


Figure 3.35: Prediction error per subject, Proficiency assessment, Exercise 6: Golden Rooster Stands on One Leg.

Chapter 4

Discussion

In this Chapter, the results presented in Chapter 3 will be discussed.

4.1 Interpretation of Biomechanical Findings

For both exercises, the binary motor behavior classifiers successfully extract biomechanically informative features from 3D kinematic data, enabling reliable detection of movement errors.
Across Exercise 1 (Raising the Power) and Exercise 6 (Golden Rooster Stands on One Leg), the Random Forest classifiers achieved Macro F1-scores ranging from 0.71 to 0.90 and from 0.79 to 0.89, respectively, indicating good predictive performance despite the heterogeneity of the motor behaviors analyzed.
These results are especially encouraging given the complexity of the task: the classifier must generalize to held-out recordings from subjects with different body morphologies and proficiency levels, as ensured by the Leave-One-Group-Out cross-validation strategy.
Performance variation across motor behaviors is not uniform, and depends on both the nature of the kinematic signal associated with each behavior and the reliability of the pose estimation for the body segments involved. Motor behaviors defined by simple, geometrically direct measures tend to show the highest classification performance.
In Exercise 1, MB1 (Wrist Height) achieved the highest Macro F1-score (0.90), reflecting the fact that wrist elevation during bilateral arm raising produces a clear vertical displacement that is easily captured by height-based features. On the contrary, the lowest-performing motor behaviors in Exercise 1 can be attributed to two distinct factors. MB5 (Feet Parallelism, F1 = 0.73) relies on the precise localization of foot and toe joints, which are among the body segments most affected by inaccuracies in the SMPLest-X pose estimation. MB6 (Arm Three-Dimensionality,

F1 = 0.71) presents a different challenge: it describes a complex, three-dimensional property of the arm trajectory that is difficult to capture through scalar features and whose presence or absence is harder to define clearly even at the labeling stage. A similar pattern is observed in Exercise 6, where the lowest-performing motor behavior (MB5, Ankle–Knee Alignment, F1 = 0.79) also involves distal joint estimates, while more directly measurable motor behaviors generally achieve higher scores.

The comparison between Random Forest and Logistic Regression provides additional insight into the nature of the classification problem. For Exercise 1, Logistic Regression achieved competitive or, in some cases, even higher F1-scores than Random Forest for certain motor behaviors, such as MB3, where LR reached an F1-score of 0.97. This suggests that, for specific motor behaviors in Exercise 1, the classification problem is approximately linearly separable.
In contrast, for Exercise 6, Logistic Regression performed worse for most motor behaviors, with F1-scores between 0.41 and 0.47, indicating that the corresponding feature spaces are not linearly separable. An exception is represented by MB6 (Hip–Knee Height Alignment, LR = 0.83) and MB7 (Elbow–Knee Alignment, LR = 0.81), which remained competitive even with a linear classifier, in line with the more linear nature of the kinematic patterns associated with these behaviors. The non-linear structure of Exercise 6 reflects its greater biomechanical complexity: maintaining single-leg balance while coordinating the ipsilateral arm involves several interdependent degrees of freedom. Across both exercises, Random Forest performed best at capturing these non-linear interactions and was therefore selected as the primary model for this analysis.

The Random Forest Regressor for the proficiency assessment achieved a MAE of 0.610 and RMSE of 0.759 with automated nested mRMR feature selection, and a MAE of 0.501 and RMSE of 0.638 with manual feature selection for Exercise 1. For Exercise 6, the manual approach gave a MAE of 0.540 and RMSE of 0.690. The model captures a general trend between predicted and true proficiency scores, as reflected by the positive regression slopes and R^2 values of 0.32 for the nested mRMR approach (Exercise 1) and 0.52 and 0.49 for the manual approach in Exercise 1 and Exercise 6, respectively. The prediction errors, however, remain considerable, highlighting the inherent difficulty of the task. A MAE of 0.50 corresponds to a single step on the proficiency scale, which ranges from 0 to 5 in increments of 0.5. Since the boundary between adjacent score levels is subjective (there is not a strict objective criterion distinguishing, for instance, a score of 2.0 from 2.5), this margin of error may still be acceptable in a clinical decision-support context.
Several factors likely contribute to these results. With six to seven binary motor behaviors defined per exercise, the number of possible motor behavior pattern

combinations is large (up to $2^6 = 64$ and $2^7 = 128$, respectively), yet only 42 recordings were available per exercise. As a result, each subject shows a largely unique combination of motor behavior errors, increasing inter-subject variability and limiting the model's ability to learn consistent patterns. The score distribution is also uneven: for Exercise 1, most subjects are concentrated between 3.0 and 4.0, while for Exercise 6 the majority cluster around 2.5, leaving the extremes of the scale poorly represented in both cases.
A further limitation concerns the nature of the features. The static and dynamic features used in this study were primarily designed to capture specific motor behavior deviations, and may not fully represent more global aspects of movement quality, such as fluidity, rhythm, or overall coordination, that likely also contribute to the expert's proficiency rating.
The better performance of the manual approach compared to nested mRMR suggests that domain knowledge provides a more stable and clinically relevant feature set, particularly when data are scarce. The inclusion of Ground Truth labels as additional predictors likely contributed to this improvement, as they provide the model with explicit information about which motor behaviors were present, directly linking movement errors to the proficiency score. The automated selection is methodologically rigorous but more sensitive to the variability introduced by the small sample size.

4.2 Impact on Tele-Tai Chi Rehabilitation

A system capable of automatically assessing movement correctness and monitoring proficiency over time could meaningfully advance tele-Tai Chi and tele-rehabilitation more broadly.
Tele-Tai Chi has been shown to help prevent and mitigate fall risk [3], particularly in the elderly population. For this reason, improving its effectiveness, especially in home-based settings, is important. As a low-impact form of physical activity, Tai Chi is well suited for remote delivery, as it does not require specific equipment (e.g., weights, elastic bands, or mats) or access to dedicated facilities such as gyms or training centers.
Group-based online sessions also make it possible to reach individuals living in remote or underserved areas, where access to in-person training facilities or qualified instructors may be limited or unavailable.
The goal of the proposed automatic assessment system is to support physiotherapists during tele-rehabilitation sessions by providing an additional reliable tool for evaluating patient performance. Such systems are not intended to replace the role of the physiotherapist or Tai Chi instructor, since human interaction remains

a fundamental component of rehabilitation. Practicing Tai Chi, especially in group settings, has been shown to positively influence mood and reduce symptoms of depression [53].
The proposed system is designed as a complementary tool to assist clinicians and physiotherapists in monitoring patients' movement quality. When this methodology is applied to video recordings collected over different sessions (e.g., weekly sessions), it becomes feasible to identify specific movement errors for each subject. This allows patients to receive targeted feedback on their performance, along with constructive suggestions for improvement in subsequent training sessions.
By comparing proficiency levels over time, it is possible to quantitatively evaluate whether a subject improves their execution across sessions or throughout the duration of the rehabilitation protocol.
Such a system could be integrated into a user-friendly platform available to both physiotherapists and Tai Chi instructors, supporting them during remote sessions. The proposed approach could also be extended to other disciplines similar to Tai Chi, broadening its potential applications in the field of tele-rehabilitation.

4.3 Study Limitations

This work is subject to several limitations. A large portion of these constraints originates from the video acquisition process, while others are related to dataset characteristics, such as limited sample size, uneven class distribution, and the methodology used to define and assign scores.
A key limitation concerns the video data collection process. Recordings were not performed under consistent conditions, both in the laboratory and in the home setting. The camera positioning frequently changed, in terms of both distance from the subject and viewing angle. Some recordings were strictly frontal or semi-lateral, while others involved intermediate or non-standardized viewpoints, introducing unwanted variability into the dataset.
For laboratory-based recordings in particular, a substantial number of videos had to be excluded from the analysis. This occurred either because the subject was already out of frame at the start of the recording or because they moved outside the camera field of view during the task execution. This issue mainly affected the visibility of body extremities, with upper and lower limbs not fully captured. This suggests that a greater distance between the subject and the camera would have been necessary to ensure complete capture of the body during the entire recording.
Another issue is related to the quality of the home-recorded videos. In several cases, the recordings exhibited minor performance quality issues, such as lag or lack of smoothness, which affected the continuity of the motion capture. As a result,

some videos did not meet optimal quality standards. More generally, video quality was often insufficient due to uncontrolled recording conditions, including poor lighting, cluttered environments, people moving in the background, and low-quality recording devices. In some instances, the combination of low illumination and complex backgrounds made it difficult to clearly distinguish the subject from the surroundings, further compromising data quality.

Another limitation concerns how the proficiency scores were assigned. The score is an overall, qualitative evaluation of the subject's execution, and the absence of a fully objective criterion for assigning it represents a limitation. Although there is a clear relationship between the number of observed motor behavior errors and the resulting proficiency score, this relationship is not formally defined through a consistent or quantifiable rule, which limits the objectivity of the scoring process. The accuracy of the SMPLest-X model in tracking extremities, especially the feet and hands, is a further constraint. In several cases, these body parts were not reliably captured, leading to less precise or unstable estimations of their positions. Finally, another limitation is the reduced size of the final dataset, which consists of 42 videos per exercise analyzed, after having excluded problematic recordings. This limited sample size restricts the ability to generalize the results obtained in this work. This limitation could be addressed in future work by expanding the dataset through the acquisition of additional recordings, thereby improving both the robustness and the generalizability of the results.

4.4 Future Developments

Future developments should first focus on extending the current analysis of motor behaviors to the remaining four exercises included in the protocol. This would allow for a more complete evaluation framework and test the generalizability of the proposed approach across a wider range of movement patterns and complexities.

Another important aspect is increasing the size of the available dataset. Increasing the number of participants and recordings would improve the statistical robustness of the results and reduce the risk of bias associated with the current limited sample size. This effort should be combined with the definition of clear and standardized acquisition protocols. In particular, recording conditions should be kept consistent across sessions by fixing parameters such as camera-to-subject distance, ensuring adequate and uniform lighting, and enforcing a frontal and well-aligned subject positioning. These measures would significantly reduce variability introduced by environmental and setup differences, leading to more reliable and comparable data.

An additional relevant improvement concerns the scoring methodology. Future work should aim at developing a more objective and formalized system for assigning proficiency scores. This could involve defining explicit criteria that quantitatively

link observed motor behavior errors to the final score, thereby reducing subjectivity and improving reproducibility. A more structured scoring framework would also improve the reliability of any machine learning model trained on these labels.
Future work could also explore richer, movement-quality-oriented features, capturing aspects such as fluidity, rhythm, or overall coordination, that go beyond the detection of individual motor behavior deviations and may better align with how proficiency is assessed by human experts.
Further research should focus on improving the accuracy of hand and foot tracking, which could be addressed through two complementary directions: adopting more advanced pose estimation models with stronger extremity tracking, or fine-tuning SMPLest-X itself on domain-specific data with more accurate hand and foot annotations. Enhancing the precision of distal joint estimation could enable a more complete and accurate reconstruction of human motion, improving the quality of the extracted kinematic features.
Finally, a promising direction for future work is a two-level architecture: a set of expert classifiers (one per motor behavior) first estimate the probability of each individual error, and these probabilities are then combined by a second-level model to estimate overall proficiency. This cascade mirrors how a clinician reasons, moving from specific, well-defined error detection to an overall judgment, and could simplify the proficiency prediction step by combining a small number of error probabilities, rather than mapping directly from high-dimensional kinematic features, an approach more compatible with the limited sample size available.

Chapter 5

Conclusions

Tai Chi is one of the most extensively validated forms of physical activity for fall prevention in older adults. In the context of developing Tele-Tai Chi interventions aimed at reducing fall risk in this population, the goal of this thesis is to develop Machine Learning algorithms to automatically assess Tai Chi movement quality and proficiency score, starting from video data processed using SMPLest-X for Human Pose Estimation.

Eighty-four videos (forty-two per exercise) of older adults performing Tai Chi exercises in both home and laboratory settings, collected prior to this study, were analyzed. The videos were first processed using SMPLest-X to extract joint trajectories and pose parameters. After trajectory normalization and video segmentation, biomechanically relevant features were extracted. From the initial pool, a specific subset of features was manually selected for each movement. These features were then used as input for a dedicated Random Forest classifier, trained to distinguish between subjects performing a given movement correctly or incorrectly.

This process was repeated across all movements for each exercise.

Finally, a Random Forest Regressor was developed to predict the continuous proficiency score of each subject.

The motor behavior classifiers achieved a mean Macro F1-score of 0.81 for Exercise 1 and 0.85 for Exercise 6, demonstrating reliable detection of movement errors across a heterogeneous dataset. Proficiency assessment proved to be more challenging: the Random Forest Regressor achieved a MAE of 0.501 and R^2 of 0.52 for Exercise 1, and a MAE of 0.540 and R^2 of 0.49 for Exercise 6, both with manual feature selection, reflecting the inherent difficulty of predicting a continuous expert score from a limited dataset.

This study demonstrated the feasibility of combining computer vision and machine learning for the assessment of Tai Chi proficiency, thereby enabling and supporting the future development of tele-rehabilitation, despite the inherent challenges of video-based pose estimation.

The results of this thesis were presented in a poster session at the 2026 Science of Tai Chi & Qigong as Whole Person Health Conference, held in Boston (USA) on April 30 and May 1, 2026.

Appendix A

Appendix

This appendix contains, for each motor behavior of the two exercises, the ranking of extracted features obtained via Random Forest mean decrease in impurity (Gini importance), averaged across cross-validation folds, together with boxplots of the top-ranked features. The boxplots are complemented by Linear Mixed-Effects (LME) statistical tests, with the motor behavior label as a fixed effect and a random intercept per subject to account for the nine repetitions per subject. These analyses were used only to confirm the discriminative value of the features, and not as a feature selection criterion, which was performed manually.

A.1 Ex 1

A.1.1 Motor Behavior 1

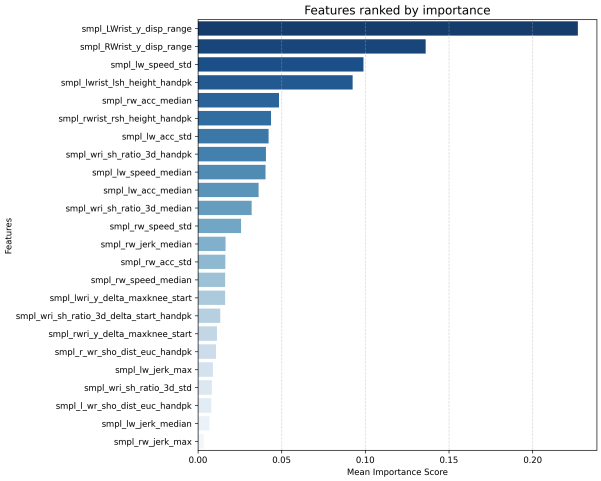


Figure A.1: Feature ranking — Exercise 1, Motor Behavior 1.

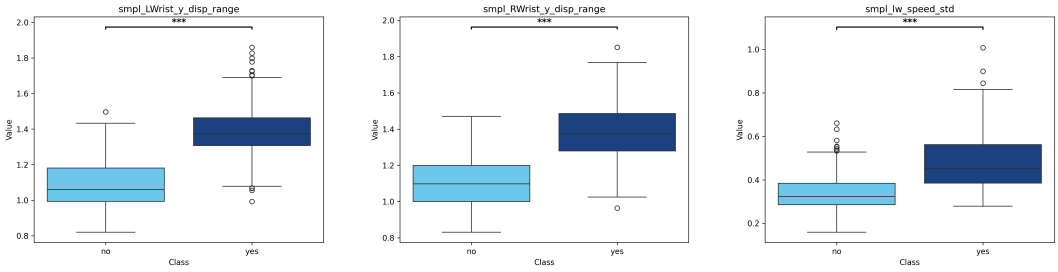


Figure A.2: Top 3 features by importance — Exercise 1, Motor Behavior 1.

A.1.2 Motor Behavior 2

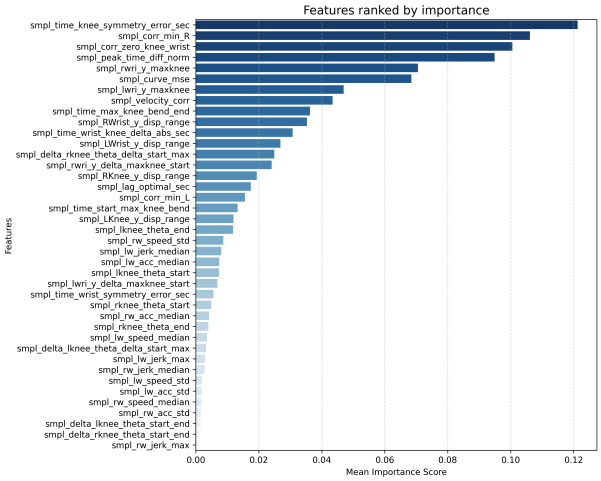


Figure A.3: Feature ranking — Exercise 1, Motor Behavior 2.

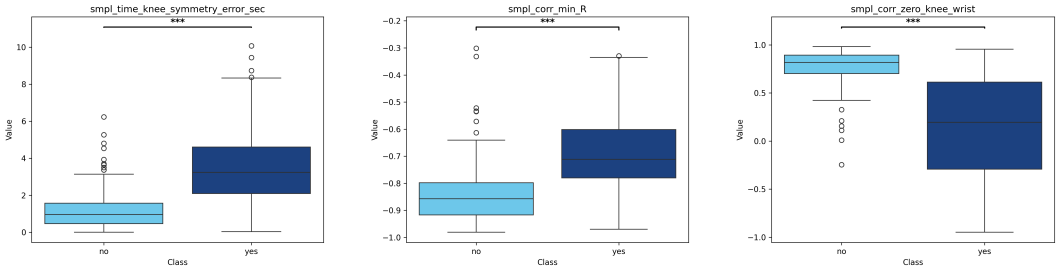


Figure A.4: Top 3 features by importance — Exercise 1, Motor Behavior 2.

A.1.3 Motor Behavior 3

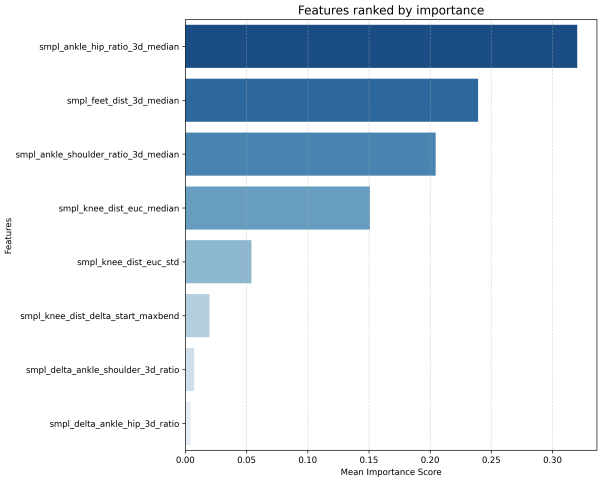


Figure A.5: Feature ranking — Exercise 1, Motor Behavior 3.

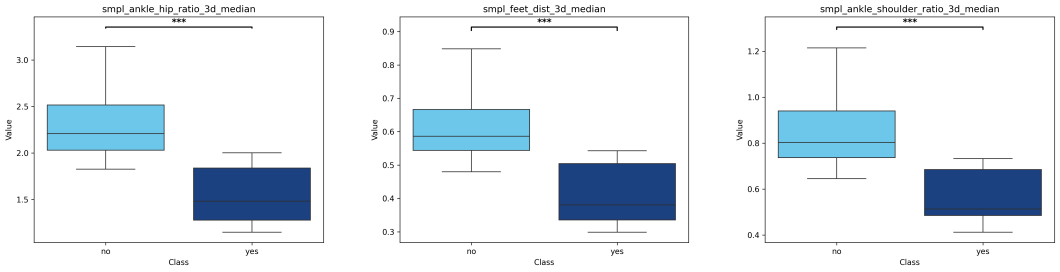


Figure A.6: Top 3 features by importance — Exercise 1, Motor Behavior 3.

A.1.4 Motor Behavior 4

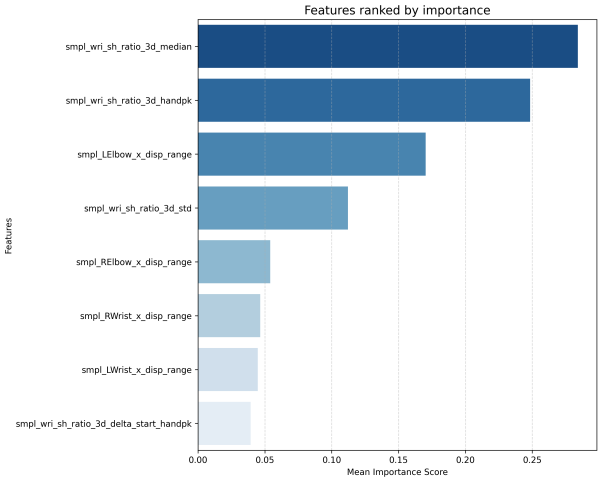


Figure A.7: Feature ranking — Exercise 1, Motor Behavior 4.

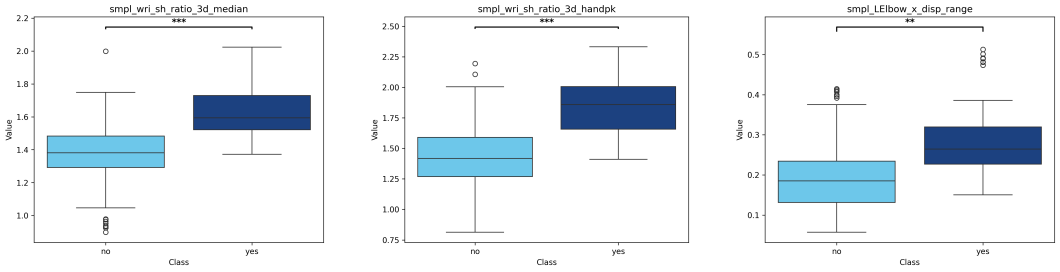


Figure A.8: Top 3 features by importance — Exercise 1, Motor Behavior 4.

A.1.5 Motor Behavior 5

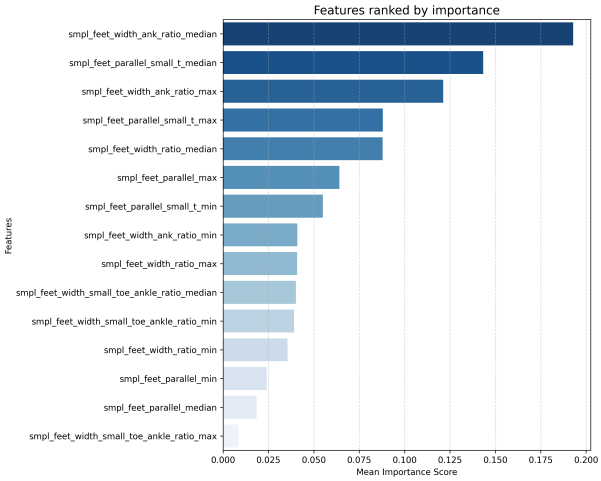


Figure A.9: Feature ranking — Exercise 1, Motor Behavior 5.

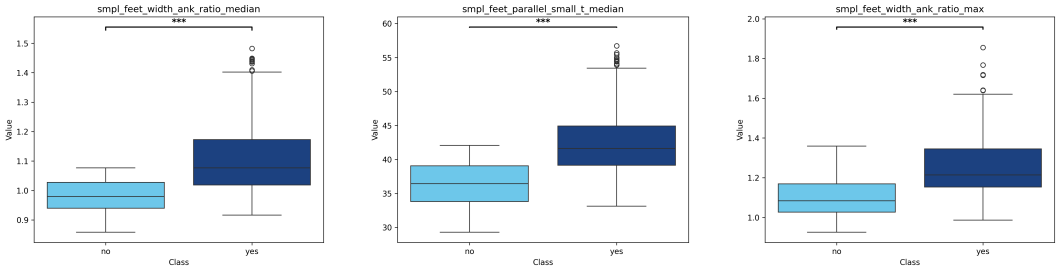


Figure A.10: Top 3 features by importance — Exercise 1, Motor Behavior 5.

A.1.6 Motor Behavior 6

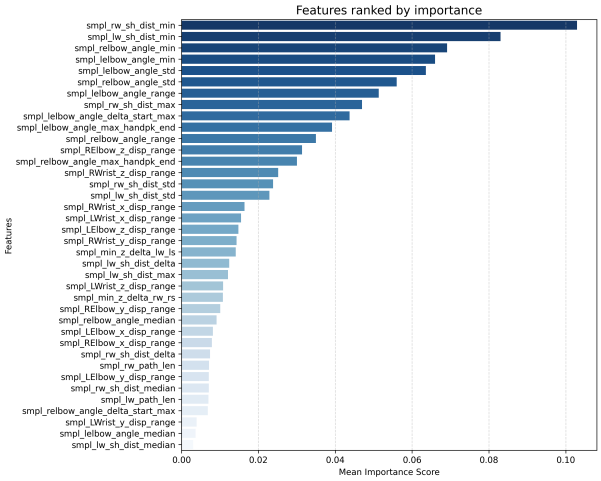


Figure A.11: Feature ranking — Exercise 1, Motor Behavior 6.

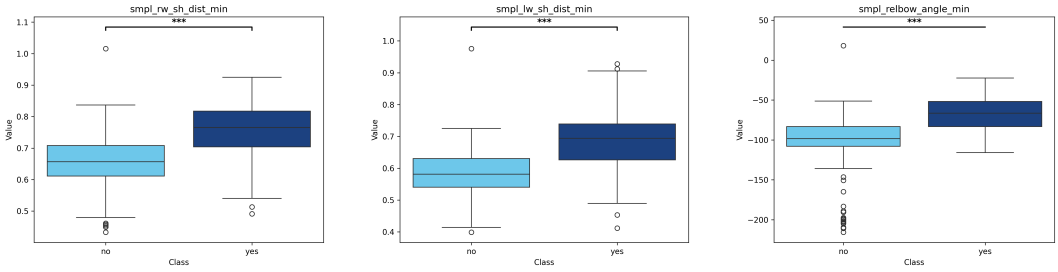


Figure A.12: Top 3 features by importance — Exercise 1, Motor Behavior 6.

A.2 Ex 6

A.2.1 Motor Behavior 1

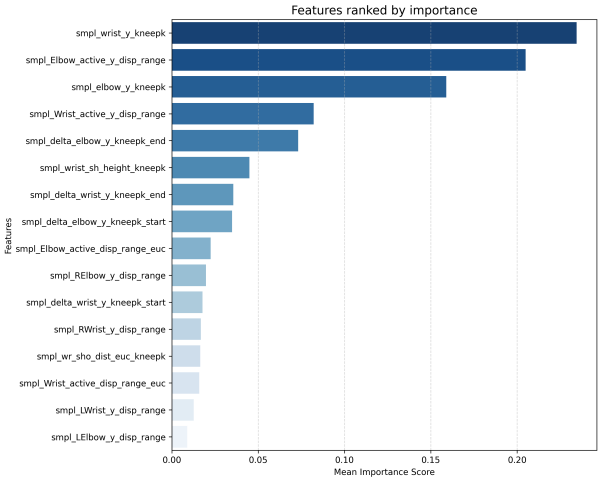


Figure A.13: Feature ranking — Exercise 6, Motor Behavior 1.

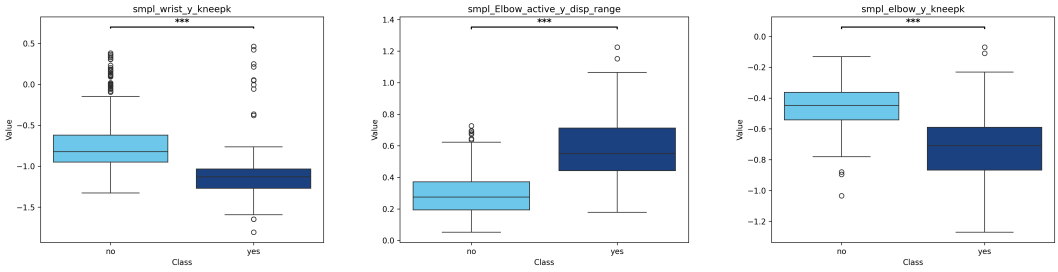


Figure A.14: Top 3 features by importance — Exercise 6, Motor Behavior 1.

A.2.2 Motor Behavior 2

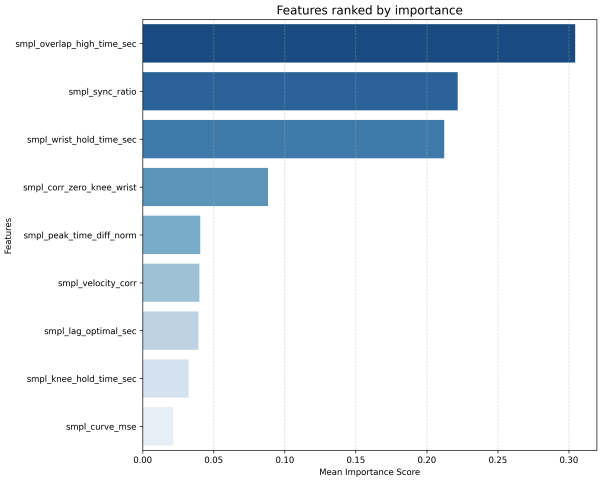


Figure A.15: Feature ranking — Exercise 6, Motor Behavior 2.

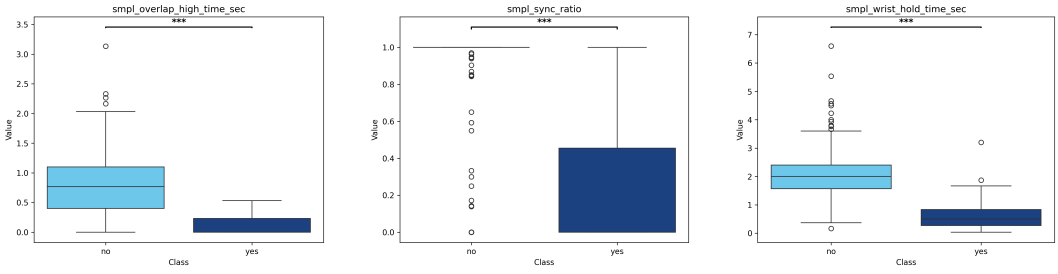


Figure A.16: Top 3 features by importance — Exercise 6, Motor Behavior 2.

A.2.3 Motor Behavior 3

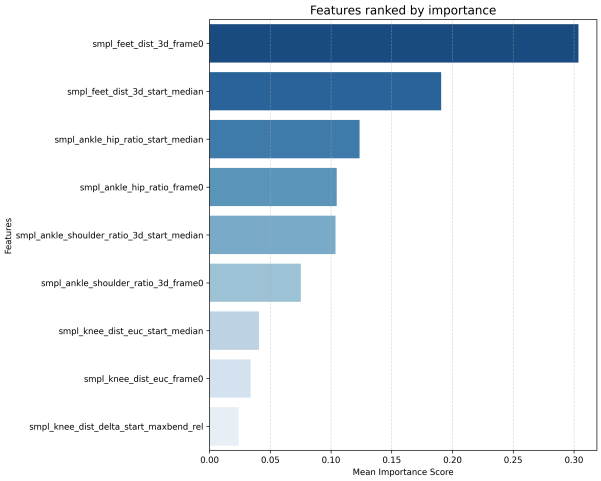


Figure A.17: Feature ranking — Exercise 6, Motor Behavior 3.

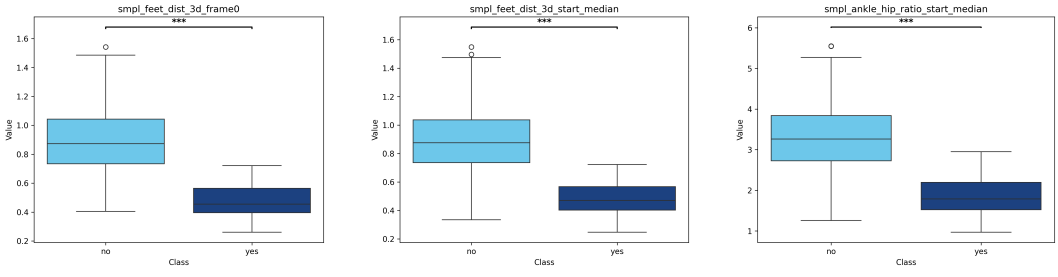


Figure A.18: Top 3 features by importance — Exercise 6, Motor Behavior 3.

A.2.4 Motor Behavior 4

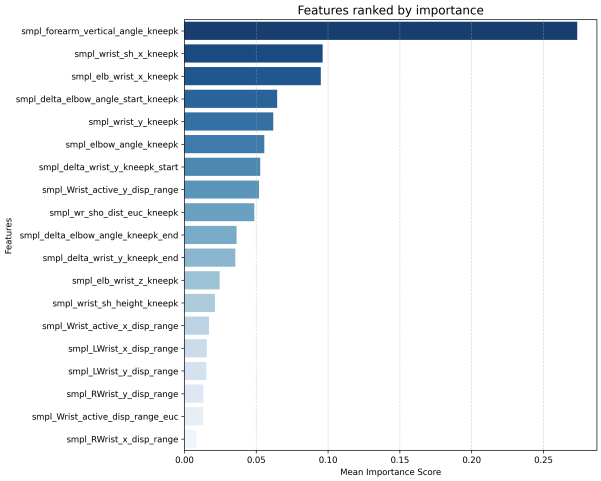


Figure A.19: Feature ranking — Exercise 6, Motor Behavior 4.

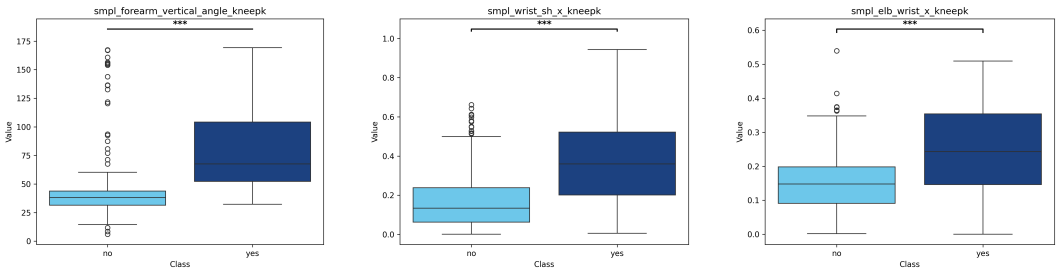


Figure A.20: Top 3 features by importance — Exercise 6, Motor Behavior 4.

A.2.5 Motor Behavior 5

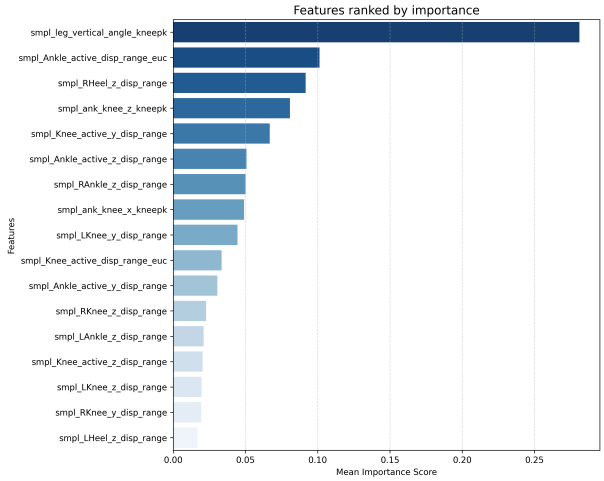


Figure A.21: Feature ranking — Exercise 6, Motor Behavior 5.

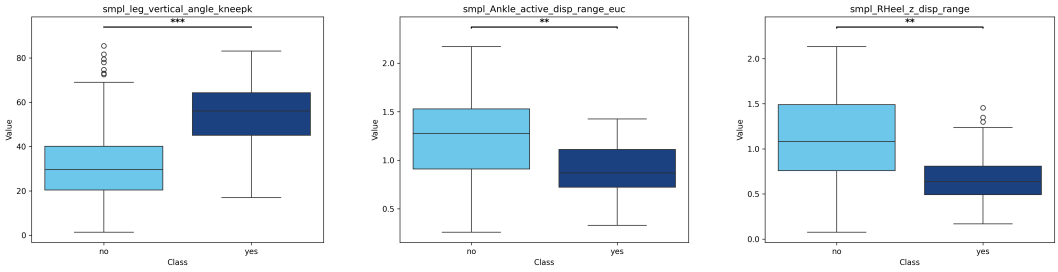


Figure A.22: Top 3 features by importance — Exercise 6, Motor Behavior 5.

A.2.6 Motor Behavior 6

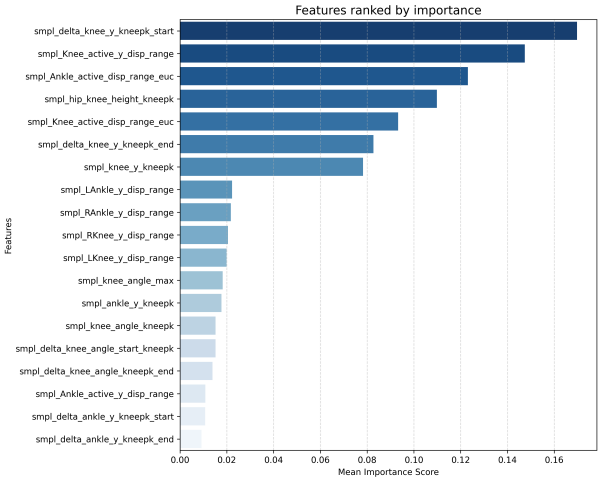


Figure A.23: Feature ranking — Exercise 6, Motor Behavior 6.

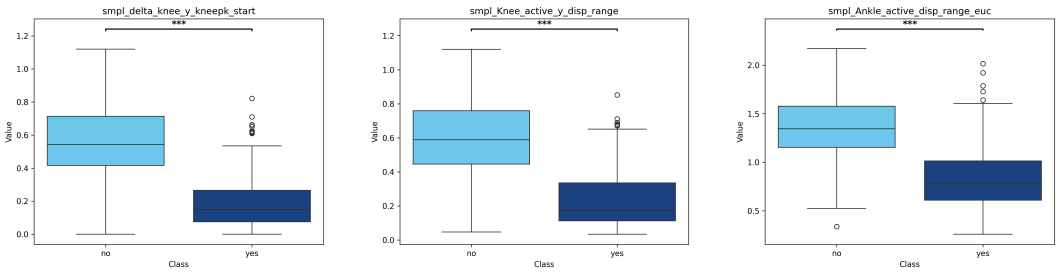


Figure A.24: Top 3 features by importance — Exercise 6, Motor Behavior 6.

A.2.7 Motor Behavior 7

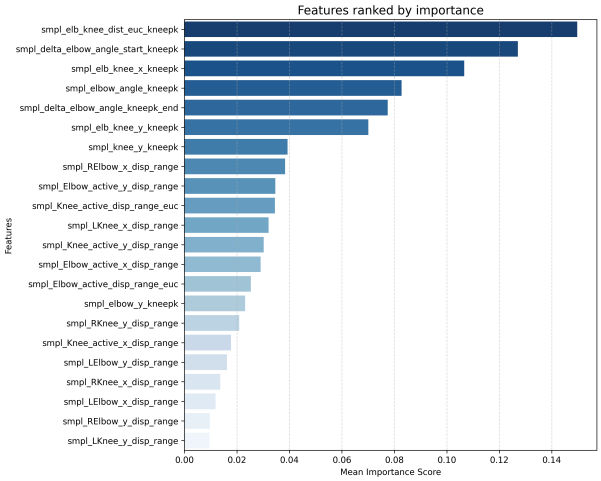


Figure A.25: Feature ranking — Exercise 6, Motor Behavior 7.

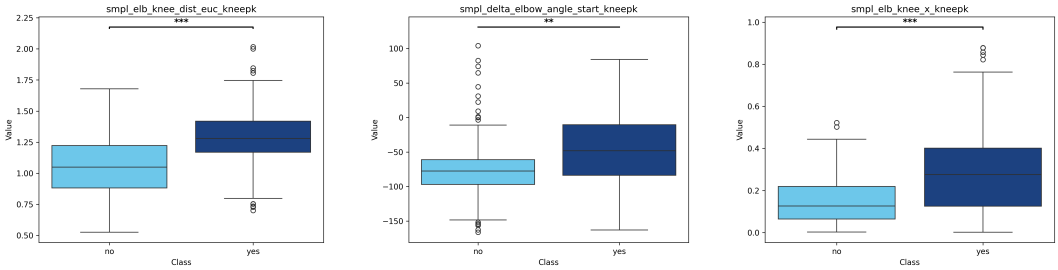


Figure A.26: Top 3 features by importance — Exercise 6, Motor Behavior 7.

Bibliography

[1] Rachel E. Ward, Lien Quach, Sarah A. Welch, Suzanne G. Leveille, Elizabeth Leritz, and Jonathan F. Bean. «Interrelated Neuromuscular and Clinical Risk Factors That Contribute to Falls». In: *Journals of Gerontology* (2019) (cit. on p. 7).

[2] Weidong Chen, Min Li, Hai Li, Yanzhao Lin, and Zhoushan Feng. «Tai Chi for fall prevention and balance improvement in older adults: a systematic review and meta-analysis of randomized controlled trials». In: *Frontiers in Public Health* (2023) (cit. on p. 7).

[3] Yu-Ning Hu, Yu-Ju Chung, Hui-Kung Yu, Yu-Chi Chen, Chien-Tsung Tsai, and Gwo-Chi Hu. «Effect of Tai Chi Exercise on Fall Prevention in Older Adults: Systematic Review and Meta-analysis of Randomized Controlled Trials». In: *International Journal of Gerontology* (2016) (cit. on pp. 7, 87).

[4] Rafael Lomas-Vega, Esteban Obrero-Gaitán, Francisco J. Molina-Ortega, and Rafael Del-Pino-Casado. «Tai Chi for Risk of Falls. A Meta-analysis». In: *Journal of the American Geriatrics Society* (2017) (cit. on p. 7).

[5] Rachael D. Seidler, Jessica A. Bernard, Taritonye B. Burutolu, Brett W. Fling, Mark T. Gordon, Joseph T. Gwin, Youngbin Kwak, and David B. Lipps. «Motor Control and Aging: Links to Age-Related Brain Structural, Functional, and Biochemical Effects». In: *Neuroscience and Biobehavioral Reviews* (2011) (cit. on p. 8).

[6] Yurun Cai, Qu Tian, Alden L. Gross, Hang Wang, Jian-Yu E, Yuri Agrawal, Eleanor M. Simonsick, Luigi Ferrucci, and Jennifer A. Schrack. «Motor and Physical Function Impairments as Contributors to Slow Gait Speed and Mobility Difficulty in Middle-Aged and Older Adults». In: *Journals of Gerontology* (2022) (cit. on p. 9).

[7] David A. Ganz and Nancy K. Latham. «Prevention of Falls in Community-Dwelling Older Adults». In: *The New England Journal of Medicine* (2020) (cit. on p. 9).

[8] Bruce H. Alexander, Frederick P. Rivara, and Marsha E. Wolf. «The cost and frequency of hospitalization for fall-related injuries in older adults». In: *American Journal of Public Health* (1992) (cit. on p. 9).

[9] Sari Stenholm, Annemarie Koster, Heli Valkeinen, Kushang V. Patel, Stefania Bandinelli, Jack M. Guralnik, and Luigi Ferrucci. «Association of Physical Activity History With Physical Function and Mortality in Old Age». In: *Journals of Gerontology* (2016) (cit. on p. 9).

[10] Eric Yiou, Teddy Caderby, Arnaud Delafontaine, Paul Fourcade, and Jean-Louis Honeine. «Balance control during gait initiation: State-of-the-art and research perspectives». In: *World Journal of Orthopedics* (2017) (cit. on p. 9).

[11] Liwei Xing et al. «Falls caused by balance disorders in the elderly with multiple systems involved: Pathogenic mechanisms and treatment strategies». In: *Frontiers in Neurology* (2023) (cit. on p. 9).

[12] Anne Felicia Ambrose, Geet Paul, and Jeffrey M. Hausdorff. «Risk factors for falls among older adults: A review of the literature». In: *Maturitas* (2013) (cit. on p. 9).

[13] Catherine Sherrington et al. «Exercise to prevent falls in older adults: an updated systematic review and meta-analysis». In: *Br J Sports Med* (2017) (cit. on p. 9).

[14] Alexander Voukelatos, Robert G. Cumming, Stephen R. Lord, and Chris Rissel. «A Randomized, Controlled Trial of Tai Chi for the Prevention of Falls: The Central Sydney Tai Chi Trial». In: *Journal of the American Geriatrics Society* (2007) (cit. on p. 10).

[15] Fei-Chi Yang, Aishwarya B. Desai, Pelareh Esfahani, Tatiana V. Sokolovskaya, and Doreen J. Bartlett. «Effectiveness of Tai Chi for Health Promotion of Older Adults: A Scoping Review of Meta-Analyses». In: *American Journal of Lifestyle Medicine* (2022) (cit. on p. 10).

[16] Ching Lan, Steven L. Wolf, and William W. N. Tsang. «Tai Chi Exercise in Medicine and Health Promotion». In: *Evidence-Based Complementary and Alternative Medicine* (2013) (cit. on p. 10).

[17] Boonsri Kittichittipanich and Wichitra Kusoom. «The Effectiveness of Tai Chi Exercise in Improving Balance and Preventing Falls Among Older Adults». In: *The Open Public Health Journal* (2019) (cit. on p. 11).

[18] Fuzhong Li, Peter Harmer, Kathleen Fitzgerald, Elizabeth Eckstrom, Ronald Stock, Johnny Galver, Gianni Maddalozzo, and Sara S. Batya. «Tai Chi and Postural Stability in Patients with Parkinson's Disease». In: *The New England Journal of Medicine* (2012) (cit. on p. 11).

[19] Pamela Seron et al. «Effectiveness of Telerehabilitation in Physical Therapy: A Rapid Overview». In: *Physical Therapy & Rehabilitation Journal* (2021) (cit. on p. 11).

[20] Benoit Nicolas, Emilie Leblong, Bastien Fraudet, Philippe Gallien, and Patrice Piette. «Telerehabilitation solutions in patient pathways: An overview of systematic reviews». In: *DIGITAL HEALTH* (2024) (cit. on p. 11).

[21] Muhammad Abrar Ashraf, Shaheryar Najam, Touseef Sadiq, Shabbab Algamdi, Hanan Aljuaid, Hameedur Rahman, and Ahmad Jalal. «A Novel Telerehabilitation System for Physical Exercise Monitoring in Elderly Healthcare». In: *IEEE Access* (2025) (cit. on p. 12).

[22] Ge Wu and Lawrence Malcolm Keyes. «Group Tele-Exercise for Improving Balance in Elders». In: *TELEMEDICINE AND e-HEALTH* (2006) (cit. on p. 12).

[23] Elena T. Remillard, Tracy L. Mitzner, and Kara T. Mumma. «Tele Tai Chi for people aging with mobility disabilities: Novel methodology and structured adaptation approach». In: *Contemporary Clinical Trials Communications* (2025) (cit. on p. 12).

[24] Dimitrios-Sokratis Komaris, Georgia Tarfali, Brendan O'Flynn, and Salvatore Tedesco. «Unsupervised IMU-based evaluation of at-home exercise programmes: a feasibility study». In: *BMC Sports Science, Medicine and Rehabilitation* (2022) (cit. on p. 12).

[25] Andrea Cereatti, Reed Gurchiek, Annegret Mündermann, Silvia Fantozzi, Fay Horak, Scott Delp, and Kamiar Aminian. «ISB recommendations on the definition, estimation and reporting of joint kinematics in human motion analysis applications using wearable inertial measurement technology». In: *Journal of Biomechanics* (2024) (cit. on p. 12).

[26] Federico Roggio, Bruno Trovato, Martina Sortino, and Giuseppe Musumeci. «A comprehensive analysis of the machine learning pose estimation models used in human movement and posture analyses: A narrative review». In: *Heliyon* (2024) (cit. on p. 13).

[27] Kishor Das, Thiago de Paula Oliveira, and John Newell. «Comparison of markerless and marker-based motion capture systems using 95% functional limits of agreement in a linear mixed-effects modelling framework». In: *Scientific Reports* (2023) (cit. on p. 13).

[28] Valentin Bazarevsky, Ivan Grishchenko, Karthik Raveendran, Tyler Zhu, Fan Zhang, and Matthias Grundmann. *BlazePose: On-device Real-time Body Pose Tracking*. arXiv preprint arXiv:2006.10204v1. 2020 (cit. on p. 13).

[29] Camillo Lugaresi et al. *MediaPipe: A Framework for Building Perception Pipelines*. arXiv preprint arXiv:1906.08172. 2019 (cit. on p. 13).

[30] Ronny Votel and Na Li. «Next-Generation Pose Detection with MoveNet and TensorFlow.js». In: *TensorFlow Blog* (2021). Google Research. URL: `https://blog.tensorflow.org/2021/05/next-generation-pose-detection-with-movenet-and-tensorflowjs.html` (cit. on p. 14).

[31] Zhe Cao, Gines Hidalgo, Tomas Simon, Shih-En Wei, and Yaser Sheikh. «OpenPose: Realtime Multi-Person 2D Pose Estimation using Part Affinity Fields». In: *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2019) (cit. on p. 15).

[32] Mukhiddin Toshpulatov, Wookey Lee, Suan Lee, and Arousha Haghighian Roudsari. «Human pose, hand and mesh estimation using deep learning: a survey». In: *The Journal of Supercomputing* (2022) (cit. on p. 16).

[33] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J. Black. «SMPL: A Skinned Multi-Person Linear Model». In: *ACM Transactions on Graphics* (2015) (cit. on p. 16).

[34] Georgios Pavlakos, Vasileios Choutas, Nima Ghorbani, Timo Bolkart, Ahmed A. A. Osman, Dimitrios Tzionas, and Michael J. Black. *Expressive Body Capture: 3D Hands, Face, and Body from a Single Image*. arXiv:1904.05866v1. 2019 (cit. on p. 17).

[35] Wanqi Yin et al. *SMPLest-X: Ultimate Scaling for Expressive Human Pose and Shape Estimation*. arXiv:2501.09782v1. 2025 (cit. on pp. 17, 28).

[36] Peter Wayne and Mark Fuerst. *The Harvard Medical School Guide to Tai Chi*. Shambhala, 2013 (cit. on pp. 21–25).

[37] Harvard Health Publishing. *Traditional Tai Chi Elements*. `https://youtu.be/PqxNQmdcp1w`. YouTube video. 2017 (cit. on p. 25).

[38] Livia Colucci. «Development of Human Pose Estimation and Machine Learning-based algorithms for assessing physical exercise proficiency». MA thesis. Turin: Politecnico di Torino, 2023 (cit. on p. 25).

[39] Chris Ding and Hanchuan Peng. «Minimum Redundancy Feature Selection from Microarray Gene Expression Data». In: *Journal of Bioinformatics and Computational Biology* 3 (2005) (cit. on p. 41).

[40] Aasim Ayaz Wani. «Comprehensive review of dimensionality reduction algorithms: challenges,limitations, and innovative solutions». In: *PeerJ Computer Science* (2025) (cit. on p. 41).

[41] Leland McInnes, John Healy, and James Melville. *UMAP: Uniform Manifold Approximation and Projection for Dimension Reduction*. arXiv preprint arXiv:1802.03426. 2018 (cit. on p. 42).

[42] Craig Starbuck. «The Fundamentals of People Analytics». In: Springer, 2023, pp. 223–238 (cit. on p. 43).

[43] Maher Maalouf. «Logistic Regression in Data Analysis: An Overview». In: *International Journal of Data Analysis Techniques and Strategy* (2011) (cit. on p. 43).

[44] Leo Breiman. «Random Forests». In: *Machine Learning* (2001) (cit. on p. 44).

[45] Leo Breiman. «Bagging Predictors». In: *Machine Learning* (1996) (cit. on pp. 44, 45).

[46] Leo Breiman, Jerome Friedman, Charles J. Stone, and R.A. Olshen. *Classification and Regression Trees*. Taylor & Francis, 1984 (cit. on p. 44).

[47] Adele Cutler, D. Richard Cutler, and John R. Stevens. «Random Forests». In: *Machine Learning* (2011) (cit. on p. 44).

[48] Gerard Biau. «Analysis of a Random Forests Model». In: *Journal of Machine Learning Research* (2013) (cit. on pp. 44, 46).

[49] Sudhir Varma and Richard Simon. «Bias in error estimation when using cross-validation for model selection». In: *BMC Bioinformatics* (2006) (cit. on p. 46).

[50] Gareth James, Daniela Witten, Trevor Hastie, and Robert Tibshirani. «An Introduction to Statistical Learning». In: Springer, 2021. Chap. 5 (cit. on p. 48).

[51] Hendrio Bragança, Juan G. Colonna, Horácio A. B. F. Oliveira, and Eduardo Souto. «How Validation Methodology Influences Human Activity Recognition Mobile Systems». In: *Sensors* (2022) (cit. on p. 48).

[52] Juri Opitz and Sebastian Burst. *Macro F1 and Macro F1: A Note*. arXiv preprint arXiv:1911.03347. 2021 (cit. on p. 50).

[53] Chenchen Wang, Raveendhara Bannuru, Judith Ramel, Bruce Kupelnick, Tammy Scott, and Christopher H. Schmid. «Tai Chi on psychological well-being: systematic review and meta-analysis». In: *BMC Complementary and Alternative Medicine* (2010) (cit. on p. 88).