PhD Thesis of

CY Cergy Paris Université

Doctoral School N° 405
Economics, Management, Mathematics,
Theoretical Physics, and Computer Science
Speciality: *Sciences and Technologies of Information and Communication*

By

Claire BÉRANGER

Gait Recognition and Fall Prediction with Deep-Learning on Micro-Doppler Signatures

Thesis presented and defended in Cergy, on the 12th of January 2026
Research Laboratory: Équipes Traitement de l'Information et Systèmes (ETIS) - UMR 8051

Pre-defense reviewers:

Iness AHRIZ	Associate Professor, HDR, Conservatoire National des Arts et Métiers Paris, CEDRIC
Catherine DEZAN	Associate Professor, HDR, Université de Bretagne Occidentale, Lab-STICC

Composition of the Jury:

Chair of the Defense:	Hedi TABIA	Professor, Université Paris-Saclay, IBISC
Examiner:	Rémi DUBOIS	Professor, Université de Bordeaux, IHU LIRYC
	Thanh Phuong NGUYEN	Professor, Université of Côte d'Azur, I3S
	Petr DOBIÁŠ	Associate Professor, CY Cergy Paris Université, ETIS
	Ngoc-Son VU	Associate Professor, HDR, ENSEA, ETIS
Thesis Supervisor:	Olivier ROMAIN	Professor, CY Cergy Paris Université, ETIS

Guest(s):

David GUYARD	CEO, BlueLinea
Julien LE KERNEC	Senior Lecturer, University of Glasgow, ETIS

Résumé

Le risque de chute lié au vieillissement de la population constitue un enjeu majeur de santé publique, en particulier dans les pays développés. En France, plus de 10 000 décès et 136 000 hospitalisations par an concernent des personnes âgées de plus de 65 ans. Ces événements exercent une pression importante sur le système de santé, tant sur le plan de la mobilisation du personnel que de la gestion des capacités d'accueil, et affectent fortement l'autonomie et la qualité de vie des personnes concernées.

Pour prévenir ces incidents à domicile et rassurer les familles, des systèmes de téléassistance ont été déployés. Toutefois, leur efficacité reste limitée par un taux élevé de fausses alertes, générant une charge supplémentaire pour les services et une anxiété chez les utilisateurs. Des études récentes ont montré que l'analyse des mouvements à l'aide de capteurs non intrusifs pouvait fournir des informations précieuses sur les activités quotidiennes, permettant de détecter des comportements inhabituels ou des anomalies dans la marche et les mouvements.

Parmi ces technologies, le radar suscite un intérêt croissant, car il permet d'extraire des signatures micro-Doppler visualisables sous forme de spectrogrammes, offrant une riche source d'information pour différencier les activités tout en respectant la vie privée. Dans ce cadre, cette thèse développe des méthodes de reconnaissance d'activités quotidiennes à partir de données radar non conventionnelles, en collaboration avec l'entreprise de téléassistance BlueLinea. Les travaux combinent l'identification de signatures Doppler caractéristiques avec le développement et l'évaluation de modèles de classification capables de reconnaître l'activité courante.

Un jeu de données expérimental multimodal a été constitué, comprenant 100 participants et plus de trois séquences par activité. Trois modèles de classification ont été développés : un réseau neuronal convolutif (CNN), ResNet-18, pour l'extraction automatique des caractéristiques, une machine à vecteurs de support (SVM) appliqué aux descripteurs géométriques et un K plus proches voisins (KNN) pour les descripteurs texturaux. Une stratégie d'optimisation complète a été mise en œuvre, incluant le prétraitement des spectrogrammes (normalisation et redimensionnement), le réglage des hyperparamètres de chaque modèle et la fusion des résultats par vote majoritaire pour combiner les forces

des approches et renforcer la robustesse du système.

Les performances obtenues sont quantitativement significatives : chaque modèle atteint une précision comprise entre 90% et 99%, tandis que la fusion des résultats permet d'atteindre un taux global de reconnaissance des activités d'environ 95%. Ces résultats démontrent à la fois la fiabilité des méthodes développées et le potentiel des capteurs radar pour le suivi non intrusif des activités humaines. Ils ouvrent ainsi des perspectives concrètes pour l'amélioration des systèmes de téléassistance et la prévention des chutes, tout en offrant de nouvelles pistes pour la recherche en suivi comportemental et analyse biomécanique.

Abstract

The risk of falls associated with an ageing population is a major public health issue, particularly in developed countries. In France, more than 10,000 deaths and 136,000 hospitalisations per year involve people over the age of 65. These events place significant pressure on the healthcare system, both in terms of staff mobilisation and capacity management, and have a substantial impact on the independence and quality of life of those affected.

To prevent these incidents at home and reassure families, teleassistance systems have been deployed. However, their effectiveness remains limited due to a high rate of false alarms, generating additional workload for services and anxiety among users. Recent studies have shown that movement analysis using non-intrusive sensors can provide valuable information about daily activities, enabling the detection of unusual behaviour or abnormalities in gait and movement.

Among these technologies, radar is attracting growing interest as it can extract micro-Doppler signatures visualised as spectrograms, providing a rich source of information for differentiating activities while preserving privacy. In this context, this thesis develops methods for recognising daily activities from unconventional radar data, in collaboration with the teleassistance company BlueLinea. The work combines the identification of characteristic Doppler signatures with the development and evaluation of classification models capable of recognising daily activities.

A multimodal experimental dataset was compiled, comprising 100 participants and more than three sequences per activity. Three classification models were developed: a Convolutional Neural Network (CNN), ResNet-18, for automatic feature extraction, a Support Vector Machine (SVM) applied to geometric descriptors, and a K-Nearest Neighbour (KNN) for textural descriptors. A comprehensive optimisation strategy was implemented, including spectrogram pre-processing (normalisation and resizing), hyperparameter tuning for each model, and majority voting to combine the strengths of the approaches and enhance system robustness.

The results are quantitatively significant: each model achieves an accuracy of between 90% and 99%, while the fusion of results reaches an overall activity recognition rate of

approximately 95%. These findings demonstrate both the reliability of the methods developed and the potential of radar sensors for non-intrusive monitoring of human activities. They also offer concrete prospects for improving teleassistance systems and fall prevention, while providing new avenues for research in behavioural monitoring and biomechanical analysis.

ACKNOWLEDGEMENTS

Completing a thesis requires great personal discipline, but above all, it is made possible by the encounters, collaborations, and support received throughout the journey.

First, I am deeply grateful to my thesis supervisors: Professor Olivier Romain, Associate Professor HDR Son Vu, and Associate Professor Petr Dobiáš. Thank you for your support, trust, and advice, which were essential to the smooth progress of my work. I am especially grateful to Petr Dobiáš for joining my supervisory team midway through the project and for his daily guidance, which was crucial to the continuity and completion of my thesis.

I also wish to thank the project collaborators, Senior Lecturer Julien Le Kernec and Dr. David Guyard, for their follow-up, advice, and external supervision. Their support provided valuable insight into the subject and a better understanding of real-world conditions.

My thanks go to the members of the jury for the time they devoted to evaluating my work: my reviewers Associate Professor HDR Iness Ahriz and Associate Professor HDR Catherine Dezan, as well as Professor Hedi Tabia, Professor Rémi Dubois, and Professor Thanh Phuong Nguyen. I am also grateful to the Île-de-France Region, BlueLinea, and CY Cergy Paris Université for their institutional, technical, and material support.

I am grateful to the members of the ETIS laboratory for their kindness, warm welcome, and the enriching exchanges over the years. Although I cannot mention everyone by name, these encounters will remain solid benchmarks for the future. Special thanks to Camille, Inbar, Aymeric, Stéphane Z., and Guillaume R. for their support to PhD students and the moments we shared.

To the PhD colleagues and friends, thank you for the moments we shared and for your trust. Representing you on the laboratory council has been a real honour. I am thinking particularly of: Arnau, Étienne, Tarek, Pauline, Théo, Louis LH., Louis D., Louis A., Antoine, Ishak, Nicolas, Ouliana, Marco, Mathilda, and Mathieu. I greatly appreciate everything we have shared, your ideas, your attention, and your support.

A special mention goes to my office neighbours, Cécile and Sekou. Sharing the daily grind of the thesis with you has been invaluable. Thank you, Cécile, for our conversations

during difficult times. I wish you all the best, Sekou, at the beginning of your journey.

My experience of the laboratory council would not have been the same without Quitterie. Thank you for everything we shared and for your support. I am delighted with the team we formed and all that we achieved for the PhD students and the laboratory. Your encouragement and general support also helped me greatly.

I am sincerely grateful to everyone who contributed to the creation of our dataset: Lili, Konstantina, Antoine Junior, and Omar, without you, everything would have been different. I also wish to thank all the participants from various backgrounds who helped collect the data. The exchanges with you were among the most memorable moments of this thesis.

Thanks also go to the members and PhD colleagues and friends of the SmartGait project, particularly Antoine Junior, Omar, Alexandre, and Anouar, the project engineer. Your good humour and everything we built together were a real source of support.

Special thanks to Alexandre and Laure. Without you, this thesis would certainly not have had the same dimension. Your presence, support, advice, and everything we went through together will remain unique. Having become almost like the musketeers of the laboratory, we defended our theses a few months apart. Thank you for all those moments.

I am also grateful to my friends for their support despite my silences and difficult periods. Particular thanks to Clara and Sylvain, former PhD colleagues, for your advice, our discussions, walks, meals, and everything we shared. Thanks also to François-Laurent, Laurent, Pierre-Louis, and everyone else who supported me.

Finally, I thank my family for their daily support throughout all these years. I am grateful to my parents for their presence, listening, and patience, and to my brother for repeatedly trying to pull me away from work. I am happy to have shared all these activities and moments with you.

To everyone who supported me, showed kindness, and encouraged me throughout my thesis: your presence has been a source of strength and inspiration.

Table of Contents

Résumé iii

Abstract v

Acknowledgements vii

Glossary xix

1 Introduction 1

1.1 Research Context: Fall Risk Assessment and the SmartGait Project 2

1.2 Development and Contributions 4

1.3 Manuscript Structure 7

2 Clinical Context and Medical Problem Statement 9

2.1 Falls in Elderly People: Risks and Interventions 10

2.1.1 Ageing and Fall Risk: Consequences and Strategies 11

2.1.2 Structure for Fall Prevention and Management 17

2.1.3 Patient Monitoring 24

2.2 Medical Challenge and Innovations 29

2.2.1 Clinical Solutions and Emerging Approaches 29

2.2.2 Evolution of Healthcare and Support Systems in France 32

2.2.3 Caregivers and Assistive Technologies 36

2.3 Biomechanics, Behavioural Analysis, and Clinical Strategies 37

2.3.1 Gait Analysis: Biomechanical Principles and Applications 37

2.3.2 Assessment of Risk Through Activity and Behaviour Analysis . . . 41

2.3.3 Institutional Vision and Developments in Clinical Practice 44

2.4 Problem Statement and Conclusion 45

3 Radar-Based Human Activity Recognition: Review of Concepts, Techniques, and Datasets 49

3.1 Human Activity Recognition: Foundations and Technologies 51

3.1.1 Activity Recognition: Concepts and Applications 51
3.1.2 Sensor-Based Systems 56
3.1.3 Radar-Based Activity Recognition 60
3.2 Data Processing and Learning Techniques 65
3.2.1 Data Processing and Evaluation 66
3.2.2 Learning Approaches for Human Activity Recognition 71
3.2.3 Advanced Data Processing Techniques 75
3.3 Radar Human Activity Datasets: Review and Study Dataset 77
3.3.1 Available Radar-Based Human Activity Datasets 78
3.3.2 Data Challenges and Limitations in Radar Human Activity Recognition 83
3.3.3 Dataset Description and Pre-Processing 86
3.4 Conclusion 90

4 Performance Evaluation of Feature Extraction and Recognition Methods 93
4.1 Experimental Evaluation of CNN for Radar-Based Human Activity Recognition 95
4.1.1 Preliminary Evaluation of Classification Architectures 95
4.1.2 Implementation and Performance Evaluation of ResNet-18 96
4.1.3 Impact of Image Resizing 102
4.1.4 Evaluation on Laboratory Data 105
4.2 Feature Extraction and Recognition from Spectrograms 107
4.2.1 Processing and Extraction of Activity from Images 107
4.2.2 Spectrogram-Based Geometric Form Recognition 113
4.2.3 Spectrogram-Based Textural Feature Recognition 118
4.3 Optimisation and Evaluation of Learning with Diverse Features 123
4.3.1 Impact of Data and Pre-Processing on Learning 124
4.3.2 SVM Optimisation and Application for Geometric Forms 128
4.3.3 KNN Optimisation and Application for Textural Features 132
4.3.4 Evaluation of the Fusion by Majority Vote 136
4.4 Conclusion 139

5 MAD: Creation of a Multimodal Dataset and Classification Analysis 141
5.1 Movement Platform and Data Processing 143

5.1.1 Movement Platform Configuration and Equipment . . . 143
5.1.2 Pre-Processing Chain . . . 149
5.1.3 Motion Capture as a Ground Truth: Equipment, Acquisition and Processing . . . 152
5.2 Preparation and Creation of a Multimodal Dataset: MAD . . . 156
5.2.1 Ethical and Organisational Preparation of the Study . . . 157
5.2.2 Experimental Planning and Protocol . . . 158
5.2.3 *MAD* Dataset: Management, Processing and Content . . . 161
5.3 Evaluation of Models and Pre-Processing Strategies on the MAD Dataset . 166
5.3.1 Data Description . . . 166
5.3.2 Evaluation of Model Performance and Pre-Processing Strategies on the MAD Dataset . . . 168
5.3.3 Learning Chain: Model Parameter Optimisation and Performance Evaluation . . . 173
5.3.4 Fusion of Models by Majority Vote . . . 177
5.4 Conclusion . . . 178

6 Conclusion and Perspectives **181**

Publications **187**

Bibliography **189**

A Ethical and Organisational Aspects of the Dataset Creation **213**
A.1 Submission of the Dataset Creation Project to the Ethics Committee . . . 213
A.2 Team and Session Management . . . 214
A.3 Communication and Call for Participants . . . 215
A.4 Participant Selection Criteria . . . 216
A.5 Sessions Planning . . . 217
A.6 Research Ethics Committee Approval . . . 219

B Complementary Details on Experimental Protocol and MAD Dataset **221**
B.1 Experimental Setup: Additional Logistical Details . . . 221
B.2 Experimental Accessories . . . 222
B.3 Session Schedule and Participant Protocol: Operational Details . . . 223
B.4 Data Management Interface and API . . . 224

B.5 Detailed Participant Metadata and Survey Responses 226

List of Figures

1.1 Overview of SmartGait . . . 3
1.2 Pre-Processing Chain and Classification for Activity Recognition . . . 5
1.3 Overview of the Full Thesis Chain . . . 6

2.1 Overview of Chapter 2 . . . 10
2.2 Population Pyramid of France Projected for 2025 with 2010 for Comparison 13
2.3 Collaborative Network of Healthcare Professionals in Fall Prevention . . . 19
2.4 Examples of Regional Fall Prevention Pathways and Professional Involvement 20
2.5 Illustration of the *Berg Balance Scale* (BBS) Test . . . 22
2.6 Illustration of the Timed Up and Go (TUG) Test . . . 23
2.7 Preliminary Fall Risk Assessment Algorithm . . . 28
2.8 Location of Sensors on the Body . . . 30
2.9 Historical Evolution of Elderly Care Institutions in France . . . 33
2.10 Walk Phases . . . 38
2.11 Summary of Societal Context, Chapter 2 . . . 47

3.1 Overview of Chapter 3 . . . 50
3.2 Skeletal Representations of the Same Posture Captured from Different Viewpoints, Positions, and Camera Angles . . . 53
3.3 Taxonomy of Human Activity Recognition . . . 54
3.4 Generic Steps in Human Activity Recognition . . . 56
3.5 Richness VS Perceived Privacy . . . 61
3.6 Overall Operation of Radar Data Acquisition . . . 62
3.7 Overview of Radar Processing Chain . . . 63
3.8 Spectrogram Correspondance with Walk Phases . . . 67
3.9 Comparison of Walk and Fall Signatures . . . 67
3.10 Micro-Doppler Spectrogram Representation of Limbs . . . 68
3.11 Binarisation of Rice Image . . . 69
3.12 Contour Detection of Cube Image . . . 69

3.13 Data Extraction from MoCap for Skeletal Reconstruction and Motion Analysis . . . 71
3.14 Set of Human Aquatic Data . . . 80
3.15 Set of Military and Diverse Human Activities Data . . . 81
3.16 Set of Hospital and Daily Human Activities Data . . . 83
3.17 Pre-Processing Chain Method . . . 88
3.18 *Radar Signatures of Human Activities* Spectrogram Matrix . . . 90
3.19 Summary of Technical Context, Chapter 3 . . . 91

4.1 Overview of Chapter 4 . . . 94
4.2 Simplified Diagram of the ResNet-18 Architecture with Residual Blocks . . . 97
4.3 Sample Images from the *CIFAR-10* Dataset . . . 99
4.4 ResNet-18 Normalised Confusion Matrix: 6 Activities . . . 101
4.5 Image Resizing Procedure for ResNet-18 Training . . . 102
4.6 Accuracy per Activity as a Function of Input Size . . . 103
4.7 Overall Accuracy as a Function of Input Size . . . 104
4.8 Normalised Confusion Matrix on Four Micro-Doppler Classes with a Model Pre-Trained on Six Different Activities . . . 106
4.9 Noise Removal in Spectrogram Data . . . 108
4.10 Incorrect Spectrogram Data . . . 108
4.11 Projection of the Energy of the Spectrogram Image . . . 109
4.12 Crop by Energy Projection . . . 109
4.13 Crop Problem Using Min and Max Method . . . 110
4.14 Correction of the Crop Issue Using the Min and Max Method . . . 110
4.15 Complex Cases with the Crop Method . . . 111
4.16 Representation of Integral Calculation Under Curve . . . 112
4.17 Spectrogram Analysis, from the Original to the Parameters Extraction . . . 114
4.18 Geometrical Form Extraction for SVM Training . . . 116
4.19 Grayscale Spectrogram . . . 119
4.20 Gray Level Coocurence Matrix of Spectrogram: Walk Activity . . . 122
4.21 Haralick Parameters Extraction for KNN Training . . . 122
4.22 ResNet-18 Normalised Confusion Matrix: 5 Activities . . . 125
4.23 Denoise Spectrogram Example . . . 125
4.24 ResNet-18 Normalised Confusion Matrix: 5 Activities with Denoise . . . 126
4.25 Crop Spectrogram Example . . . 127

4.26 ResNet-18 Normalised Confusion Matrix: 5 Activities with Crop 127
4.27 Explained Variance per PCA for SVM Learning 130
4.28 SVM Model with RBF Kernel 131
4.29 SVM Model with Polynomial Kernel and PCA 132
4.30 Explained Variance per PCA for KNN Learning 134
4.31 KNN Normalised Confusion Matrix 135
4.32 KNN Normalised Confusion Matrix with PCA 136
4.33 Global Chain for Training with Majority Voting 137
4.34 Majority Voting Normalised Confusion Matrix 138
4.35 Summary of Data Processing and Classification Methods, Chapter 4 140

5.1 Overview of Chapter 5 142
5.2 Set-Up of the *Movement Platform* 144
5.3 SDR 980AD2 Ancortek FMCW Radar 144
5.4 Vicon Vero V2.2 Camera 145
5.5 Vicon Vue Camera 145
5.6 GAITRite Pro Mat 146
5.7 EVK1-VGA from Prophesee Event-Based Camera 146
5.8 Illustrations of the Components in the Instrumented Chair Setup 147
5.9 Radar Setup Diagram 148
5.10 Magic Wand V2 148
5.11 Pre-Processing Chain with FMCW Radar 150
5.12 Our Velocity Time Graph (Spectrogram) 150
5.13 Motion Capture Markers in Real and their *Nexus*'s Representation 152
5.14 Top-View Schematic of Activity Layout within the *Movement Platform* . . 153
5.15 *MoCap* Velocity-Time Graph of *Walk* 154
5.16 Range-Time Representation of *Walk* Activity Detected by *Radar* 154
5.17 Overlaid Spectrogram of *Walk* and *MoCap* Data (Velocity-Time Graph) . . 156
5.18 *Movement Platform*'s Dimensions Measured in January 2024 159
5.19 Sensor's Configuration for Synchronisation 160
5.20 *MAD* Data Tree Structure 162
5.21 Top-View Schematic of Activity Directions within the *Movement Platform* 164
5.22 Mosaic of our 9 Activities, Performed in All Directions (T, B, RL, D) . . . 165
5.23 Normalised Confusion Matrix of *Radar Signatures of Human Activities* Dataset as Input to our Pre-Processing Chain 168

5.24 Spectrograms Generated with Different Filters and Biases 169
5.25 Normalised Confusion Matrix of *ResNet-18* with *MAD5* 2sscateg 170
5.26 Normalised Confusion Matrix of *ResNet-18* with *MAD5* 3sscateg 171
5.27 Normalised Confusion Matrix of *ResNet-18* with *MAD9* 2sscateg 172
5.28 Normalised Confusion Matrix of *ResNet-18* with *MAD9* 3sscateg 172
5.29 Normalised Confusion Matrix of *SVM* with *MAD5* 2sscateg 175
5.30 Normalised Confusion Matrix of *KNN* with *MAD5* 2sscateg 176
5.31 Normalised Confusion Matrix of *Fusion* with *MAD5* 2sscateg 177
5.32 Summary of MAD Constitution and Learning Chain, Chapter 5 179

A.1 Communication Poster for the MAD Creation Project 215
A.2 Registration Process for Participants 217

B.1 Dimensions of the *Movement Platform* Measured in March 2023 222
B.2 List of Equipment Used for Dataset Activities 223
B.3 Interface Used for Completing the Form 225
B.4 Functioning of the API Utilised for the Dataset 226

List of Tables

2.1 Distribution of the French Population by Age Group in 2010 12
2.2 Iso-Resource Groups (GIR): 6 Levels of Dependence and Required Assistance 15
2.3 Roles of Healthcare Professionals in Fall Prevention, Ordered by Level of Patient Contact 18
2.4 Summary of the Main Clinical Tests Used for Fall Risk Assessment 24
2.5 Comparison of Elderly Care Institutions in France 34
2.6 Evolution of Gait Analysis Methods over Time 39

3.1 General Comparison of Sensor Categories for Activity Recognition 57
3.2 Overview of Denoising Methods Applied to Radar Spectrograms and their Reported Benefits 70
3.3 Classification Models Using the Dataset from *Radar Signatures of Human Activities* 73
3.4 Spectrogram Datasets 79
3.5 Training Set Distribution Details 89

4.1 Classification Models for Radar-Based System 96
4.2 Comparison of Training and Testing Protocols for CIFAR-10 and *Radar Signatures of Human Activities* 98
4.3 Accuracy Comparison Between State-of-the-Art and our Implmentation on CIFAR-10 99
4.4 Cross-Validation Results 100
4.5 Comparison of Accuracy and F1-score Rates Achieved with Thumbnail Resizing 102
4.6 Summary of Acquired Test Activities 105
4.7 Comparison of Loss Rates Between the Last Two Crop Methods 113
4.8 Training Set Distribution Details 117
4.9 Description of Haralick's 14 Parameters 121
4.10 Accuracy Comparison Between Original, Denoise and Crop Methods 128

4.11 Accuracy for RBF Kernel Hyperparameter Combinations 129
4.12 Accuracy for Polynomial Kernel Hyperparameter Combinations 129
4.13 Optimum Hyperparameters of KNN 133
4.14 Summary of the Highest Accuracies Achieved on the *Radar Signatures of Human Activities* Dataset 138

5.1 Description of *MAD* Activity List 158
5.2 Participant Distribution 163
5.3 Dataset Overview 166
5.4 MAD5 and MAD9 Specifications 167
5.5 Activity Distribution for 2 and 3 Subcategories Classification 167
5.6 Comparison of Accuracies for Different Biases for MAD5 170
5.7 Performance Comparison Between Image Processing Methods 173
5.8 *SVM* Performance Comparison Between Combination Search and *PCA* Methods 174
5.9 *KNN* Performance Comparison Between Combination Search and *PCA* Methods 176
5.10 Summary of the Highest Accuracies Achieved on the *MAD5* 2sscateg Dataset 177

B.1 Protocol Stages with Approximate Durations 224
B.2 Percentage of Sedentary and Physically Active among the Population of MAD 226
B.3 Population Characteristics 227

GLOSSARY

➢ ADL: **A**ctivities of **D**aily **L**iving
➢ AlexNet: Convolutional Neural Network
➢ AOI: **A**rea **o**f **I**nterest
➢ API: **A**pplication **P**rogramming **I**nterface
➢ BBS: **B**erg **B**alance **S**cale
➢ BMI: **B**ody **M**ass **I**ndex
➢ CNN: **C**onvolutional **N**eural **N**etwork
➢ CW: **C**ontinuous **W**ave
➢ DCNN: **D**eep **C**onvolutional **N**eural **N**etwork
➢ DNS: **D**omain **N**ame **S**ystem
➢ EHPAD: **É**tablissement d'**H**ébergement pour **P**ersonnes **Â**gées **D**épendantes
➢ Evento: Event Management and Planning Tool
➢ FFT: **F**ast **F**ourier **T**ransform
➢ FIR: **F**inite **I**mpulse **R**esponse
➢ FMCW: **F**requency **M**odulated **C**ontinuous **W**ave
➢ FSR: **F**orce **S**ensing **R**esistor
➢ GAN: **G**enerative **A**dversarial **N**etwork
➢ GIR: **I**so-**R**esource **G**roups (French dependency classification)
➢ GLCM: **G**ray **L**evel **C**o-occurrence **M**atrix
➢ GoogLeNet: Convolutional Neural Network
➢ GPIO: **G**eneral-**P**urpose **I**nput/**O**utput
➢ RGB Image: **R**ed **G**reen **B**lue Image
➢ IMU: **I**nertial **M**easurement **U**nit
➢ INSEE: **N**ational **I**nstitute of **S**tatistics and **E**conomic **S**tudies
➢ KNN: **K**-**N**earest **N**eighbour
➢ MAD: **M**ultimodal **A**ctivity **D**ataset
➢ MLP: **M**ulti-**L**ayer **P**erceptrons
➢ MoCap: **Mo**tion **Cap**ture
➢ MTI: **M**oving **T**arget **I**ndication

- PCA: **P**rincipal **C**omponent **A**nalysis
- Tinetti POMA: Tinetti **P**erformance **O**riented **M**obility **A**ssessment
- PSNR: **P**eak **S**ignal to **N**oise **R**atio
- Radar: **Ra**dio **D**etection **a**nd **R**anging
- RBF: **R**adial **B**asis **F**unction
- ResNet-18: Convolutional Neural Network
- RF: **R**andom **F**orest
- Rx1/Rx2: Radar Receiving Antennas
- SMA: **S**ub**M**iniature version **A** (connector)
- SPWVD: **S**moothed **P**seudo-**W**igner–**V**ille **D**istribution
- STFT: **S**hort **T**ime **F**ourier **T**ransform
- SSH: **S**ecure **Sh**ell Protocol
- SVM: **S**upport **V**ector **M**achine
- Teleassistance: Remote Assistance System for Dependent People
- TUG: **T**imed **U**p and **G**o test
- UWB: **U**ltra-**W**ide**b**and
- Vicon: Motion Capture System
- VM: **V**irtual **M**achine
- WHO: **W**orld **H**ealth **O**rganisation

Chapter 1

Introduction

The health of a population represents a major challenge for governments across the world. Every country must contend with the emergence of both chronic and acute diseases, as well as the gradual *ageing* of its population. This *ageing* process is often accompanied by signs of frailty, the onset and progression of which vary from one individual to another. To manage this phenomenon and safeguard the health of *elderly people*, a range of measures has been implemented, including retirement homes, nursing homes, and *teleassistance* services. Government initiatives, such as the *Anti-Fall* Plan [21], have also been introduced with the aim of promoting *prevention* and reducing associated risks.

Among the various pathologies and risks associated with *ageing*, *falls* are one of the leading causes of hospitalisation and mortality among elderly people. However, this risk could be mitigated through improved support and the use of more effective *detection* tools. At present, clinical services, nursing staff, and caregivers face excessive workloads and persistent staff shortages, which compromise the quality of care. Although *clinical tests* are available to assess the *risk of falls*, their implementation remains challenging in real-world care settings.

Conventional *detection* devices, such as cameras or pendant sensors [76, 169, 174], are widely used but present several limitations, particularly regarding *privacy* concerns and *reliability*. In recent years, new approaches based on less intrusive sensing technologies, such as *Radar*, have been explored [57]. These aim to provide solutions that are both privacy-preserving and reliable. The characterisation of *human activities* using motion information extracted from *Radar* signals shows considerable potential in this field.

The literature reports various processing chains that enable the extraction of features such as *distance* and *velocity* associated with human activities from raw *Radar* data [183]. Research has also been conducted on the use of *classification* models to recognise activities based on this piece of information [164]. However, most of these studies do not target specific contexts, such as *fall* risk *prevention*, but rather focus on demonstrating the general feasibility of *radar*-based activity *classification*. Datasets and studies specifically

dedicated to *fall* risk *detection* and *prevention* remain scarce or incomplete. The scientific objectives of this thesis are as follows:

1. To effectively *recognise* and *detect* daily activities using suitable *classification* methods (Chapters 2 to 4):
   - Designing models capable of identifying clinically relevant movements associated with *fall risk*;
   - Distinguishing between different types of daily activities with high reliability;
2. To *characterise* the extracted data through appropriate pre-processing techniques (Chapters 3 and 4):
   - Applying data processing strategies that enhance the quality and interpretability of *Radar* measurements;
   - Extracting meaningful features to support detailed activity analysis;
3. To propose an *end-to-end* processing chain adapted to *clinical* settings (Chapters 2 to 5):
   - Integrating data acquisition, pre-processing, and classification into a coherent workflow that ensures accurate, non-intrusive, and clinically relevant evaluation of mobility and fall risk;
4. To construct a *new dataset*, evaluate its performance, and compare it with existing datasets (Chapter 5):
   - Creating a multimodal dataset that captures realistic daily activities;
   - Enabling robust evaluation of learning strategies;
   - Supporting reproducibility and benchmarking.

1.1 Research Context: Fall Risk Assessment and the SmartGait Project

In *France*, approximately *20%* of the *elderly* population is at increased *risk* of falling due to age-related declines in motor skills and balance [117]. *Falls* are a major public health concern, causing both *physical* consequences (such as fractures and loss of independence) and *psychological* repercussions (including loss of confidence and social isolation). Fall risk is typically assessed using targeted *clinical tests* when concerns arise about a patient's gait or balance, evaluating postural stability and gait quality.

To complement these *clinical* assessments, various *technologies* have been developed. Wearable devices, such as gyroscope- or accelerometer-equipped watches and pendants, can support healthcare professionals [76, 169, 174], but they can be *restrictive* and exhibit very high *false alarm* rates (>90%) [107]. Vision-based systems offer an alternative, yet sensitivity to lighting and privacy concerns *limit* their acceptability [26]. By contrast, *Radar* technology is insensitive to light, non-intrusive, and capable of detecting micro-movements. *Radar* signals can be converted into micro-Doppler *spectrograms* for analysing and recognising activities within the observed environment [164].

The *ETIS* laboratory (UMR 8051, CY Cergy Paris Université, ENSEA, CNRS) conducts both theoretical and experimental research on the design of intelligent solutions, ranging from information theory to data analysis, at both the algorithmic and embedded system levels. Its research addresses major societal challenges in the fields of health, communication, security, and mobility. Within *ETIS*, the *CELL* team focuses on intelligent embedded systems for biomedical applications, with particular expertise in machine learning in constrained contexts, such as limited datasets or restricted computational resources.

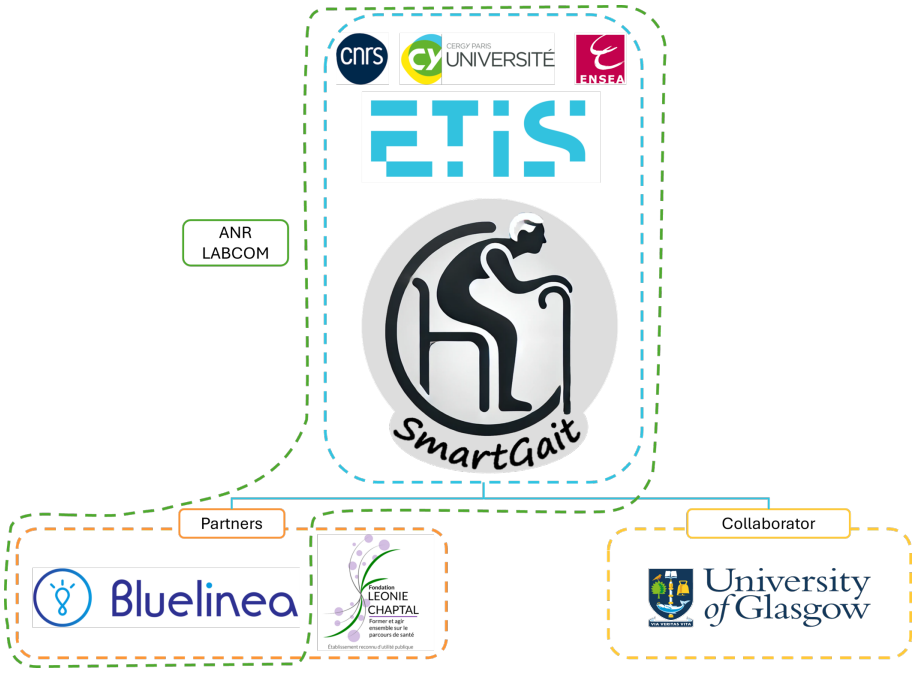

Figure 1.1 – Overview of SmartGait

In this context, the *SmartGait* project was initiated to address the needs of elderly people in terms of monitoring and support for maintaining independence, both in nursing homes and in home settings. The project brings together several stakeholders, including

the *ETIS* laboratory, the *teleassistance* company *BlueLinea*, and *partners* from the *clinical* sector, as illustrated in Figure 1.1. Close collaboration with hospitals and nursing homes provides a framework for developing the proposed solutions, while regular consultations help to better understand the needs of end users and to validate the relevance of the project's objectives.

BlueLinea, a French company specialising in remote *assistance* and *monitoring*, contributes to the project by providing its expertise in the field. The company offers solutions aimed at improving the quality of life of older adults or individuals experiencing a loss of independence, delivering assistance and monitoring on a 24/7 basis.

As part of the *SmartGait* project, and in parallel with my work, Dr. Alexandre Bordat [23] conducted a *CIFRE* thesis on *Radar* data acquisition aimed at identifying early signs of biomechanical decline and predicting falls. My thesis builds upon this work and is part of the doctoral programme for the digitisation of skills in the Île-de-France region, *Paris Région PhD*^{2}, co-funded by *CY Cergy Paris Université* and *BlueLinea*. It is being conducted at the *ETIS* laboratory, within the *CELL* team, under the supervision of Olivier ROMAIN and Son VU, with co-supervision initially provided by Amine KHELIF (from 1st June 2022 to 15th November 2023) and, since December 2023, by Petr DOBIÁŠ, both members of the *CELL* team.

1.2 Development and Contributions

My thesis focuses on recognising daily activities using *Radar* data transformed into micro-Doppler *spectrograms*. To this end, I developed *learning* methods tailored to this type of data and identified relevant features to support activity analysis and *fall-risk prevention*. The work aims to provide caregivers with objective, automated tools, with a view to the future development of a *non-intrusive* embedded system capable of supporting individuals in their daily lives.

In this context, existing *clinical tests*, such as the Tinetti and *Timed Up and Go* (TUG) assessments, were considered as a source of inspiration for identifying quantitative indicators for the objective evaluation of gait and daily activities. They were examined alongside studies of human motor skills and the *decomposition* of movements from a *clinical* perspective. This provided a deeper understanding of mobility loss, helped identify quantitative indicators, and facilitated communication with healthcare teams regarding *clinical* testing protocols and *scoring systems*.

Dr. Alexandre Bordat's work [23] led to the development of a *Radar* data pre-processing architecture at the *ETIS* laboratory. Using a *Frequency Modulated Continuous Wave* (FMCW) radar, it converts raw signals into usable *spectrograms*, reduces noise, and optimises processing time, while providing sufficient resolution to capture detailed movement *dynamics*, as illustrated in Figure 1.2.

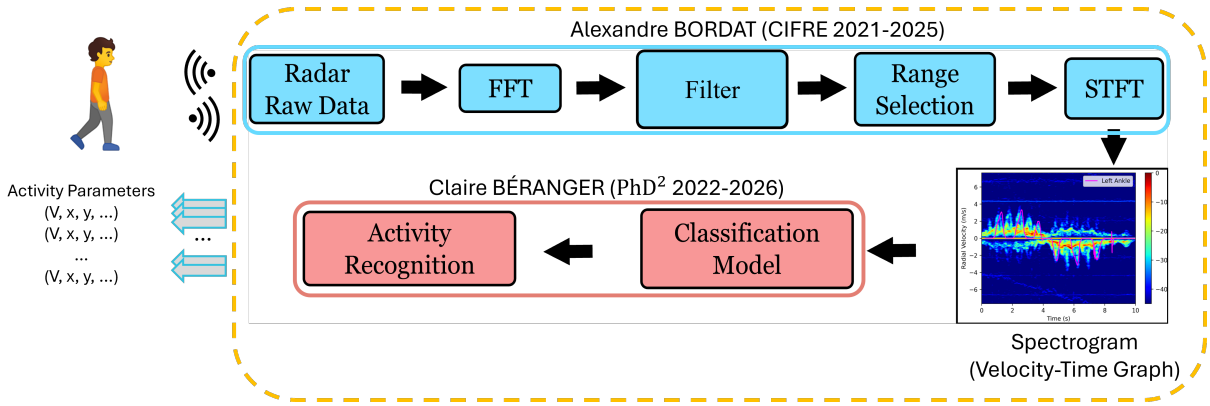

Figure 1.2 – Pre-Processing Chain (in blue) and Classification for Activity Recognition (in red)

Building on this foundation, the first months of my thesis were devoted to a literature review on *activity recognition* using *radar* data. This review aimed to identify commonly used *datasets* and *classification* models, and to analyse radar technologies applied to motion recognition within the project. It revealed that most studies still rely on accelerometers, gyroscopes, or image-based systems [76, 169, 174], while radar-based research remains limited. Moreover, available *datasets* are often simulated or derived from experimental conditions differing from the real-life situations targeted by the project.

We established an initial processing chain for recognising *daily* activities from *radar* data, as shown in Figure 1.2. A *learning* model known for its simplicity and efficiency was selected, also meeting the requirements for embedded implementation. Experiments on both open-source *datasets* and laboratory recordings validated the functionality of the chain and produced satisfactory results for activity recognition. This work led to a paper submitted to, accepted by, and presented at the *Digital System Design* (DSD) conference in September 2023 (Albania) [31], and also presented at the *IASIS and Robotics Inter-GdR Day* on 4th July 2024 in Paris, France.

Initial results, with accuracies ranging from 85% to 90%, highlighted issues with data quality and pre-processing. *Resizing* input images into small thumbnails may have altered spatial information, and some data were *noisy* or *erroneous*. To address this, we imple-

mented a method combining *noise reduction* and *dynamic cropping* to extract only the relevant regions of the *radar* signatures.

Subsequently, focus shifted to the *geometric* characteristics of the *radar* signature. An extraction programme computed descriptors such as *Area*, *Perimeter*, and *Orientation*, which were then used to train a model capable of effectively distinguishing different activities. This approach provided clear class separation while remaining suitable for moderately sized *datasets*, and led to two publications, both accepted and presented in 2025:

- a poster presentation at the *GRETSI* conference, held in Strasbourg in August 2025 [30];
- an oral presentation at the *DSD* conference in Salerno, in September 2025 [29].

The development of the initial *classification* models prompted the use of an ensemble of classifiers to improve prediction reliability. Texture-based parameters of *radar* signatures were explored, inspired by traditional image analysis, and several *learning* approaches were evaluated with broadly comparable performance.

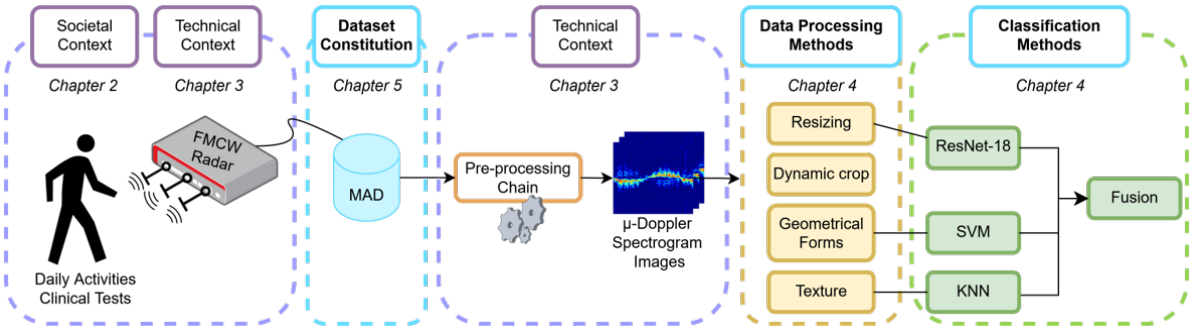

Figure 1.3 – Overview of the Full Thesis Chain

Ultimately, three independent models based on *geometric form*, *texture*, and the complete *spectrogram* were combined, and a *majority-voting fusion* method was implemented to produce a more robust final decision. These contributions are summarised in Figure 1.3, with elements in *blue* representing our work and those in *purple* representing components from the *state of the art*.

During this work, the *limitations* of the *dataset* in terms of both *quality* and *volume* were confirmed. To address these issues, a proposal was submitted to the *Ethics Committee* of *CY Cergy Paris Université* with the aim of creating a *dataset* specific to the project. This *new dataset* enables the evaluation of model robustness and the validation of a complete, functional chain for activity recognition based on *radar* data.

The work undertaken as part of my thesis continues the research initiated by the *SmartGait* project and has opened the way for further studies. Two additional theses have been launched following those of Alexandre Bordat and myself: one by Omar Ghoneim, focusing on *optimisation* and embedded processing, and another by Antoine Junior Tsagmo Denkeng, concentrating on the study of *stabilometry*. An engineer, Anouar Bengarai, also joined the team to *support* the practical *implementation* and *validation* of the complete experimental chain.

1.3 Manuscript Structure

This manuscript is organised into four chapters, as summarised in Figure 1.3, and finalised by a conclusion and perspectives.

- Chapter 2 presents the *clinical* context relevant to this thesis. It begins by describing the causes and consequences of *falls* associated with *ageing*, before discussing the role of medical innovations and support solutions for caregivers. Finally, it addresses the observation of human activity in the context of patient supervision, as well as the institutional perspectives related to such monitoring.
- Chapter 3 is devoted to the state of the art in *human activity recognition* technologies. It begins by presenting general approaches in the field, followed by a discussion of the main *learning* methods employed. Finally, it reviews the *datasets* and studies most relevant to our context, with the aim of selecting the *datasets* and *learning* models best suited to our research.
- Chapter 4 describes our initial experiments, which validated both the *dataset* and the selected *learning* methods. It also presents various approaches to *feature extraction*, including *geometric* forms and *texture*, as well as techniques for optimising the *classification* models. The chapter concludes with the introduction of a method for combining the developed models using *majority voting*.
- Chapter 5 presents the experimental *platform* and the various *sensors* available in the laboratory. It then describes the *protocol* implemented for the creation of our own *dataset*, named *Multimodal Activity Dataset* (MAD), before analysing its content and comparing the results with those reported in the *state of the art*.
- Chapter 6 provides a comprehensive *summary* of the work undertaken and highlights the prospects and *future* directions opened up by this research.

Chapter 2

Clinical Context and Medical Problem Statement

The *global demographic* context is characterised by an increasingly *ageing* population. As a result, *pathologies* are emerging, and long-term care demands on *clinical services* are rising. Despite technological advances in the medical field, there are still insufficient staff and facilities to provide adequate *support* for this growing population. Nursing homes, residential care facilities, and *teleassistance* services have been established to address the needs of patients who do not require hospitalisation or hospice care. These services are relatively recent, with most having been developed only in the 20th century [43, 52]. Consequently, numerous technological *innovations* and proposals have emerged to meet these evolving needs.

This chapter first focuses on understanding *ageing* and the *pathologies* associated with this phenomenon. There are various *stakeholders* who support elderly people in maintaining a good quality of life. These include clinical professionals as well as carers, such as family members or friends. Plans and programmes have been implemented to help *reduce* the *risk of falls*, which are one of the leading causes of morbidity in older age. *Clinical tests* have been developed to detect early warning signs, but they are not sufficient on their own. Different levels of risk are therefore identified and used to provide the most appropriate support for maintaining an individual's independence.

Faced with a shortage of staff and limited time in clinical services, technological solutions have been developed to provide support. These include initiatives at both national and regional levels, as well as within relevant services. Other solutions are still being investigated to enhance the services offered and, above all, to enable continuous monitoring of patients. To this end, an overview of the evolution of services and the development of *teleassistance* is provided, followed by a discussion of the various types of carers and the systems currently in place.

Finally, a description of *gait analysis* is provided, beginning with a clinical perspective

and progressing to human and software-based interpretation. The aim is to understand how *human activity* can be interpreted to identify indicators of *fall* risk. From *biomechanical* studies to *gait analysis*, observing specific *daily activities* is essential for extracting characteristic parameters. To support this objective, the *French government* has implemented an action plan for *fall prevention* since 2021 [21, 51, 122, 144, 66], introducing new measures, including improved communication between *medical services* and *teleassistance*.

As illustrated in Figure 2.1, the following diagram summarises the main themes introduced above and provides a visual overview of the chapter's structure, from the general context of ageing to clinical assessment and technological interventions.

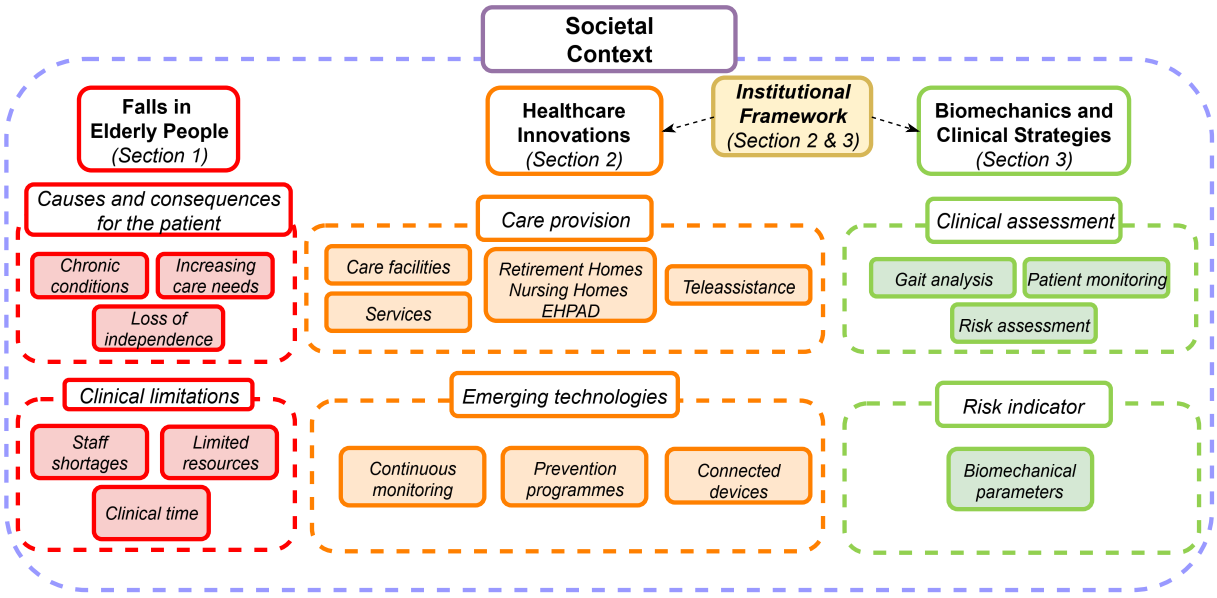

Figure 2.1 – Overview of Chapter 2

2.1 Falls in Elderly People: Risks and Interventions

Today, *healthcare services*, both in France and worldwide, are facing increasing pressures. As a result, patient care is becoming more complex, particularly in the context of demographic change. Both developed and developing countries are affected by this phenomenon, mainly due to the progressive *ageing* of the population.

This *demographic* shift is closely associated with a rise in chronic diseases, age-related disorders, and, in particular, *balance impairments*. These public health issues represent a growing challenge for healthcare systems.

In response, various initiatives have been implemented to assess the scale of the problem and to develop appropriate strategies. In France, several *action plans* [21, 51, 122, 144, 66] have been launched to raise public awareness and to prevent the risks associated with ageing. One key objective is to reduce hospital admissions and, whenever possible, to support individuals in living safely and independently in their own homes.

2.1.1 Ageing and Fall Risk: Consequences and Strategies

Ageing is a universal phenomenon that affects the entire population [116]. It can be defined in several ways. The most common is *chronological age*, which is measured by the passage of time and often linked to the onset of certain diseases, as the likelihood of developing health problems generally increases with age.

Another perspective is *biological age* [116, 190], which reflects the physiological changes associated with ageing. Unlike *chronological age*, it varies from one individual to another and is strongly influenced by bodily changes, habits, and lifestyle. This variability often makes medical cases particularly complex.

Finally, *psychological age* [116] also plays a significant role in the ageing process. It refers to an individual's mindset, reactions, and feelings in response to life situations. For example, a person who remains highly active may be considered psychologically younger than another person of the same *chronological age* who leads a *sedentary* lifestyle.

a) Ageing Population and Health Challenges

The ageing process often leads to the onset of pathologies of varying complexity, which manifest differently across individuals. These conditions may be visible or invisible, depending on their severity and the suddenness of their *consequences* for the patient. Ultimately, ageing is a phenomenon that impacts all generations.

When a person develops a *pathology*, they may require support in multiple ways. The individual is directly affected, experiencing fatigue, treatment side effects, or changes in daily habits [116, 190]. At the same time, neighbours, friends, and family members are also impacted, as they must adapt, *provide assistance*, and reorganise aspects of their own lives to ensure the well-being of the affected person. In addition, both patients and their caregivers may need to rely on healthcare services, further underlining the broad social and medical implications of ageing.

In developed countries, individuals are generally considered to show increasing signs of

ageing from the age of 65 onwards. The first manifestations may include *hearing or visual impairments*, osteoarthritis, diabetes, or lower back pain [190]. Mental health issues such as *depression* and *dementia* also become more prevalent with age. As the ageing process advances, individuals are increasingly prone to multifactorial health problems, many of which are classified as geriatric syndromes. These complex conditions encompass frailty, *falls*, confusion, and pressure ulcers, and represent a significant challenge for both patients and healthcare systems. To contextualise the discussion on ageing and associated health risks, it is important to visualise both the historical increase and the projected future growth of the elderly population in France.

The demographic context reinforces this trend. In France, the number of people aged 61 and over is expected to reach 18.3 million in 2025, representing 27% of the total population. By contrast, individuals aged 18–60 and 0–17 are projected to represent 53% and 20% of the population, respectively, as summarised in Table 2.1.

Table 2.1 – Distribution of the French Population by Age Group in 2010 (Millions and Percentage) [79]

Age Groups	**Millions**	**Pct.**
61+	18.3	27%
18-60	36.2	53%
0-17	14.1	20%
Total	**68.6**	**100%**

This evolution is illustrated in Figure 2.2, which shows the population pyramid for 2025 with a superimposed outline of 2010 for comparison. The figure highlights both the growth and the ageing of the population, particularly the increasing proportion of older adults.

All of these emerging conditions pose varying degrees of risk to an individual's health and quality of life. Each requires specific monitoring and may involve different medical services. While some conditions, such as hypertension or diabetes, are relatively straightforward to manage and follow up, others, such as *falls*, remain highly unpredictable. This *unpredictability* makes falls one of the conditions most strongly associated with morbidity among older adults, thereby representing a major concern for healthcare systems [117, 167].

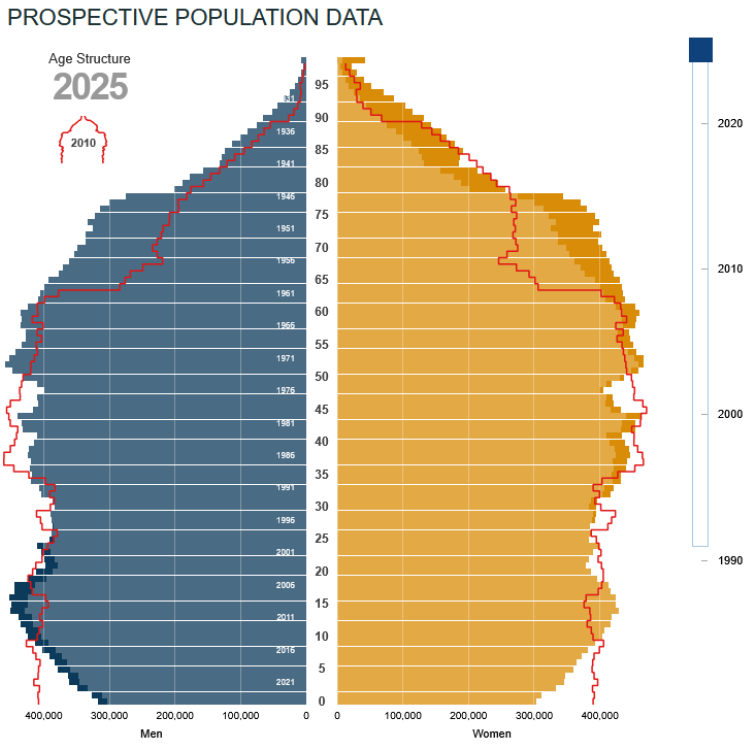

Figure 2.2 – Population Pyramid of France Projected for 2025 with 2010 for Comparison [79], Age-Sex Pyramid on *1st January 2025*

b) Epidemiology and Consequences of Falls

Falls are recognised as one of the leading health risks associated with *ageing* and constitute the second most common cause of accidental injury worldwide [189]. In the French context, falls account for approximately *10,000 deaths* and more than *136,000 hospital* admissions each year [117, 179, 177, 125]. They affect nearly 20% of individuals aged 65 and over, underscoring their major impact on public health and the pressing need for effective *prevention strategies*.

There are several types of falls and categories of fallers [33, 34]. A *fall* is commonly defined as an event in which a person unintentionally comes to rest on the ground, floor, or any surface located at a lower level than their original position [189].

Drawing on this definition and the classifications proposed in the literature [33, 34], several distinct situations can be identified:

- **Fall without witnesses:** a situation in which the patient is unable to describe the circumstances of the event, although there is clear evidence that a *fall* has

occurred;

- **Fall risk:** the probability that a person will experience a *fall*, usually determined through an objective assessment of individual risk factors;
- **Repeated falls:** defined as more than two falls within a six-month period. Individuals in this category are at greater risk of further incidents. In some cases, recurrent falls may be linked to underlying medical conditions or specific behaviours. Physicians, therefore, consider the consistency of contributing *factors* rather than simply the similarity of circumstances;
- **Post-fall syndrome or fear of falling:** a cluster of *psychological* disorders observed after a *fall*, including *anxiety*, *loss of confidence*, and *fear of recurrence*. These symptoms may lead to balance disturbances, reduced physical activity, and *decreased autonomy*, even in the absence of neurological deficits, mechanical dysfunction, or traumatic complications;
- **Almost fall or Near fall:** a situation in which the individual is about to *fall* but is prevented from doing so, for example, by being assisted to a chair or supported to the floor by another person.

According to the "*Global recommendation*" [21, 144] and the "*French national plan*" [51, 122, 66] for the prevention of falls among older adults, these accidents cost the healthcare system approximately €2 billion per year [167]. *Falls* can have numerous *consequences* for the individual. Patients may sustain *serious injuries or fractures*, often requiring hospitalisation. In addition, the *psychological* impact of a fall, such as fear of recurrence, can lead to a *loss of confidence* and *autonomy*, further compromising the independence of the individual.

Fear of falling can restrict movement and reduce physical activity. This restriction may contribute to *social isolation* and a *deterioration in the general state of health* of the individual [116]. Reduced physical activity also increases the risk of *undernutrition*, further aggravating the situation.

c) Post-Fall Decline Process

As outlined in article [116], a *self-perpetuating cycle of decline* emerges. When this decline becomes irreversible, the decision is often made to place the individual in a specialised institution, typically a residential care facility for dependent older adults, referred

to in France as an *EHPAD* (Établissement d'Hébergement pour Personnes Âgées Dépendantes). This choice is made to ensure access to medical care and a safe living environment.

In France, the *level of autonomy* of elderly individuals is often assessed using the *Iso-Resource Groups* (GIR) scale [156, 121, 120]. This classification defines six levels of dependence, from *complete dependence* to *full autonomy*, and is used to determine the type and intensity of support required for *daily activities*. *GIR 1* corresponds to the *highest level* of dependence, with individuals requiring constant assistance, whereas *GIR 6* represents the *lowest level* of dependence. Table 2.2 summarises the 6 GIR levels and the assistance they entail.

Table 2.2 – Iso-Resource Groups (GIR): 6 Levels of Dependence and Required Assistance, Inspired by [156]

GIR Level	**Level of Dependance**	**Assistance Required**
GIR 1	*Severely* impaired *physical* and *cognitive* abilities. Individuals are generally bedridden	*Continuous* presence and *daily* assistance
GIR 2	*Preserved* motor *abilities* but *partially* impaired *cognitive* abilities	*Assistance* with routine *daily activities*
GIR 3	*Partially* impaired *cognitive* and *motor* abilities	Assistance with *personal care* several times per day
GIR 4	Able to move around at home, but *difficulties* with certain *daily* tasks	*Help* with dressing and hygiene, getting up and going to bed
GIR 5	*Preserved* cognitive and motor abilities	*Occasional assistance* with hygiene and *daily activities*
GIR 6	*Fully autonomous* individuals	*Occasional assistance* with household tasks

Organisations such as *National Institute of Statistics and Economic Studies* (INSEE), the *World Health Organisation* (WHO), and *Santé publique France* agree that the risk of falling increases significantly from the age of 65 [189, 117, 78]. Both the physical and psychological consequences of falls are often severe.

The elderly are more vulnerable to accidents. Their bodies become more fragile, their reflexes decline, and injuries are often more severe. Falls can cause lasting, sometimes irreversible, trauma.

Falls have multiple *causes*. Some are related to *environmental factors* [63], such as slippery floors, poor lighting, or poorly adapted living spaces. Others are linked to the

individual's health, including *balance disorders*, *impaired vision*, *muscle weakness*, or the side effects of certain *medications*.

d) From Risk Analysis to Prevention

It is therefore essential to identify these risk factors in order to prevent falls. *Prevention* strategies include reorganising the living environment, attending regular medical check-ups, and maintaining a level of physical activity appropriate to the individual's age and condition.

Different *levels of risk* exist, which are described and defined in the "*Global recommendations*" [21]. To determine the *risk level* of an individual, careful *observation* is required to identify potential vulnerabilities. In the articles [15, 177, 125], the main risk *factors* are:

- **Sociodemographic**: age, sex, social isolation and prior falls;
- **Intrinsic**: balance, presence of chronic or acute diseases;
- **Behavioural**: daily activities, fear of falling, medication use, consumption of alcohol or other substances, level of physical activity, and nutrition;
- **Environmental**: potential hazards within and outside the home.

People who experience falls often require care in specialised clinical departments. This creates an additional workload for nursing staff. However, these services currently suffer from a *lack of time and staff*, which makes it difficult to manage these patients quickly and appropriately. Similar challenges are also observed in institutions such as residential care facilities for dependent elderly, like *EHPAD*.

According to some studies [34], 20% of elderly people who *fall* are responsible for 80% of the falls recorded in a single day. This indicates a very high relapse rate. These figures underline the importance of targeted measures for these so-called "*high-risk*" patients.

To reduce the number of accidents, *prevention campaigns* have been introduced in France [51]. Their aim is to raise public awareness of the risks associated with falling. These campaigns are being conducted both in clinical settings and in the workplace. Accidents in the workplace and certain living conditions can lead to pathologies that later increase the risk of falling.

The main aim of these campaigns is to prevent risk from an early age. Improved information allows individuals to take action in advance to *limit long-term consequences*. This involves adopting appropriate daily habits, maintaining a healthier lifestyle, and

attending regular medical check-ups, since workplace accidents or certain living conditions can lead to health problems that later increase the risk of falling.

More recently, the *French government* established the "*Anti-Fall*" Plan [21]. This plan aims to gain a better understanding of clinical needs, to take stock of individuals who fall, and to analyse the causes of their accidents. It also seeks to provide the tools and equipment necessary to improve care.

An increasing number of countries, including developing ones, are exploring sustainable solutions. The aim is to enable elderly individuals to remain independent for as long as possible, while ensuring access to appropriate *in-home services*.

2.1.2 Structure for Fall Prevention and Management

The range of conditions observed throughout the ageing process requires the involvement of multiple people to provide *optimal support* for the individual [33, 34]. Each of these individuals has a distinct *role* and set of *responsibilities*.

To the extent of their abilities, the person receiving care is the first to take an active role in managing their health [33]. The *patient* participates in *decision-making* and is actively involved in the care they require.

a) Definition of Roles in Fall Prevention and Management

As part of *fall prevention*, the patient and their relatives are encouraged to share information with healthcare professionals about the patient's lifestyle habits. The patient also expresses their preferences and goals regarding their difficulties, thereby contributing to the establishment of preventive measures and care objectives related to *fall risk* [33].

The *patient* and their family members play an essential role in accepting and preparing for the monitoring of their health condition [34]. They subsequently meet with various healthcare professionals and other key contributors involved in maintaining independence [33, 34].

Multiple healthcare professionals collaborate to support the patient's safety and independence. Table 2.3 summarises the main *responsibilities* of each professional, organised roughly according to their proximity and frequency of contact with the patient. The patient remains at the centre of these interventions, actively participating in decisions and expressing preferences and goals.

As shown in Table 2.3, the *auxiliary nurse* and the *nurse* maintain the closest and

most frequent contact with the *patient*, providing daily care and monitoring. Other professionals, including the *occupational therapist*, *nutritionist*, *specialised nurse practitioner*, and *doctor*, intervene as needed to assess specific risks and implement targeted preventive strategies. The *special education professional* contributes additional support, particularly regarding autonomy and environmental adaptations.

Table 2.3 – Roles of Healthcare Professionals in Fall Prevention, Ordered by Level of Patient Contact

Profession	Main responsibilities	Specific role in fall prevention
Auxiliary nurse	Provides *daily care*, supports the *nursing plan*, observes and reports changes	Identifies risk *factors*, informs patients/families, supports *post-fall* follow-up, applies *prevention strategies*, reports warning signs
Nurse	Assesses health, develops and monitors care plans, coordinates the team, ensures *continuity* of care	Educates patients/families, screens for risk *factors*, conducts *post-fall* analyses, adjusts strategies, liaises with other professionals
Specialised nurse practitioner	Experienced nurse with advanced clinical skills	Performs targeted *fall* risk assessments, identifies risk factors, examines patients after falls
Occupational therapist	Evaluates functional abilities, develops interventions, *adapts environment* to maintain independence	Analyses *post-fall* events, teaches safe transfers, recommends equipment, collaborates on *prevention strategies*
Nutritionist	Assesses nutritional status, develops and monitors *dietary* interventions	Screens for *malnutrition*, supports strength and *balance*, monitors diet, educates on nutrition-related *fall prevention*
Doctor	Diagnoses and treats conditions, supervises care, coordinates *multidisciplinary* interventions	Identifies *medical factors* influencing risk, performs examinations, plans treatment, guides the team in *prevention*
Special education professional	Supports patients with *social*, *educational*, *psychological*, or *health challenges*.	Identifies risk *factors*, reports concerns, applies preventive measures for at-risk patients

Figure 2.3 illustrates the collaborative network of healthcare professionals involved in *fall prevention*. The *patient* is positioned at the center, with professionals who provide daily care (*nurse* and *auxiliary nurse*) placed nearby, while those with more specialised or supervisory roles (*specialised nurse practitioner*, *occupational therapist*, *nutritionist*, *doctor*, and *special education professional*) are positioned further away. Solid arrows indicate direct interactions with the *patient*, while dashed arrows represent *coordination* and *communication* between the professionals themselves. This layout highlights the integrated, team-based approach required to support the *patient*'s safety and independence.

Together, these professionals form a coordinated support system, ensuring that preventive measures are tailored to each patient's individual needs, and that interventions are adapted as their condition evolves.

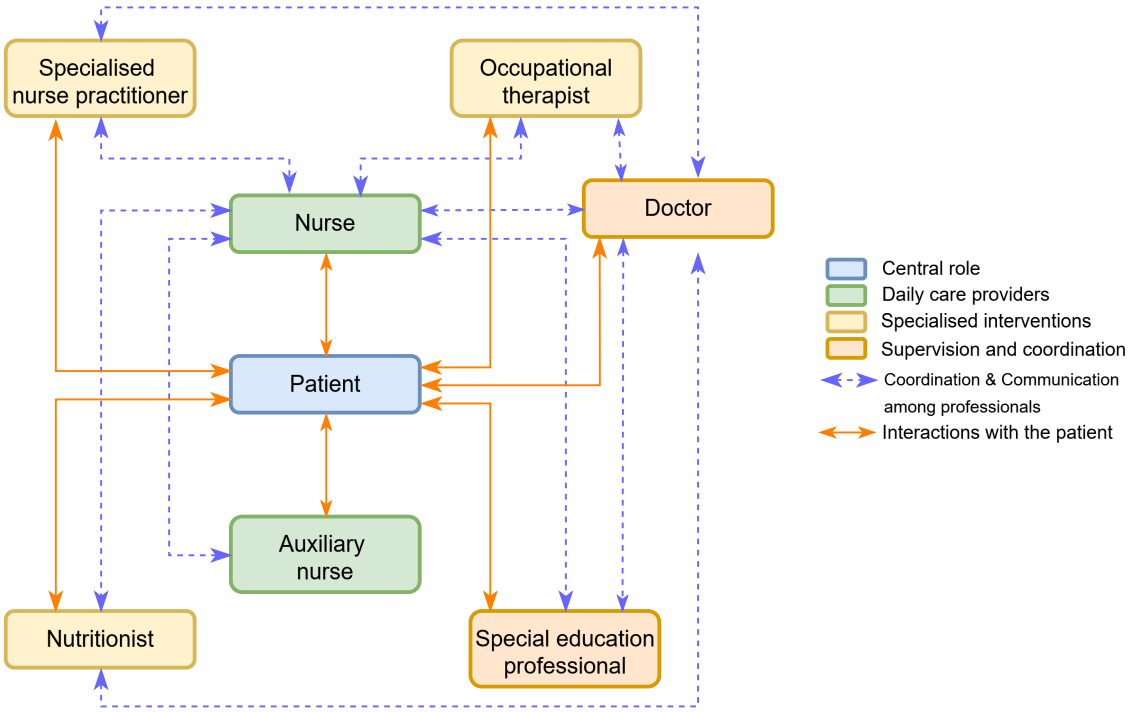

Figure 2.3 – Collaborative Network of Healthcare Professionals in Fall Prevention

b) National and Regional Strategies

Various plans and programmes [21, 51, 122, 144, 66] have been discussed in recent years. *Falls* account for a significant number of deaths among individuals over the age of 65, with more than 10,000 fatalities per year in France [117, 179, 125]. Governments, both collectively and individually, are considering the implementation of different measures to provide *optimal support* for care services [21, 51, 122, 144, 66].

In 2022, with the introduction of its *fall prevention* plan [51, 122], the *French Ministry of Health* aims to address needs more effectively and better organise the following five areas:

- **Early risk detection**: identify risks to take action as early as possible;
- **Home adaptation**: suggest modifications to the home;
- **Mobility support**: provide mobility and technical assistance;
- **Physical activity**: set up physical activities adapted to the patient;
- **Teleassistance**: use teleassistance devices.

The aim here is to provide new support tools for both *nurses* and *patients*. By implementing these new measures, the government hopes to reduce falls causing serious injury

or death [80] among the over the age of 65 by 20% [51].

Regions are working to implement solutions to achieve this objective at the local level [4] . *Prevention* assessments have been established, consisting of a *multidisciplinary* evaluation. *Patients* complete a questionnaire and undergo a clinical examination. They perform functional tests, and a home diagnosis is conducted, followed by workshops tailored to their specific needs. This approach allows for a comprehensive evaluation of the *patient*'s health, particularly in terms of *balance* and *nutrition*.

For example, in the *Auvergne-Rhône-Alpes* region [4], a *mobile showroom* has been set up to provide *information* and *communication* aimed at improving and raising awareness about the quality of life for patients who wish to remain independent in their own homes.

In *Annecy* [4], *prevention platforms* are offered with the objectives of raising awareness, identifying risks, assessing needs, *guiding patients*, and coordinating interventions. This is an integrated organisation, involving the mobilisation of the 15 centres, *Territorial Professional Health Communities*, and remote alarm systems. Individual and group prevention workshops are held throughout the region, along with training and awareness-raising activities.

A programme entitled "*Ageing Healthily at Home*" is also being offered, run by the Versants d'Aime *Multi-professional Health Centre* [4]. This programme targets the entire population aged over 65, focusing on identifying vulnerabilities and implementing preventive and corrective measures in a multi-professional manner. These regional approaches are summarised in Figure 2.4, which illustrates the main stages and actors involved in *fall prevention* programmes.

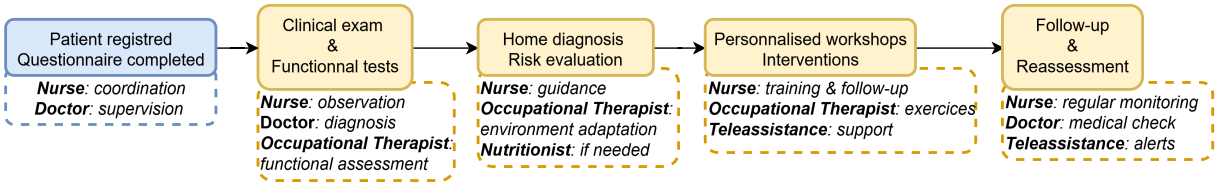

Figure 2.4 – Examples of Regional Fall Prevention Pathways and Professional Involvement. The figure illustrates the main stages implemented in certain French regions, from patient registration to follow-up, and the associated roles of healthcare professionals such as nurses, doctors, occupational therapists, and teleassistance providers. Based on the descriptions provided in [4, 119].

Among all these initiatives, *nurses* play a crucial role [119], serving as *communicators*, *observers*, *intermediaries* between *doctors* and *patients*, trainers, and implementers of care. *Nurses* help detect early warning signs, such as changes in *vision*, a *sedentary* lifestyle, or

malnutrition. They can provide guidance on home adaptations, including the installation of grab bars, the use of adapted furniture, and the rearrangement of living spaces. They also assist *patients* in using *teleassistance* services and coordinate information with *doctors*, *occupational therapists*, and other healthcare professionals.

c) Fall Risk Assessment Tests

Pathologies associated with, or unrelated to, falls are not new challenges. Over the years, clinical practice has evolved alongside these societal issues. Various training programs, protocols, and tests have been implemented to monitor early warning signs [137, 60, 28, 141, 178, 18, 82, 22]. However, a *lack of staff and time* often prevents proper monitoring, resulting in the loss of crucial information regarding the *patient*'s health status. Regardless of the condition, timely intervention, or even preventive measures before the development of a new stage or complication, reduces the risk of relapse or deterioration. Conversely, if treatment is delayed, the patient's condition may at best be stabilised or slightly improved, but full recovery to a stable state is no longer achievable.

The five *clinical tests* most commonly used for *fall* risk [28, 141, 178, 18, 82, 22] are:

- **Morse Fall Scale** [82]: covers six assessment factors
  * History of falls;
  * Presence of comorbidities;
  * Walk aids;
  * Presence of an infusion;
  * Patient's gait;
  * Mental health status.

  This test is based primarily on history and various factors, with a greater emphasis on medical, cognitive and societal aspects.
- **Tinetti Performance Oriented Mobility Assessment (POMA) test** [178]: consists of several types of activities to be performed by the patient under the supervision of a clinician. It is divided into two parts:
  * The first part assesses *balance* using a chair without armrests;
  * The second part assesses "*Walk*" on a course of at least 3 metres.

  The overall score is used to determine the risk of falling if it is below the threshold.

- **Berg Balance Scale (BBS)** [18]: is used to analyse a person's *balance* using 14 types of movements performed in everyday life within a set time, as illustrated in Figure 2.5. Clinicians use equipment to recreate these situations, such as:
  * Stopwatch;
  * Tape measure;
  * Chair with and without armrests;
  * 19.5 cm high step stool.



Figure 2.5 – Illustration of the *BBS* from [85], (A) Sit to stand, stand unsupported; (B) Stand to sit, sit unsupported; (C) Transfers between surfaces; (D) Stand with eyes closed, the *Romberg* test with eyes open; (E) Reach forward while standing; (F) Pick up an object from the floor; (G) Turn and look behind; (H) Perform a 360º turn to the right and left; (I) Touch toes to a stool; (J) Perform the sharpened *Romberg* test; and (K) Stand on a single leg.

- **Timed Up and Go (TUG) test** [141]: is simple to set up and widely used across *clinical* settings. It involves a sequence of activities performed under the supervision

of a healthcare professional, such as a *nurse*, *physiotherapist*, or *clinician*. The *patient* is asked to *stand up* from a chair with armrests, *walk* a *distance* of 3 metres, *turn around*, *walk back* to the chair, and *sit down*, as detailed in Figure 2.6. During the test, the *clinician* observes the *patient*'s performance on each sub-activity and records the total time taken. A completion time of less than *20 seconds* generally indicates a *low risk of falling*. Although variations exist depending on the institution, this is the most commonly used version of the test.

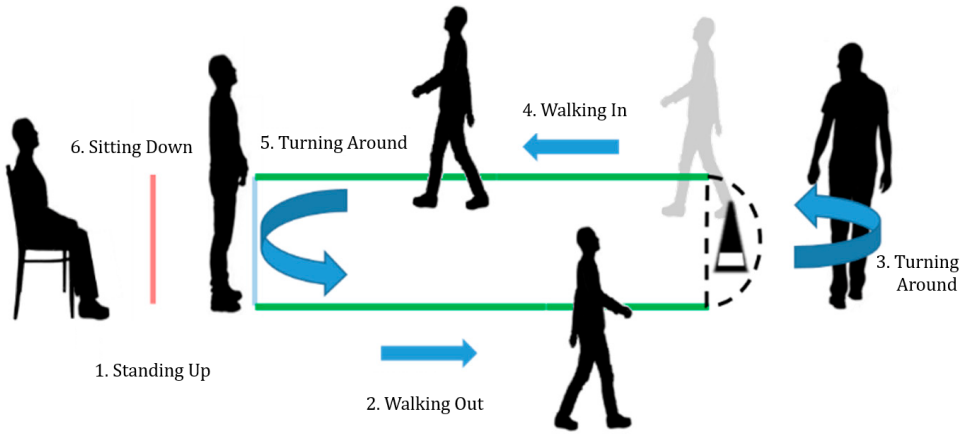

Figure 2.6 – Illustration of the Timed Up and Go (TUG) Test from [73]

- **Single-leg stance test** [22]: is a *static balance* assessment conducted under the supervision of a healthcare professional. The *patient* is asked to *stand on the leg* of their choice for as long as possible. An inability to maintain the stance for at least *5 seconds* indicates a *high risk of falling*, while a duration of *30 seconds* or more suggests a *lower risk*.

These tests together provide a comprehensive framework for assessing *fall risk* among older adults. Table 2.4 summarises their main features and evaluation *criteria*.

The purpose of these tests is to identify signs of fall risk or potential relapse based on the *patient*'s history. All of these assessments rely on observation by a healthcare professional, which means that interpretations may vary between *clinicians*, as conclusions are based on both *objective* and *subjective criteria*. Current research is exploring solutions to support *patients* even before hospital admission.

Certain aspects of these tests can be automated while still allowing the *clinician* to make the final judgment. Such automation can help compensate for the persistent *limitations of time and staffing* in healthcare departments.

Table 2.4 – Summary of the Main Clinical Tests Used for Fall Risk Assessment

Test	**Main focus**	**Key features / criteria**
Morse Fall Scale [82]	Medical and cognitive factors	History of falls, comorbidities, walk aids, infusion, gait, mental status
Tinetti (POMA) [178]	Balance and gait	Divided into two parts: static balance and walk over $3m$, risk determined by score threshold
Berg Balance Scale (BBS) [18]	Functional balance	14 everyday movements evaluated using tools such as stopwatch and tape measure
Timed Up and Go (TUG) [141]	Mobility and dynamic balance	Stand up, walk $3m$, turn, walk back, sit; total time $< 20s$ indicates low risk
Single-Leg Stance [22]	Static balance	Time maintaining single-leg stance; $< 5s$ = high risk, $\geq 30s$ = low risk

One question remains regarding the implementation of patient monitoring in *real-life* situations. Specifically, in cases where an individual is at risk of falling but is not under the supervision of a healthcare team, it is unclear how such risks are addressed. For those who are detected and monitored, it is important to consider how both the *patient* and their *family* are supported and reassured throughout the process.

2.1.3 Patient Monitoring

Patient care and health monitoring rely primarily on the *patient*. Individuals who take the initiative to learn about the risks associated with *ageing* are taking the first step toward understanding and engaging with the medical field. In most cases, the attending *physician* provides initial information on age-related preventive measures during a routine visit. The *patient* can then discuss these measures with their doctor and, if desired, with family and friends.

If a *patient* has already experienced a *fall*, it is important to discuss this with the *physician*, who must determine the underlying cause(s) in order to provide appropriate treatment. The ultimate goal is to prevent further falls and reduce the risk of hospitalisation.

a) Independence of Elderly People Living at Home in France

In France, the *government* emphasises the importance of maintaining overall health [63]. Engaging in regular physical activity is recommended to preserve muscle tone and maintain a basic level of flexibility. Individuals should also pay attention to their diet to prevent *malnutrition* and ensure adequate nutritional status. Regular eye and hearing examinations are necessary to avoid discomfort and maintain confidence in one's abilities. Finally, to reduce the risk of *imbalance*, it is recommended that individuals wear appropriate footwear and take proper care of their feet.

For individuals who continue to live independently at home, support can be provided to optimise and adapt their living environment according to their health needs. An *occupational therapist* or *specialist physician* may be consulted to advise on home modifications [63]. Assessments of the home's practicality, particularly regarding *fall risk*, focus on the following aspects:

- **Lighting**: Ensure good lighting to avoid dark areas or sudden changes in brightness. For example, installing a lighted path for night-time trips to the toilet may be recommended;
- **Assistive equipment**: Add suitable aids such as grab bars, shower seats or handrails to enhance safety;
- **Home modifications**: Carry out work to facilitate safe movement. This may include adding non-slip strips under rugs, installing non-slip stair nosings, replacing slippery flooring, or adjusting the height of shelves;
- **Bathroom and toilet adaptations**: Modify wet areas, for example by installing a walk-in shower or a raised toilet to improve accessibility and reduce the risk of falls.

Technical support may also be recommended to the individual. This can include devices such as a walking stick, a walking frame, a wheelchair, an access ramp, or even an electric scooter to facilitate mobility both at home and outdoors.

Depending on their health status, individuals may request *home-based* services, such as visits from *nurses* or *personal care assistants*. In addition to, or as an alternative to, these services, they may also arrange for *teleassistance* support.

b) Assessment and Monitoring in a Hospital Setting

In a hospitalised setting, such as a nursing home or hospital, individuals remain at risk of falling even if they have not previously experienced a *fall*. As in the home environment, the surroundings can constitute a risk factor. Flooring, surfaces, and equipment in corridors and bedrooms can all affect a person's *balance*. Additionally, individuals who have recently left their home may experience disorientation due to the change in environment. Illness or recent surgery can also contribute to instability. Low staffing levels and time constraints may prevent staff from immediately noticing when a patient falls. Nonetheless, falls require prompt attention to minimise the risk of recurrence [124].

If staff witness a *fall*, they will attempt to break the *fall* using their body while ensuring that the *patient*'s head does not strike the floor or any other surface. They will remain with the *patient* and, if necessary, call for additional assistance [127]. If the *fall* is not witnessed, staff will seek information from the *patient* or any bystanders who observed the incident. They will check the *patient's breathing and blood pressure* and assess for injuries, including potential *fractures*. Depending on the patient's condition, staff may place them in a chair, on a backboard, or on a stretcher. Monitoring continues after the fall, with *staff* performing various checks, such as assessing alertness, *measuring blood pressure*, and *monitoring blood glucose levels*.

Patients may be asked to undergo assessments to identify and prevent potential risk factors for falls before they occur. For this reason, *physicians* are advised to ask individuals in *higher-risk* age groups whether they have experienced a *fall* in the previous year. Simply asking this question provides an opportunity to discuss preventive measures with the patient. The goal is to respect each individual's desire to maintain independence and make their own life choices, thereby supporting "*healthy ageing*" [69].

The procedure then involves assessing potential risk factors related to the individual's health, behaviour, and environment. These factors can be both *intrinsic* and *extrinsic* and may be detected through clinical examination and patient interviews. The number of medications and the presence of illnesses also influence the risk of falls, as they can impair sensory, cognitive, or motor functions. The risk of injury following previous falls is particularly high, especially when conditions affecting bone strength, vision, or balance are present. Improving the management of these health issues and reducing polypharmacy are central objectives of *multifactorial intervention programmes* aimed at *fall prevention* [69]. Other factors must also be considered, particularly those related to behaviour and the environment.

c) Level of Fall Risk Assessment

In addition to interviews and the analysis of various *factors* that may initially appear to influence a patient's risk of falling, specific *tests* may also be conducted. These *tests* involve interacting with the *patient* by asking them to perform certain tasks while observing their ability to do so. The attending *physician* can perform an initial assessment using simple *tests*, such as asking the *patient* to balance on one leg [69, 70]. They may also ask the patient to walk while simultaneously answering a question. If the *patient stops* to respond, this indicates a *higher risk* of falls.

If any of these initial screening tests indicate a potential risk, a *multifactorial* and *interdisciplinary* assessment must be carried out. This broader evaluation integrates the results of the preliminary tests and patient interviews, providing the foundation for developing targeted strategies to prevent falls and fractures. Following these observations and assessments, appropriate interventions can then be recommended to both the patient and their relatives.

Based on this comprehensive assessment, different levels of fall risk severity have been identified [70]. During the interview, if an individual has previously fallen and responds affirmatively to any of the additional questions, their *fall* is considered *serious*, and the *patient* is classified as at *high risk* of falling.

If this is not the case, further assessment may be conducted based on the individual's responses, particularly regarding their confidence in maintaining balance. Should a balance or gait disorder be detected, the patient is classified as at intermediate risk of falling. Finally, if *no such signs* are present, the *patient* is considered at *low risk*. Patient care can be summarised using the diagram in Figure 2.7.

For patients classified at *high risk*, the healthcare team provides aetiological treatment and secondary prevention [70]. The goal is to assess all risk factors and implement personalised interventions. This follow-up generally lasts between 30 and 90 days.

For patients at *intermediate risk*, the focus is on secondary prevention and modifying major risk factors. Personalised balance and muscle-strengthening exercises are recommended, and the patient receives education on fall prevention.

Patients classified at *low risk* receive primary prevention, which includes education on fall prevention and advice to maintain or initiate regular physical activity. For both *intermediate- and low-risk* patients, follow-up typically extends over a period of 12 months [70].

It should be noted that while some *tests*, such as the *TUG* test [141] and *BBS* [18], are widely used *internationally*, others, like the *Morse Fall Scale* [82] or certain versions

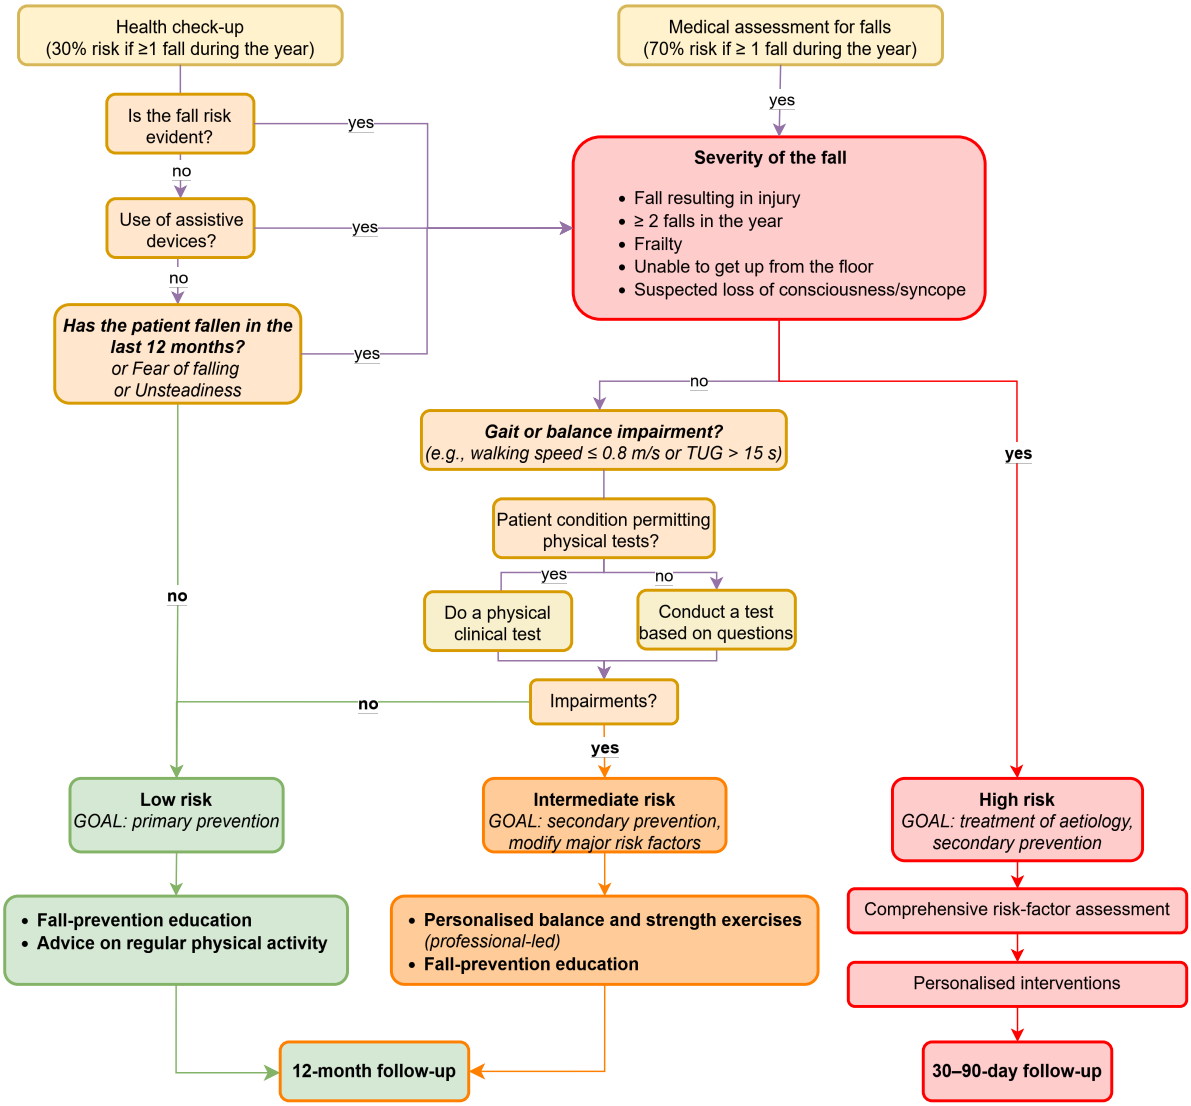

Figure 2.7 – Preliminary Fall Risk Assessment Algorithm, Inspired by [35, 70]

of the *Tinetti* test [178], may *vary* in use and scoring across countries. This variability underscores the ongoing *challenge of standardising fall risk assessment* globally and highlights the importance of implementing coordinated support measures to ensure consistent monitoring and preventive care.

In parallel with these clinical and medical procedures, countries are implementing support measures. In France, local authorities, regions, and territories are launching initiatives to facilitate access to services [4]. Monitoring requirements may be established at certain key ages or depending on whether the patient has previously undergone specific examinations.

For example, if a *patient* has already experienced a *fall*, a follow-up assessment is initiated to ensure a minimum level of service and to evaluate the consequences of the incident [4]. However, this assessment often occurs only one year later, leaving the patient without structured support during the intervening period.

Some regions have implemented initiatives such as *mobile local platforms* [4] to provide *assistance* and organise *workshops* for the broader population. Despite these efforts, communication regarding these measures and patients' acceptance of frequent health checks remains inconsistent.

2.2 Medical Challenge and Innovations

Despite these solutions and the proposed monitoring protocols, adequately tracking every patient or at-risk individual remains challenging. Some individuals may be identified late if they have not been made aware of the risk in time or do not adhere to their physician's recommendations. Additionally, implementing these measures often requires significant changes in the individual's *daily life*. Modifications such as home adaptations, relocation, or changes in routines can be sudden and disruptive.

Technological advances have made it possible to complement these interventions while allowing a more gradual transition to new habits, thereby reducing the potential disruption to the individual's life.

The advent of various *sensors* and connected devices has enabled their use in the *clinical* field, with the intention of providing more detailed and continuous monitoring. Advances in *signal processing*, including the *analysis of images* and sound, have significantly contributed to the assessment of pathologies and the monitoring of health status in identified conditions. Remote *assistance technologies* have also emerged, allowing devices to be installed directly in a patient's bedroom or home. These innovations offer potential solutions for mitigating the risk of falls.

2.2.1 Clinical Solutions and Emerging Approaches

Today, to limit the risk of falls, the medical profession recommends the use of assistive devices [119]. The most commonly recommended items include *walk sticks*, *walkers*, *grab bars*, *footstools*, and *mobility scooters*. These devices can be found on many websites specialising in this field [180, 46]. Some equipment can be installed by the user, while

other devices require professional installation by companies that provide such services. For example, *Prevenchute* [143] offers various solutions for adapting private homes.

Many devices are motorised and are becoming increasingly sophisticated by integrating various technologies that can be tailored to the patient's needs. Examples include motorised chairs, more practical shopping bags, and increasingly well-equipped walkers. All of these devices provide *support* and *reassurance* to *patients* and their families. However, they do not enable monitoring of the individual's health status.

a) Clinical Solutions and Recommendations for Fall Prevention

Alert systems have also been implemented. Initially introduced in clinical departments to allow patients to call for a healthcare professional, they have since been adopted by certain hospitals. Depending on the facility, some use multiple call buttons, while others have begun to combine detection technologies with the call function. Several types of alert systems exist. As shown in [126], in 2025, several devices with *fall detection* capabilities were introduced.

Some examples of such systems include *Medical Guardian*, *MobileHelp*, and *Bay Alarm Medical* [126]. These systems assist in *fall detection*, either when the patient presses a button or when abnormal movement or speed is detected. The underlying technologies often rely on *accelerometers* or *gyroscopes*. They frequently exhibit high false alarm rates is around 90% [107]. Generally, the devices are designed to be worn on the neck, belt, in a pocket, on the wrist, or in the hand, as illustrated in Figure 2.8.

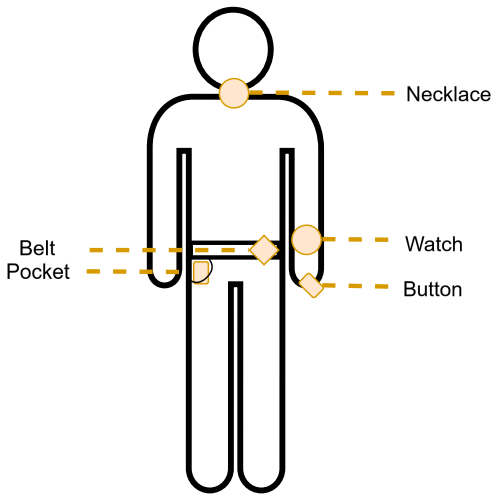

Figure 2.8 – Location of Sensors on the Body

The implementation of these solutions is costly and still lacks precision. "*Fall*" is a *multifactorial phenomenon*, as discussed in Section 2.1.1. This highlights a gap in medical knowledge, particularly in the availability of data on falls. Currently, detection relies solely on the observation of factors or indicators that are assumed to be associated with a fall. For instance, a sudden increase in the person's movement speed may indicate a fall. However, this alone is insufficient. Thresholds vary from person to person, and a fall can also occur gradually. A drop in *blood glucose levels* may also serve as an *indicator* [127], but it could be related to another health issue. All of these factors complicate analysis, whether performed by a machine or through direct observation by a clinician.

The solutions studied and available on the market tend to generate many *false alarms* rather than identifying only actual falls [107]. This leads to two main conclusions: first, these solutions are effective in drawing attention to a *patient* who may require *assistance*, allowing for *rapid intervention* if the alert is genuine. Second, if the alert is false, it consumes *staff time* and *resources* and may occupy space that could be used for another patient.

b) Additional Solutions to Support Clinical Practice

Local authorities, regions, and territories aim to provide more viable and relevant solutions for *clinical services* [4]. Research has been conducted using various *sensors* to evaluate the feasibility and obtain new *metrics* that could better support caregivers. *Image processing* has been studied for many years, and knowledge in this field is highly advanced. Research has particularly focused on the use of video cameras.

Cameras are indeed valuable tools for detecting activity in multiple domains, such as defence and security. However, these *sensors* are highly *sensitive* to lighting conditions, which poses challenges for monitoring falls that may occur at night [26]. Furthermore, *privacy considerations* must be taken into account. In a medical setting, having access to images of patients in all the rooms they use, including toilets and bathrooms, is not a viable solution.

Other research proposes the use of connected insoles [170]. The concept involves *embedding sensors* in the insoles worn by patients to study their *balance*, which allows for calculations such as determining the location of the *centre of gravity*. However, this approach is costly and requires the insoles to be worn at all times, including when lying down. Current studies aim to develop solutions that are inexpensive, easy to implement, respect privacy, and are minimally restrictive for the patient. Wearing a sensor may be

perceived as intrusive or inappropriate, and the visibility or tactile presence of the device may influence the patient's behaviour in daily life.

c) Innovative Clinical Monitoring Solutions

Other studies focus on the use of *Inertial Measurement Unit* (IMUs) [187, 185]. These devices consist of a small unit worn by the patient while performing a test. The issue is that such technologies still impose significant constraints on the patient. For this reason, the most commonly used solutions remain cameras, necklaces, and pendants. Despite their *high false alarm rates* and associated *limitations*, they are still the easiest to use and the most familiar to the general population.

A device capable of detecting the movements involved in performing an activity is required, while remaining insensitive to changes in lighting and non-intrusive for the user. This is why *radar* technology is increasingly being considered as a solution to this challenge [74, 160, 58, 50].

The *radar* approach does not require the patient to wear any equipment, ensuring that *no intrusive* procedures are involved [50]. The *radar* emits waves and captures their echoes, which are then processed into speed-time images known as *spectrograms*. No visual image of the patient is recorded, only the movement itself is captured. These micro-Doppler *spectrograms* are subsequently analysed to interpret the performed activity and detect potential risks.

2.2.2 Evolution of Healthcare and Support Systems in France

Originally, there were no dedicated retirement homes. Elderly people lived with their families and relied on them for their livelihood [43]. Since ancient times, some civilisations have had hospices that took in the most destitute or those without family support. Later, the Church and other religious institutions established care services for the sick, the destitute, and the elderly. Their role was primarily to provide shelter and food. These establishments were often rudimentary and functioned more as a form of assistance than as a means of supporting the ageing process.

a) History of Healthcare and Medical Accommodation Institutions

In France, assistance for the elderly began to take shape in the 17th century, when Louis XIV established the "*Hospices Généraux*" (general hospices). However, these insti-

tutions remained largely asylums, where the quality of life was precarious. With industrialisation in the 19th century, family ties began to weaken, and the issue of housing for the elderly became increasingly urgent [43]. The state started to intervene by creating specific shelters for the elderly. Over time, these institutions gradually evolved into organised living spaces that provided better care. In the 20th century, living conditions continued to improve, particularly with advances in *medicine*. As life expectancy increased, many elderly individuals were no longer able to remain in their own homes, making the question of where to house them ever more pressing.

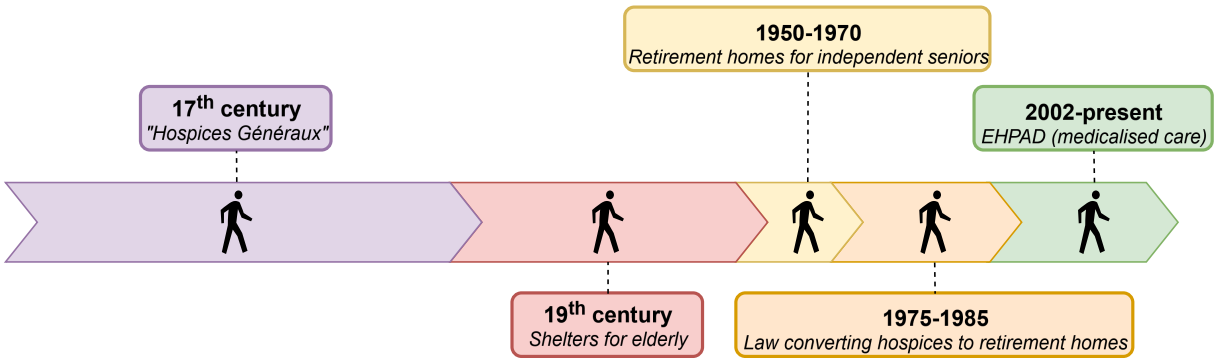

Figure 2.9 – Historical Evolution of Elderly Care Institutions in France

The first *retirement homes* emerged after the *Second World War*, during the 1950s and 1960s. They differed from traditional hospices and were created to provide accommodation for independent senior citizens, sometimes offering limited services such as meals and *assistance* with *daily activities* [43]. It was only in the 1970s and 1980s, with growing awareness of the ageing population and the increasing number of elderly people losing their independence, that public authorities recognised the need to adapt these structures. In 1975, a law was introduced requiring public authorities to convert hospices into retirement homes, whether or not they provided *medical care* [52]. A ten-year period was granted to complete this transition. Some establishments employed *nursing staff* and began offering medical care.

In the 1990s, the *medicalisation* of retirement homes became a priority. In 1997, a personalised autonomy allowance was introduced, amending the 1975 law [43]. Flat rates, based on each resident's needs and condition, were proposed to finance their care. In 2002, residential care homes for dependent elderly people (*EHPAD*) were established. These institutions differ from traditional retirement homes in several ways:

- **Medical care**: each establishment must sign an agreement guaranteeing medical

presence and appropriate care for residents;

- **Multidisciplinary team**: presence of carers, nursing assistants, nurses, coordinating doctors and sometimes psychologists or occupational therapists;
- **Accommodation adapted to dependency**: rooms and common areas designed to ensure the safety and comfort of people with reduced autonomy;
- **Regulated framework**: *EHPAD* rates and services are regulated, and establishments are subject to regular inspections to guarantee the quality of care and support.

To visualise this historical evolution, Figure 2.9 presents a timeline of the main types of elderly *care institutions in France* from the *17th century* to the establishment of *EHPADs*. Table 2.5 summarises the main characteristics of these institutions, highlighting their target populations, services provided, and specific features.

Table 2.5 – Comparison of Elderly Care Institutions in France

Structure	**Period**	**Target Population**	**Services**	**Particularities**
Hospices Généraux	17th century	Dependent elderly	Basic shelter	Low quality of life
Shelters	19th century	Elderly needing support	Shelter	Limited care
Retirement Homes	1950-1970	Independent seniors	Shelter, meals	Limited care
Post-1975 Retirement Homes	1975-1985	Seniors losing autonomy	Shelter, some care	Conversion of hospices, regulation
EHPAD	2002-present	Dependent elderly	Medical care, multidisciplinary team, adapted rooms	Regulated rates, inspections

The advent of *nursing homes* has enabled a more structured approach to the care of dependent elderly people by *professionalising* and *standardising* services [53]. However, it has also introduced new challenges, such as the high cost of accessing these facilities and the shortage of *nursing staff*. Today, the main issue is the need to rethink the support provided to elderly people in a society that increasingly favours *home-based* care.

b) From TéléAlarm to Teleassistance

Alongside this overhaul of the elderly care system, the *TéléAlarm* concept was launched in 1975 [1]. A French company introduced the first solution for providing assistance to elderly people in their own homes. The *TéléAlarm* device consisted of a cord with a push-button alarm connected to the user's telephone line [44, 1]. In case of need, the person could press the button, which would alert a call centre and print a ticket with their name and telephone number. The beneficiary was then contacted by an advisor. This type of device was intended for individuals who were unable to live independently without short-term human assistance and was primarily used to notify call centres of potentially risky situations.

The implementation of this solution highlighted the need to develop systems that would allow people with reduced independence to remain in their own homes [1]. The principle of calling the beneficiary back proved too restrictive and caused delays in providing care. In 1982, *teleassistance* was introduced. This new solution offered a device that could adapt to the beneficiary's lifestyle. Instead of a cord, it featured a remote control system in the form of a watch or pendant. The remote control activated a transmitter equipped with a loudspeaker and microphone, allowing the beneficiary to communicate within a range of 50 to 60 m^2 without needing to pick up the phone. Today, *teleassistance* functions as a *daily* support service, providing a permanent connection between the beneficiary and:

- **Their social environment:** including family, friends, and community contacts;
- **A team of specialist advisers:** providing expertise and guidance in fall prevention;
- **A medical platform:** ensuring coordination and monitoring of the patient's health.

This device is available 24 hours a day, 7 days a week. It enables users to feel reassured at all times, to receive regular monitoring, and, in critical situations such as a fall, to contact the platform directly. The beneficiary, therefore, has access to a personalised system that can be activated with a simple gesture using the remote control, or automatically through detection in the event of alarming behaviour, such as a *fall*.

2.2.3 Caregivers and Assistive Technologies

Before considering future solutions for supporting elderly people at risk of falling, it is essential to understand the role of caregivers within this framework. *Clinical services*, *nursing homes*, *retirement homes*, and *teleassistance* are all external actors in the life of a person experiencing a loss of independence. However, *family members*, *friends*, and *neighbours* also play a crucial role in providing effective *support* to individuals as they become increasingly dependent.

There are two main types of caregivers: *family caregivers* and *informal caregivers* [54]. These are individuals who assist a dependent person, either as part of their close circle or because they have been chosen by the person himself or herself. As long as the individual remains at home, the caregiver plays an essential role in ensuring that their needs are properly met. *Teleassistance* can serve as a valuable complement to this role, providing a form of presence even when the *caregiver* is not physically there. In this way, *teleassistance* becomes a reassuring tool for both the beneficiary and their loved ones.

In a *hospital* setting, *doctors* and *nursing staff* are able to *reassure* caregivers and explain the patient's condition [69]. In *nursing homes* or *residential care homes*, any staff member who has interacted with the resident may provide information, but this requires that they have monitored and observed potential abnormalities in the preceding days. When a person is already being cared for by *clinical services*, both the patient and their loved ones expect that every effort will be made to ensure *optimal support*. By contrast, when an individual still lives independently at home, he or she is alone again as soon as the caregivers leave. This is why *teleassistance* has gained importance in recent years: it offers reassurance while enabling long-term monitoring.

Teleassistance remains in its early stages, and the solutions currently available are not yet sufficiently reliable. Consequently, recent research has focused on enhancing existing systems while also exploring new alternatives [126]. The shortage of clinical staff has further accelerated the search for innovative approaches. Pendant devices, for instance, are now also used in medical facilities, where they not only allow patients to call for assistance but can also alert staff to unusual behaviour.

Innovation and new approaches are required in this field to address the challenge of continuously monitoring a patient's condition. Equally important is ensuring effective communication among all parties involved in providing care and support to elderly individuals.

2.3 Biomechanics, Behavioural Analysis, and Clinical Strategies

The use of new technologies to support healthcare teams requires a thorough understanding of the relevant warning signs to be monitored. In the context of falls, specific risk indicators can be identified. As mentioned in the Section 2.1.2, *tests* are carried out to assess the patient's weaknesses as accurately as possible. However, these *tests* remain *subjective* in several respects. For this reason, current research focuses on developing tools to monitor and automate certain assessments, with validation by healthcare professionals. The ultimate goal is to improve monitoring and provide stronger support for clinical staff.

2.3.1 Gait Analysis: Biomechanical Principles and Applications

During clinical assessments, the healthcare team must carefully observe the patient's movements. For example, during the *TUG test* presented in Section 2.1.2, clinicians may observe whether the patient uses the armrests of the chair when standing. This is one of several indicators reflecting lower-limb confidence and *balance* during the test. However, due to repetition and staff shortages, certain details may be overlooked or misinterpreted. Moreover, each patient is unique and may display behaviours that complicate interpretation. Nevertheless, there are specific observable indicators that can still serve as reliable points of analysis.

In the medical field, *gait analysis* is one of the most widely used methods for assessing pathologies. "*Walk*" can be described as a controlled sequence of falls [101]. More precisely, *biomechanics* defines it as the translation of the entire human body resulting from the rotational movements of the lower limbs. Although "*Walk*" may appear simple and automatic, it is in fact a learned process that develops during childhood and becomes natural through daily repetition. Locomotion is made possible by the alternating pressure exerted on the ground by the legs, a process referred to as the *gait cycle*.

a) Biomechanics of the Locomotor Cycle

Biomechanical analysis [101, 115, 149] of walking covers several areas:

- **Kinematics**: spatio-temporal description of movement;
- **Dynamics**: study of the forces involved in movement;
- **Electromyography**: muscle actions involved;

- **Energy**: energy expended during the movement studied.

The *gait cycle* is commonly described in terms of successive phases that repeat continuously [101]. A *single cycle* consists of alternating periods of *double support* and *single-limb support* [115]. It can be divided into two main phases: the *stance phase*, during which the limb supports the body, and the *swing phase*, during which the limb moves forward.

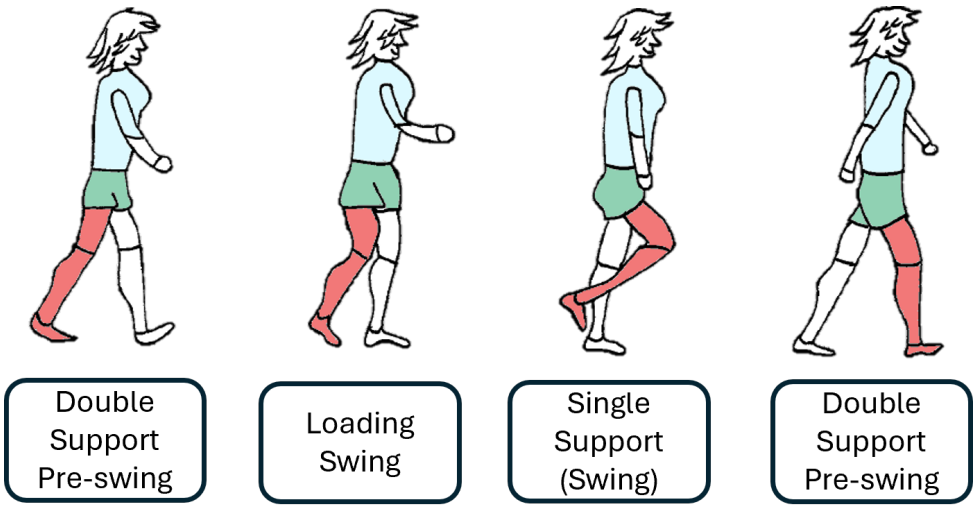

Figure 2.10 – Walk Phases

Each of these *phases* can be further subdivided into four key periods, corresponding to the phases shown in Figure 2.10. During the *stance phase*, the *cycle* begins with *double support (pre-swing)*, when the limb prepares to leave the ground while the opposite limb begins stance, followed by *loading swing*, as the limb starts to transfer weight and moves slightly forward. The *swing phase* then consists of *single support (swing)*, during which the limb is entirely in the air propelling the body forward, and ends with *double support (pre-swing)*, when the limb contacts the ground again, completing the *cycle*.

The cycle of the right lower limb is used as a reference. To enable comparison between individuals, the *gait cycle* is standardised as a percentage of its total duration. *Normal gait* refers to walking unaffected by any pathology. By contrast, *pathological gait* or *hemiplegic gait* refers to autonomous walking that is influenced by a medical condition or hemiplegia.

b) Developments in Gait Analysis Methods

Gait analysis is not a recent development in *medicin*, the observations of human locomotion date back to ancient times. In the 15th century, Leonardo da Vinci showed a

particular interest in the *anatomy* and *mechanics* of walking, studying the role of muscle insertions and the concept of the *centre of gravity* [59]. In the 19th century, experimental tests were developed using mechanical devices and platforms to analyse walking, particularly in relation to the pressure exerted on the ground. By the 20th century, advances in technology enabled the integration of *dynamics*, *kinematics*, and *centre of mass* analysis, providing a deeper understanding of muscular function during locomotion.

One of the earliest tools for observing gait is the *human eye*, a method still employed by healthcare professionals today [59]. However, this approach has clear *limitations*. While it allows clinicians to form initial impressions or identify potential concerns, it is insufficient for accurately assessing a patient's condition. For this reason, a variety of instruments and technologies have been developed to *quantify gait* more precisely. Table 2.6 summarises the main methods used for gait analysis throughout history, highlighting the parameters measured and key remarks for each technique.

Table 2.6 – Evolution of Gait Analysis Methods over Time

Period	**Method**	**Parameters Measured**	**Remarks**
15th century	Visual observation	Gait form, posture	Qualitative, limited accuracy
19th century	Mechanical platforms	Ground pressure, force	Early quantitative assessment
Early 20th century	Stopwatch, tape measure	Speed, cadence, stride length	Basic spatio-temporal parameters
Mid 20th century	Pedometers, photoelectric cells	Speed, step count	Simplified measurement, low preparation
Late 20th century	Instrumented mats, cable-based systems	Spatio-temporal parameters, plantar pressures, ground reaction forces	Improved precision, more complex setup
2000s	Accelerometers, potentiometer goniometers	Joint angles, kinematics	Portable, lightweight, but can interfere with natural gait
21st century	Vision-based systems, inertial sensors (IMUs)	Kinematics, ground reaction forces, centre of mass, detailed motion capture	Miniaturised, passive markers preferred, highly precise

Initially, a stopwatch and tape measure were sufficient to measure basic *spatio-temporal*

gait parameters such as *speed*, *cadence*, and *stride length*. Pedometers were later introduced to simplify the process; although they did not increase accuracy, they required less preparation. To measure walking *speed* over short *distances*, photoelectric cells were sometimes employed, as they offered greater reliability [59]. However, relying on a single metric quickly proved insufficient for comprehensive *gait analysis*. In response, systems involving cables in shoes were developed, eventually leading to the adoption of instrumented mats. These mats enable the measurement of *spatio-temporal* parameters as well as plantar pressures and ground reaction forces.

The capture of movement initially emerged within *biomechanics* for the study of changes in joint angles, typically using potentiometer goniometers. In the 2000s, medical services began to adopt *accelerometers*, as explained in Section 2.2.1, due to their small size and lightweight design [59]. This approach allows for the study of an individual's *kinematics* during movement. However, it soon became clear that *accelerometers* alone were insufficient, given the complexity and volume of data collected. Furthermore, these devices can interfere with natural gait due to their bulkiness.

Within this framework, *kinematic* analysis is now conducted using vision-based systems. Measurement systems can be classified into four categories: *optical*, *electromagnetic*, *ultrasonic*, and *inertial*, with inertial systems being the most recent development [59]. Thanks to the miniaturisation of technology, it is now possible to *combine accelerometers, gyroscopes, and magnetometers* for detailed motion capture. In *gait analysis*, two types of markers are commonly used: *active* and *passive* [59, 115]. *Active* markers may interfere with natural movement, as the individual must wear equipment to emit signals. *Passive* markers, by contrast, are tracked by cameras, minimising mechanical constraints on walking. For this reason, current systems tend to rely more heavily on passive markers for gait observation.

c) Biomechanical Approaches to Fall Prevention

In recent years, studies have focused on the use of *force platforms* and treadmills in combination with other measurement tools [59, 139, 160, 153, 60, 149]. The goal is to allow healthcare professionals to observe the *biomechanical* phases of walking across all limbs that influence an individual's gait. This includes monitoring the behaviour of the hips, knees, and ankles, assessing the quality of the swing and stance phases, and evaluating the individual's ability to maintain their centre of mass during locomotion.

Several theories have been proposed to characterise and define walking in a generalised

manner. Among them, two modelling approaches are widely discussed: the *six-determinant theory* and the *inverted pendulum model*. Recent studies have shown that the *inverted pendulum model* is particularly effective at predicting various characteristics of gait [59]. Pathological walking can lead to atypical segmental movements, which can be analysed, at least to some extent, using the *inverted pendulum model*, for example, by adjusting parameters such as unsteady walking speed.

Finally, *clinical gait analysis* remains a challenging task, requiring the presence of an experienced observer capable of detecting *overt* gait disorders. Today, healthcare services can use video recording to review and clarify uncertain aspects of a patient's gait. They also have the option of performing *three-dimensional gait analysis*, either in motion laboratories or using reflective markers with infrared cameras. Despite these advancements, the most commonly employed tools continue to be *force platforms* and *analysis software*.

To reduce the risk of developing pathologies, individuals are strongly encouraged to engage in regular physical activity, as it helps maintain both physical and mental health [8]. As noted in Section 2.1.2, when assessing a patient's physical condition, healthcare professionals analyse the individual's movements, paying particular attention to whether each limb performs the movements required for balance correctly. *Activities of daily living* (ADLs) are also crucial. These include *Walk*, *Shop*, *Climb stairs*, and independently *Sit Down* and *Stand Up* from a chair.

The *clinical tests* described in Section 2.1.2 incorporate elements of *daily activities*. By asking individuals to perform simple, repetitive tasks that are familiar to them, clinicians can observe their behaviour in conditions that closely resemble real-life situations. The primary goal is to assess the movements of the legs, including the feet, ankles, pelvis, and knees. Understanding what the clinician observes is essential for developing tools that effectively support their assessment and meet their needs.

2.3.2 Assessment of Risk Through Activity and Behaviour Analysis

Since the establishment of clinical services and facilities to support ageing, technologies have continued to evolve. Studying movement in individuals' daily lives has become essential for maintaining their full or partial independence. Regional authorities and public institutions tend to prioritise people's well-being to prevent premature admission to care facilities. This approach addresses both economic and health-related concerns [21].

The expansion of such facilities underscores the persistent *shortage of staff and limited time* available to provide adequate care for each patient or resident, creating significant costs for local authorities tasked with ensuring their proper operation and sustainability [51]. At the same time, research has shown that the longer an individual can remain at home, the better their outcomes in terms of *physical* and *psychological* health. Prolonged home living has been associated with lower rates of depression and fewer urgent medical interventions.

a) Identification of Activities Relevant to Risk Assessment

As highlighted in Section 2.1.2, practitioners must observe specific activities in order to identify potential pathological signs. This includes not only "*Walk*" but also daily activities. To monitor these observations effectively, it is necessary to clearly define them, enabling the development of coherent solutions for detecting and preventing health risks. Future technologies in this field will need to establish a point of correspondence between existing metrics and those that could be newly introduced. In other words, tools must be adapted to current clinical observations, while healthcare teams must be prepared to integrate and make use of these new technologies.

In the context of detecting and preventing *fall* risks, it is important to identify the *parameters* that could contribute to a reliable and explainable *fall risk score*. The study of *human activity recognition* has gained increasing importance in recent years across various fields and applications. Research by [45] predates the development of modern *sensors* and the application of *classification models* to such problems. Before addressing the *recognition or classification of activities* to generate predictive insights, it is necessary to define what constitutes a *human activity*. *Human activity* can be described in multiple ways, with its definition varying according to the study or context. In our case, the focus is on activities of *daily living*, which include actions such as "Walk", "Sit", "Stand up from a chair", "Fall", and "Pick up an object".

The central challenge is to determine whether a person's gait has changed in a way that increases their risk of falling. This requires distinguishing between a "*normal*" gait and a "*risky*" gait. Research by [45] explored whether it is possible to recognise an individual solely by observing their gait. Such insights could inform the selection of parameters for predictive models, allowing us to identify which criteria reliably characterise a person's gait and, subsequently, to determine when changes might pose a risk.

b) Behavioural Recognition of Individuals and Activities

According to [45], psychologists have identified several relevant distinctions in movement execution, including *speed* fluctuations, the type of *movement* performed, whether it is *static* or *dynamic*, and whether secondary movements influence the recognition of the primary activity. One particularly notable concept is the interrelation of two invariants in activity analysis: first, the underlying *dynamic* aspect of movement, i.e., the specific movement of an observed limb and its *speed*; and second, the *structural* unity, i.e., how all limbs evolve together in relation to one another during the activity.

It is important to establish a *taxonomy* of relevant activities in order to define the most pertinent parameters for predicting fall risk. The study by [45] validates a method originally proposed by [84, 83], which involves identifying key points in a person's movement by placing light markers on strategically selected locations. This approach is conceptually similar to modern *Motion Capture* techniques or *pose estimation*, which are used before video-based activity recognition models.

c) Activity-Based Characterisation of Risk Indicators

Finally, the concept of basing activity recognition on movement *dynamics*, *speed*, *distance* travelled, and associated micro-movements for each limb appears to be a crucial parameter for accurate recognition. However, the study highlights the importance of the *temporal* aspect, which has often been overlooked by psychologists, and suggests that it may also be relevant to consider the different stimuli applied during the performance of the activity.

As part of this thesis, direct interviews with practitioners provided insight into concrete observations made in a clinical context. The tests conducted, as described in Section 2.1.2, included the *TUG* test and the *Tinetti* test. Both assessments consider a range of *criteria*, such as *instability*, *stride length*, and *gait symmetry*. One notable observation is that, despite the structured framework provided by patient assessment grids, many comments can vary, and observations often differ depending on the practitioner administering the *test*. Building on these observations and the *criteria* defined by the grids, the objective is to identify and discuss additional parameters that could be relevant for inclusion in a targeted fall-risk prediction model.

Companies are beginning to offer solutions for future hospitals and medical facilities [158]. *Artisight* [11] and *Virtusense* [182] are two examples. Both solutions aim to leverage arti-

ficial intelligence to monitor, supervise, and communicate with patients throughout their stay. For instance, if a patient attempts to get out of bed when their vital signs or overall health status indicate it is unsafe, the system can issue a voice alert to prevent them from doing so. There is an ongoing race to develop solutions that are reliable, cost-effective, and responsive. However, many of these technologies continue to raise *ethical* concerns and questions.

2.3.3 Institutional Vision and Developments in Clinical Practice

The implementation of the *Anti-Fall* Plan [21] in France has prompted a review of care protocols and the coordination between different services. The government has issued a set of specifications and articulated its vision for the plan's development [51]. The primary objective is to provide a form of streamlined support while maximising the sharing of information among all involved services.

Their observations on the current situation are as follows:

- **Ageing and perception**: Ageing is invisible and stigmatised in society;
- **Acceptance process**: Successful fall prevention depends on various stages of acceptance;
- **Nature of falls**: Falls are often a symptom of physical and/or social frailty, not just a cause;
- **Informal prevention**: Fall prevention is mainly carried out informally by family and friends;
- **Role of caregivers**: Caregivers are the trusted individuals closest to the person receiving care and have the greatest impact on them;
- **Local actors**: Local and everyday institutional actors, who provide a framework of trust, are the first point of entry for prevention;
- **Prevention pathways**: There is no single prevention pathway, but rather a combination of multi-actor preventive measures;
- **Access to information**: Prevention messages and activities primarily reach those who are able to access information;
- **Social connections**: Social ties are the cornerstone of prevention;
- **Teleassistance**: Often used as a tool for reassurance after a fall rather than as a prevention tool.

The *government* aims to implement *five preventive measures* to address these issues: *raising awareness*, *detecting problems*, *informing the public*, *providing support*, and *building networks* [51]. Among the planned initiatives are the enhancement of *fall prevention* strategies and the use of *teleassistance* as a preventive tool. The overarching goal is to reduce admissions to hospitals and medical facilities.

The proposed scenario involves launching a communication campaign targeting individuals approaching the age of 65, as they are at increased risk of falls [51]. This outreach could be conducted via mail. Concurrently, the *government plans* to introduce the role of caregivers to relatives using a structured roadmap. Municipalities could provide voucher booklets to encourage older adults to participate in activities that promote health and mobility. To prevent risks, particularly falls, the standardisation of operators' services will be essential. In this context, *teleassistance* would evolve beyond its traditional role of observing, reporting, and contacting relatives or emergency services.

The partnership between *teleassistance* providers and healthcare professionals will become increasingly legitimate, enabling truly personalised monitoring at home. The operator will also assume a *daily social role*, providing reassurance to both beneficiaries and their relatives. At any time, the beneficiary will be able to consult their calendar and request information from the operator [51]. Conversely, the operator may periodically check in to ensure that the individual has taken their medication or has not missed appointments. Additionally, a form of shared health record between *teleassistance* providers and the *healthcare sector* could be established to enhance *coordination* and *continuity of care*.

The theme of this thesis aligns closely with these five preventive measures. First, it contributes to raising awareness and providing information by communicating about falls and associated risks. Our goal is to detect signs that correlate with clinical observations to prevent potential incidents. Ultimately, once the care process is completed, our findings will be used to support beneficiaries. Finally, the data we aim to collect is intended to be shared with healthcare professionals to enhance patient monitoring and intervention strategies.

2.4 Problem Statement and Conclusion

The *ageing* population is accompanied by an increase in chronic conditions and *fall* risks, representing a major public health challenge. This chapter has emphasised the importance of understanding the causes, consequences, and contributing factors of *falls*

in eldrely people, while highlighting the need for appropriate *clinical* support to preserve their *autonomy*.

The evolution of care structures, from hospices to medicalised facilities and *teleassistance* services, reflects a continual adaptation to the needs of an ageing population. At the same time, advances in *gait* and locomotion analysis have allowed a more precise understanding of the *biomechanical* and clinical mechanisms underlying balance loss.

Nevertheless, despite *technological and medical progress*, several *limitations* remain: insufficient human resources, variable reliability of measurements, and difficulties in integrating these tools into patients' *daily* lives without causing *intrusion*. These observations underscore the need for more integrated, accurate, and unobtrusive approaches capable of providing continuous and personalised assessment of *fall risk*.

Addressing this gap requires novel solutions that combine *reliability*, *non-intrusiveness*, and *real-time monitoring*. Current methods are either limited to clinical environments, lack sufficient accuracy, or disrupt the *daily life* of elderly people, leaving a critical unmet need for *practical and effective fall-risk assessment*. Figure 2.11 summarises the main aspects discussed in this chapter.

In this context, the Chapter 3 focuses on an *emerging technological approach: radar-based human activity recognition*. This contactless, *privacy-preserving* technology offers new possibilities for *movement analysis and fall detection*. Chapter 3 will present the *fundamental principles of activity recognition*, the *associated data processing and learning techniques*, and a *review of radar-based datasets* used to evaluate these methods in both *clinical and experimental settings*. By addressing the *limitations* identified here, this chapter establishes the foundation for technologically informed solutions to the *clinical challenges of fall prevention*.

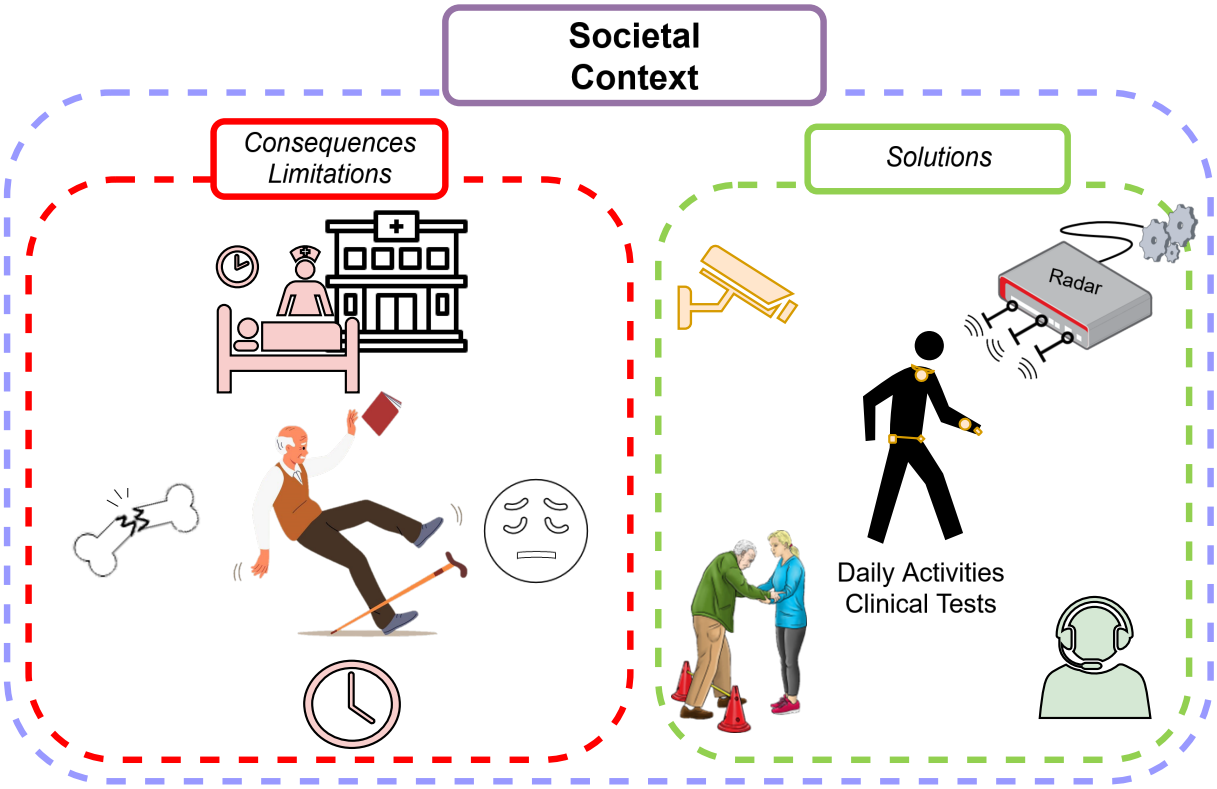

Figure 2.11 – Summary of Societal Context, Chapter 2

Chapter 3

Radar-Based Human Activity Recognition: Review of Concepts, Techniques, and Datasets

The study of *activity recognition* has been a topic of interest for several decades. Human beings have always had an innate tendency to observe their surroundings in order to adapt more effectively. With the advancement of technology, *human activity recognition* has emerged as a key area of research, with applications across diverse fields such as defence, sport, healthcare, and navigation.

As mentioned in Chapter 2, the field of healthcare encompasses a wide range of concerns. From the simple monitoring of general physiological parameters to the investigation of complex pathologies, the emergence of new sensors and advances in image processing have profoundly transformed *medical practice*.

Nevertheless, the challenges encountered in clinical practice have prompted the scientific community to develop new solutions. These challenges are manifold and demand approaches that can accommodate the *diversity of patients and pathologies*. A *shortage of time and staff*, together with the *subjectivity* inherent in certain analyses and assessments, remain major obstacles. Moreover, the large number of *criteria* and special cases makes accurate and systematic monitoring particularly difficult.

The main questions that arise can be summarised as follows:

- Which *sensors* are best suited to specific clinical needs?
- Which technological solutions are most appropriate for particular *pathologies*?
- Which computational methods can be used to identify the *relevant parameters*, both for *healthcare teams* and for *patients* themselves?

The aim of this work is to propose low-cost solutions capable of transforming simple measurements into complex indicators that are essential for enhancing medical monitor-

ing, thereby benefiting healthcare professionals, patients, and their families alike. These solutions could also be deployed in the context of home care, providing an initial level of support for individuals who remain independent.

In this chapter, we first review *activity recognition* from a technological perspective. We discuss the various sensors employed in this field and summarise the main applications of the data they collect, with the aim of highlighting the criteria that guide the selection of equipment according to the study context. In the *medical* domain, this choice is particularly critical, as it concerns not only the patient but also the healthcare team and the family, who must agree on the relevance and acceptability of the monitoring method.

We then examine the *data processing* methods necessary for *activity recognition*. This encompasses pre-processing steps as well as classification approaches, some of which have been specifically developed to enhance performance and minimise resource requirements in the context of *embedded* solutions.

Finally, we discuss the *available datasets*, which remain limited in this field. After providing an overview of existing datasets, we select one for our experiments and carry out a detailed analysis. Figure 3.1 provides an overview of the structure and methodological flow of this chapter.

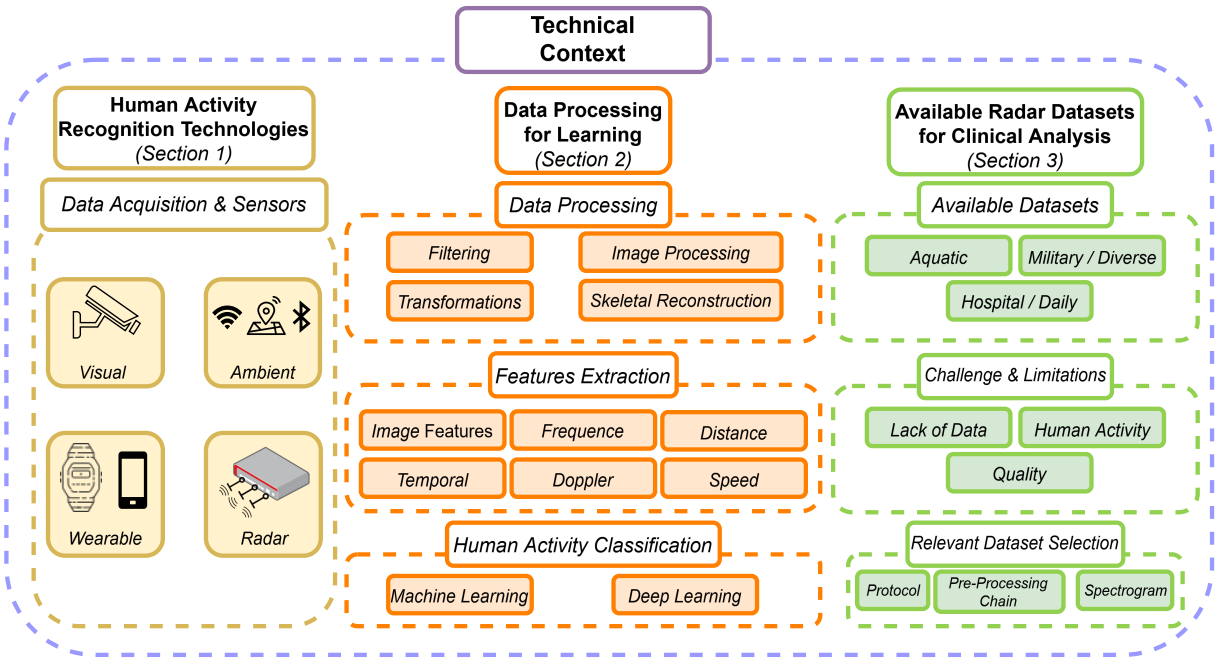

Figure 3.1 – Overview of Chapter 3

3.1 Human Activity Recognition: Foundations and Technologies

To monitor various conditions, particularly falls, which constitute the main focus of this thesis, *clinicians* consider specific *parameters*. Physicians observe patients to assess how each individual's risk evolves with age. Over time, clinical departments have, through accumulated experience, learned to identify certain signs through careful observation.

Patients are then asked to perform specific activities to determine whether these are being carried out correctly. In this sense, such observation can be related to *activity recognition* or the analysis of failures, signs of weakness, or anomalies. All of these terms *fall* within the domain of algorithmics and signal analysis that pertains to *activity recognition*, in this case *human activity*. The concept of *activity recognition* is not new [194, 193, 76]. It has been applied for several decades in fields such as defence, surveillance of critical sites, and biometrics, for instance in access security.

In signal and *image processing*, it is necessary to define *activity recognition*. This concept is highly specific and relies on several underlying principles. An activity can be interpreted in various ways, so it is essential to understand its defining characteristics in order to identify the relevant criteria for recognising it.

Various sensors have been developed and applied in the field of *human activity recognition*. However, the use of these technologies to support healthcare teams and home carers remains relatively recent. *Ethical* considerations and the *non-intrusive* nature of solutions are increasingly regarded as *important criteria* by both physicians and patients.

3.1.1 Activity Recognition: Concepts and Applications

Chapter 2 highlighted that for centuries, human beings have relied on their *eyes* to observe changes and characteristics in order to guide their actions. This makes it possible, for example, to *detect anomalies* in a person's *gait* or to identify distinctive traits of an individual. Such observations even raise questions about the ability to recognise someone based solely on their actions or movements.

With the advent of new sensors and the emergence of image processing techniques in the 1920s [157], extensive research in this area began to flourish. These studies eventually gave rise, in the 1950s, to new fields such as computer vision and the early development of artificial intelligence, as demonstrated by pioneering systems such as the Logic

Theorist [128] and the Perceptron [151].

Activity recognition began attracting the attention of several research communities in the 1980s [157, 45], owing to its potential to provide personalised support across a wide range of applications. Its strong connections to diverse fields such as sociology, medicine, and human–computer interaction have further contributed to its *interdisciplinary* appeal.

a) Activity Recognition in Computer Vision

At its core, *activity recognition* seeks to identify the actions and goals of one or more individuals by analysing sequences of observed behaviours while also considering environmental conditions. Because of its *multidisciplinary* nature, the concept is often described in various terms, including plan or goal recognition, intention or behaviour recognition, and even the estimation of an action's location [16].

This diversity presents challenges in the field, as there is no universally accepted *taxonomy* defining what qualifies as an *activity* to be observed [16, 202]. Nevertheless, most studies rely on shared principles and general definitions. The various methods typically distinguish between actions according to their intrinsic complexity. In other words, an *action* can be understood as a *sequence of continuous movements over time*, where each movement is itself characterised by information of varying complexity. For instance, its *direction*, *speed*, and impact on the type of data used, such as video recordings. A commonly accepted framework for breaking down and defining the terms and characteristics of an action [16] is as follows:

- **Gesture:** The elementary movement of a body part. It represents the most basic semantic unit, closely tied to the directly perceived motion. For example, "Node the head" or "Raise a hand";
- **Action:** The *activity* of an individual that may consist of *several gestures* organised over time, such as "Walk" or "Strike";
- **Interaction:** A *human activity* that involves either two or more individuals acting together, possibly with objects, or a single individual engaging with at least one object. For example, an interaction could be "two people pushing each other" or "Take an object to the rubbish bin";
- **Group activity:** Similar to an *interaction*, but defined by the participation of multiple individuals engaged in a shared action. It is characterised by the involvement of

a large number of people within the same scene. Examples include a "group meeting" or a "marching crowd".

Furthermore, one of the main challenges in *activity recognition* lies in the wide variety of perspectives from which *human actions* can be captured, regardless of the type of sensor employed. This phenomenon is evident even in human observation: depending on their position in a room, two observers may perceive the same *action* quite differently. The same principle applies to *sensors* such as cameras, Wi-Fi antennas, or *radar* devices, each of which records the action from its own perspective.

b) Challenges of Activity Recognition

Zhang et al. [201] illustrate this challenge by generating *skeletal representations* of the same *posture* from different camera angles, as shown in Figure 3.2. Their work highlights two main sources of variation. First, in practical scenarios, *camera* viewpoints are inherently flexible, and different perspectives can produce substantially different *skeletal representations* of the same scene. Second, the actor may perform the action in varying orientations over time, further increasing the diversity of observed views.

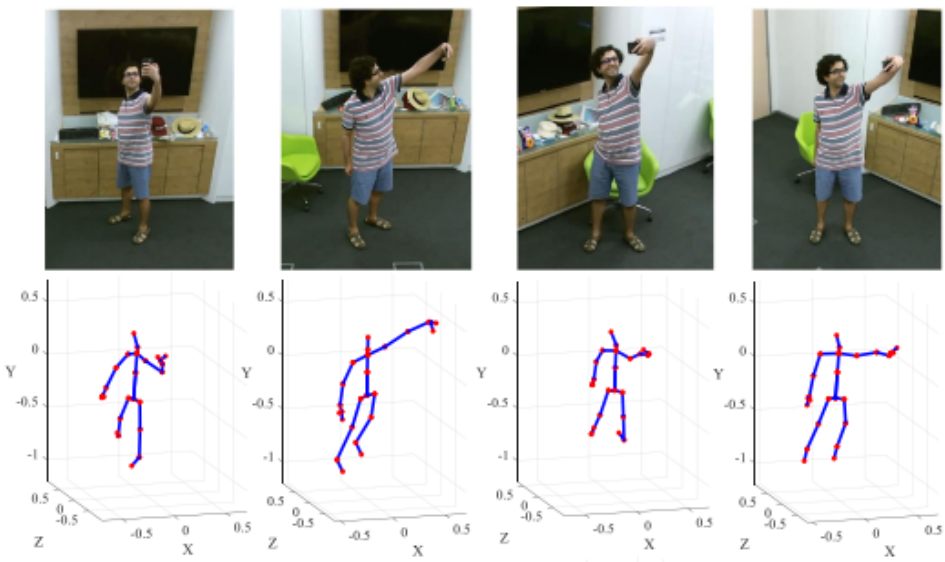

Figure 3.2 – Skeletal Representations of the Same Posture Captured from Different Viewpoints, Positions, and Camera Angles [201]

Relying solely on cameras to process images, whether through filtering, *reconstructing skeletons* from multiple viewpoints, or estimating poses, remains a challenging task. Ad-

dressing this difficulty would require the collection of larger *datasets* encompassing a wide variety of poses, gestures, and viewpoints. Alternatively, activity recognition could be approached by developing classification methods based on simple yet highly discriminative concepts that remain independent of sensor placement.

In *clinical context*, the most frequently studied activities are those referred to as *ADLs* [163, 132, 111]. These encompass both *actions* and *interactions*, with research often focusing on fundamental movements such as "Walk", "Sit", "Stand up", "Climb stairs", and "Pick up an object". Across different fields, there is consensus that these activities provide a robust foundation for validating results and enabling meaningful comparisons.

c) Activity Recognition: from the General Framework to the Clinical Context

In the health domain, such *activities* serve as essential *indicators* for comparison with direct *human observations*. Research that addresses pathological conditions typically emphasises activities closely related to a person's everyday life. While group activities are also of interest, they are considerably more complex to analyse and implement. Consequently, studies often begin with the simplest scenario, the observation of a single patient, and gradually extend toward multi-person settings.

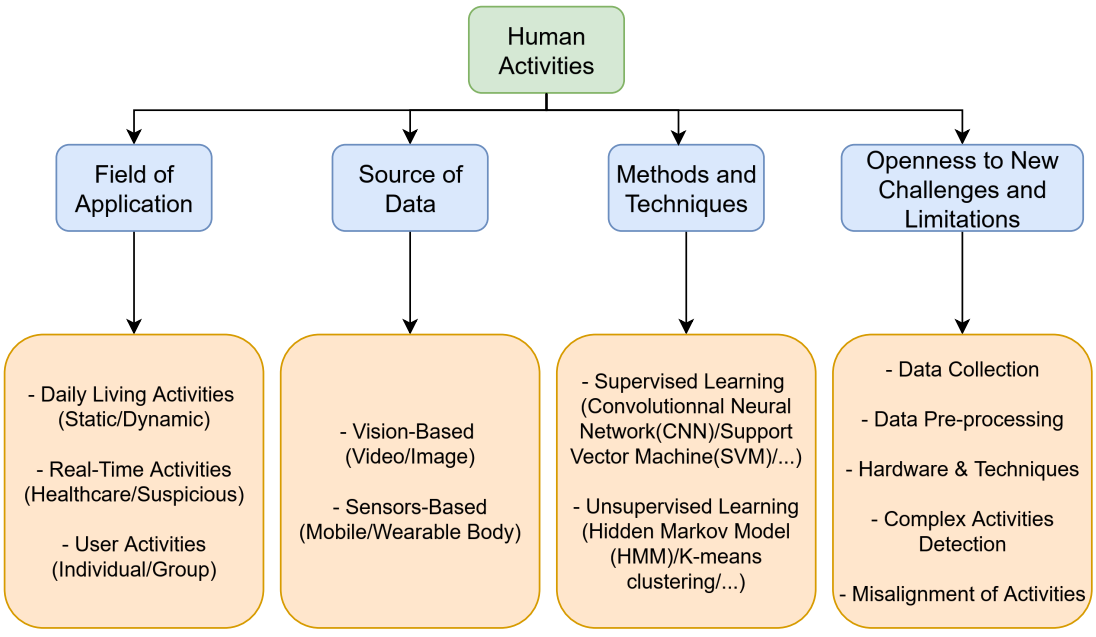

Figure 3.3 – Taxonomy of Human Activity Recognition Inspired by [10]

The study of *human activities* spans multiple areas of expertise and can be organ-

ised into several key categories [10]. Specifically, *activity recognition* relies on four main dimensions of study, as illustrated in Figure 3.3:

- **Field of Application:** This includes the types of activities being studied, such as *daily living* tasks, real-time applications, or scenarios involving different contexts, whether observing an individual, a group, or a crowd;
- **Source of Data:** Data may come from *vision-based systems*, such as images or videos captured by sensors or software capable of recording movement. Alternatively, it may originate from *body-worn sensors*, such as smartphones, smartwatches, or wearable devices;
- **Methods and Techniques:** These involve processing and enhancing the collected data to optimise it as input for classification models, which may be either *supervised* or *unsupervised*;
- **Openness to New Challenges and Limitations:** Each advancement in the field, based on the previous 3 *categories*, opens new perspectives. These include considerations such as data collection strategies, pre-processing techniques, the complexity of detecting specific movements, and the diversity of activity definitions.

The study of *human activities* encompasses a wide range of expertise and can be organised into several categories. According to a review conducted in 2022 [10], only 5% of studies focus on *health*, 24% on *dynamic activities*, and 16% on the *recognition* of a specific individual's *actions*. The field is sufficiently broad to allow the combination of multiple research approaches to address new challenges. Most reviews [10, 131, 159] indicate that so-called *daily activities*, such as "*Walk*", "*Sit*", "*Stand up*", "*Climb stairs*", "*Descend stairs*", and "*Stand*" still, remain the most frequently observed, regardless of the type of sensor employed.

For health applications, real-time and *dynamic activity* monitoring is generally preferred, as the collected data can be processed to support medical observations. In such cases, *body-worn sensors* are predominantly used, as they provide detailed information about limbs, organs, and other *relevant parameters*. However, for long-term monitoring, studies consistently highlight a key *limitation*: relying on wearable sensors significantly restricts the ability to monitor patients effectively over extended periods.

Human activity recognition can ultimately be described as a sequence of processing stages [131, 159]. The first step involves *Data Collection*, which can be achieved either by reusing existing datasets or by capturing new data using specialised equipment. The

collected data must then be analysed, validated, and categorised according to criteria relevant to the specific study. Following this, *Data Pre-processing* techniques can be applied to *filter*, *denoise*, interpret, or enhance the data, preparing it for subsequent analysis or *classification*.

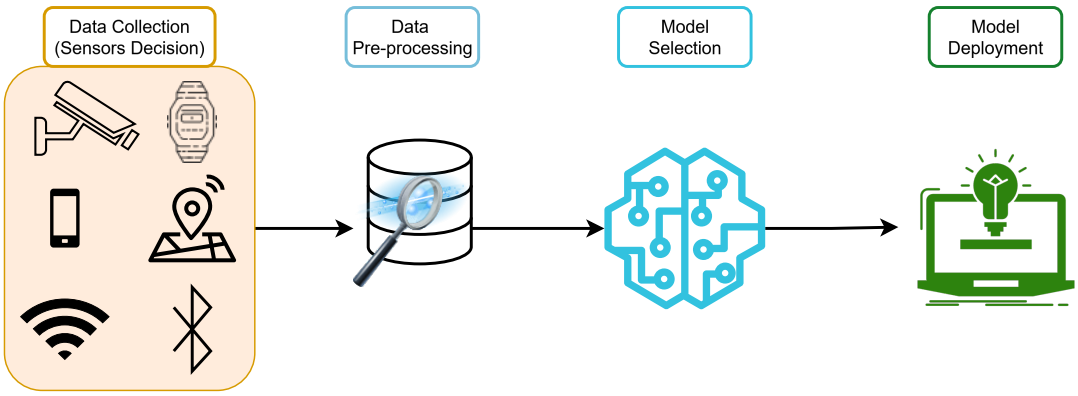

Figure 3.4 – Generic Steps in Human Activity Recognition, Inspired by [131]

As summarised in Figure 3.4, a *Model Selection* process is then performed, typically based on *machine learning algorithms*. The selection of the model depends on several *factors*, including the complexity of the data, available computational resources, and desired performance *metrics*. Once all processing stages are operational, the complete system is integrated in the *Model Deployment* phase as a unified framework to recognise *human activities* according to specific predefined *criteria*.

3.1.2 Sensor-Based Systems

One of the initial steps in *activity recognition* is the collection or selection of data that will ultimately be used to develop the complete system. As discussed in Section 3.1.1, numerous studies [10, 131, 159, 76, 169, 174, 108] have already explored specific types of sensors. Some investigations focus on a *single* sensor, while others adopt a *multimodal* approach. The choice of sensors depends largely on the application. For instance, *conventional cameras* may be sufficient for surveillance purposes, whereas monitoring a patient's behaviour typically requires wearable devices such as accelerometers.

The roles of users and other relevant stakeholders must also be taken into account. The objectives of the users, as well as the subject of observation, need to be clearly defined. In the medical field, healthcare professionals possess the expertise required to observe the *relevant parameters*. The patient's *health status* and *underlying pathologies*

impose certain constraints. The *ethical* considerations of the proposed solution are of paramount importance in the development of any intervention. Finally, the processing of the collected or utilised data must be carefully tailored to explore as many potential solutions as possible and to identify the most optimal approach for the given context.

The combination of *medical knowledge* and *pathological constraints* gives rise to specific requirements. Individual cases can then be grouped into different categories, enabling solutions to be generalised to address multiple needs using the same device. Only the interpretation of the collected data will vary, depending on the medical requirements.

Table 3.1 provides an overview of the main categories of sensors commonly used in *activity recognition*. It summarises their key *advantages*, *limitations*, and typical *applications*, offering a comparative perspective that facilitates the selection of the most suitable technology for a given clinical or home monitoring context.

Table 3.1 – General Comparison of Sensor Categories for Activity Recognition

Sensor Category	**Advantages**	**Limitations**	**Applications**
Vision-Based [48, 27]	• Well-established • Easy to interpret • Provides rich visual data	• Raises privacy concerns • Sensitive to lighting and occlusions • Requires line of sight	• Activity recognition • Pose estimation • Validation studies
Wearable [48, 27, 47]	• Reliable for motion tracking • Provides physiological data • Generally low cost	• Must be worn continuously • Can be restrictive or uncomfortable • Battery dependence	• Gait analysis • Fall detection • Clinical and rehabilitation monitoring
Ambient [197, 57]	• Non-intrusive • Preserves privacy • Enables continuous monitoring	• Installation constraints • Limited accuracy • Requires signal processing	• Home monitoring • Fall detection • Context-aware environments
Radar-Based [100, 152, 74]	• Contactless • High spatial and temporal resolution • Robust to lighting conditions • Preserves privacy	• Complex signal processing • Emerging technology • Higher cost than Wi-Fi	• Gait analysis • Fall detection • Biomechanical monitoring • Privacy-sensitive environments

a) Vision-Based Systems

In computer vision, there are various types of cameras. Cameras facilitate the collection of a sequence of information which, when analysed together, allows changes between frames to be observed [7]. These include infrared, *conventional* (RGB), *event-based*, and *thermal cameras*. Each type possesses specific characteristics that enable it to capture very different kinds of information. In the context of *activity recognition*, RGB and infrared cameras are the most commonly employed [48].

- **RGB cameras**: widely used due to early studies in computer vision [48, 27]. Numerous *datasets* are available, facilitating comparison of results. However, they present some *limitations*: *visibility* to patients may alter behaviour and affect *clinical parameters*. Sensitivity to light variations can lead to information loss, and *ethical* concerns arise regarding *privacy*, especially in private spaces or during *falls*.
- **Infrared cameras**: primarily used for validation, particularly in *Motion Capture* (MoCap) systems [31, 168]. Advantages include recovery of *ground truth* for *radar* or other sensors, the ability to superimpose *MoCap* data on *radar* measurements (*velocity*, *acceleration*, *distance*), and usefulness for *skeleton reconstruction* and movement studies. *Limitations* are high cost, substantial equipment requirements, and potential *intrusiveness* due to multiple cameras.
- **Event-based cameras**: capture discrete events at each pixel rather than full images [188]. They offer *low-latency*, *high-dynamic-range* data but remain *visible* to the patient, are relatively *expensive*, and still at an early research stage.
- **Thermal cameras**: detect body *temperature* variations, useful in specialised contexts such as surgery or post-operative monitoring. However, in many routine cases, simpler instruments (e.g., thermometers) are sufficient. They are also *visible* to patients and relatively *costly*.

b) Wearable Sensors

Accelerometers and gyroscopes remain the most commonly used sensors in computer vision [48, 27, 47]. As some of the earliest devices to be studied and implemented in services, they are inexpensive and supported by extensive *datasets*, reflecting their widespread use. These sensors must, however, be worn in order to collect data, enabling the analysis of a person's *health status* [139]. This requirement imposes constraints on the individual,

as their consent is essential for the device to be used. Moreover, because these sensors are battery-powered, periods of recharging inevitably result in data loss.

A further *limitation* lies in their performance: the *false alarm rate* for *fall detection* remains very high, at around 90% [107]. Within the family of wearable sensors, the *Inertial Measurement Unit* (IMU), which integrates accelerometers, gyroscopes, and magnetometers, has emerged as a more advanced alternative. *IMUs* have proven more reliable in detecting changes in *speed*, but they remain more costly and impose similar constraints on the wearer [187, 185].

c) Ambient Sensors

The third category of sensors used in *activity recognition* corresponds to the *ambient sensor* approach. These systems are generally regarded as *non-intrusive*, as they do not need to be worn by the individual and can be deployed in a manner that minimally interferes with *daily life* and *respects* the user's privacy. Examples include ultrasonic sensors, microphones, *pressure-sensitive* floors or *mats*, and Wi-Fi [197, 57]. Such technologies have been developed with the aim of supporting continuous monitoring within the home environment while minimising disruption and intrusion into users' *daily lives*.

Among *ambient sensors* that have attracted the most attention in the clinical field in recent years, *pressure-sensitive mats* and Wi-Fi are the most widely studied [100, 152]. *Mats* or floors equipped with *pressure sensors* are capable of collecting rich data on an individual's *centre of gravity* while walking [26, 170]. Parameters such as walking *speed*, step count, *distance* travelled, and *gait asymmetry* can be extracted from these measurements. Despite their potential, such solutions are prohibitively expensive for large-scale deployment in healthcare services and are therefore limited to specific use cases. Moreover, the installation of these systems in private homes is often impractical due to the considerable constraints involved.

Wi-Fi [197] is a wireless communication technology that is inexpensive and already widely deployed in homes and healthcare facilities, most commonly for internet connectivity. It consumes little energy and, unlike cameras, is not affected by changes in light conditions. Similar to *pressure-sensitive mats*, Wi-Fi-based systems generally raise fewer privacy concerns, making them a potentially attractive solution for unobtrusive monitoring. However, Wi-Fi provides lower spatial resolution and accuracy compared to *radar*-based systems.

Given the growing importance of radar in contactless monitoring, this technology is

discussed separately in the following section.

d) Radar-Based System

Radar technology has recently emerged as a promising approach for *contactless* and *privacy-preserving* activity recognition [50]. In particular, *Frequency-Modulated Continuous Wave* (FMCW) radar offers high temporal and spatial resolution, enabling the capture of fine-grained motion information without requiring any physical contact or visual data [74, 160, 55].

In contrast to Wi-Fi, *radar* systems can operate reliably in various lighting and environmental conditions and can be discreetly installed, for example behind furniture. This non-intrusive nature enhances user acceptability and scalability, particularly in healthcare and home-based monitoring scenarios [50].

The data obtained through *radar* are typically more precise than those collected via Wi-Fi-based systems but require advanced signal processing and machine learning techniques for interpretation. Despite research being relatively recent, *radar* sensing shows strong potential for *robust* and *privacy-conscious* biomechanical monitoring, including *fall detection* and *gait analysis*.

As summarised in Table 3.1, *Vision-based* solutions are easily interpretable but raise privacy concerns and are sensitive to environmental changes. *Wearable sensors* allow detailed physiological monitoring but remain intrusive for long-term use. *Ambient sensors* such as Wi-Fi and pressure mats are less invasive but can lack accuracy. *Radar-based* systems, however, combine the advantages of *privacy preservation*, *robustness*, and *contactless sensing*, making them particularly promising for future developments in *activity recognition* and *fall risk assessment*.

3.1.3 Radar-Based Activity Recognition

The *clinical context* described in Chapter 2 highlights the need for a reliable sensor for *fall detection* and *prevention*. In addition, caregivers require support in their decision-making and scoring processes. Ideally, the sensor should be *non-intrusive*, causing no disturbance to the patient, while providing sufficient information on the movements performed. *Radar* technology, represented by the *red square* in Figure 3.5, appears to best meet these criteria. Moreover, the system can be designed to be easily portable.

The use of *radar* offers an effective means of supporting the identification and cate-

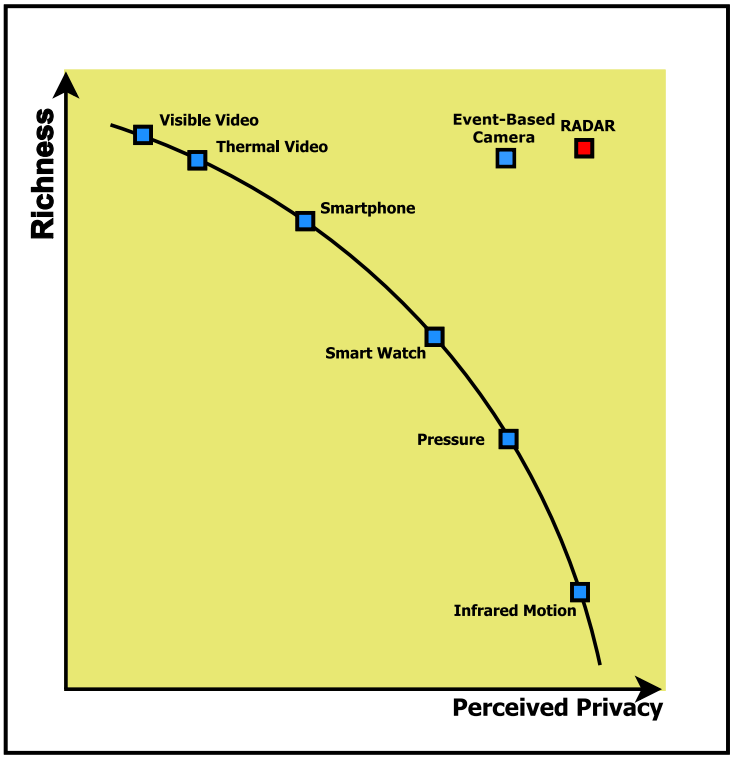

Figure 3.5 – Richness VS Perceived Privacy (Adapted from [49])

gorisation of *human activities*. *Radar*, which stands for "*Radio Detection and Ranging*", is a technology that employs radio waves to detect and track nearby objects. It is applied across a variety of fields, including defence, navigation, weather monitoring, and healthcare.

a) Radar in Activity Recognition

In recent years, extensive research [100, 152, 104, 65, 114, 72, 109, 32, 186] has focused on the use of *radar* for *fall-risk prevention and human activity recognition*. Various *pre-processing methods* have been investigated to *extract information* relevant for clinical analysis. The primary advantage of *radar* in *activity recognition* lies in its ability to analyse *Doppler signatures* and micro-Doppler effects [37, 38, 62, 172, 32, 199].

The change in *Doppler frequency* resulting from the movement of an object is referred to as the *Doppler signature*. *Micro-Doppler* describes the *small fluctuations* in the Doppler *signature* caused by secondary motions. For example, these variations may arise from hand movements during walking, superimposed on the primary motion of the body.

To obtain these *signatures*, the activity must be recorded using a *radar* system, as

illustrated in Figure 3.6. The *radar* emits a radio wave and receives the reflections returned by surrounding objects. Analysis of the received signal provides information such as Doppler frequency, time delay, and micro-Doppler signatures.

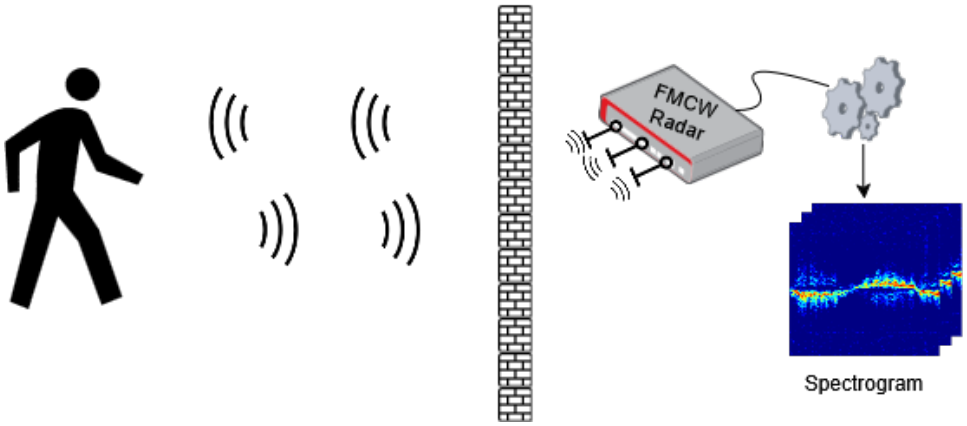

Figure 3.6 – Overall Operation of Radar Data Acquisition

Alexandre Bordat's thesis [23] presents a comparative study of *Continuous Wave* (CW), *Ultra-Wideband* (UWB), and *FMCW radars*. This work builds upon that preliminary study.

The advantages of a *CW radar* [168, 200] include low cost, privacy preservation, and ease of integration with furniture. However, it offers no range resolution, is limited to single-subject tracking, and provides restricted scene coverage. *UWB radar* [67] enables accurate 3D tracking and robust *gait event detection*, but it requires a *multi-radar setup*, is more complex, and is less portable. *FMCW radar* [112], on the other hand, allows precise limb tracking, offers high angular resolution, and is well suited for diagnostic purposes. Its drawbacks include hardware complexity, sensitivity to calibration, and a high computational load.

Our study focuses on *single-subject tracking*, though future adaptations will need to accommodate multiple individuals. For clinical testing, a single sensor may be sufficient. The system must also be easily portable and straightforward to set up. Despite the complexity of the data generated, *FMCW radar* therefore appears to be the most suitable solution for our work.

b) Extraction and Processing of Radar Data

The *radar* signal processing chain aims to extract meaningful information about target motion from the *raw data* collected by the sensor. To achieve this, a sequence of

operations is applied to progressively transform the *radar* signal into more interpretable representations [57, 20, 97, 62, 176]. As illustrated in Figure 3.7, the processing workflow comprises four main stages: *data acquisition*, *range transformation*, *clutter removal*, and *time–frequency analysis*. Each stage contributes to isolating and characterising the *dynamic* components of the observed scene.

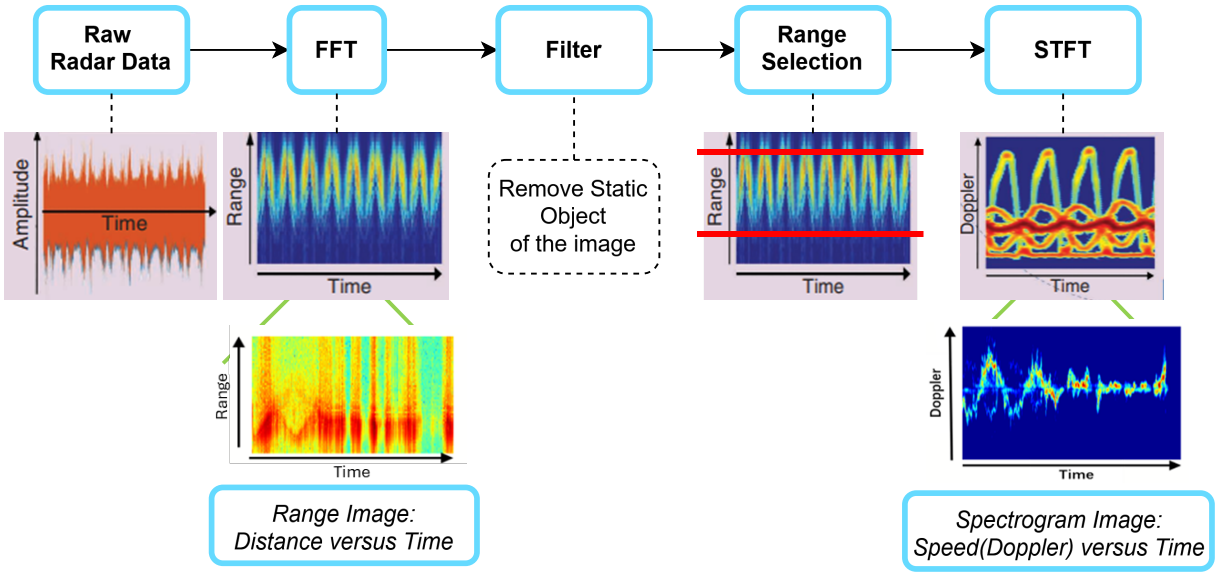

Figure 3.7 – Overview of Radar Processing Chain

- **Raw Radar Data**: The process begins with the acquisition of *raw radar* signals, typically represented as complex baseband samples containing both *amplitude* and *phase* information. These data correspond to the backscattered echoes from all objects within the *radar*'s field of view. At this stage, the signals are unprocessed and do not yet provide any direct *spatial* or *kinematic* information.
- **Range–Time Transformation (*FFT*)**: To obtain *spatial* information, a one-dimensional *FFT* is applied to each *radar* sweep (chirp). This operation converts the signal from the time domain to the *range* domain, thereby estimating the *distance* of reflecting targets. The result is a range–time map, where each row represents the range profile of a single frame, and the columns correspond to successive time samples. This transformation provides the first interpretable representation of the *radar* scene in terms of *distance versus time*.
- **Static Clutter Removal (*Filter*)**: Fixed objects such as walls, furniture, or the ground contribute strong but uninformative reflections known as *static clutter*. To

enhance the visibility of moving targets, these stationary components are suppressed using a *temporal* high-pass *filter* or mean subtraction along the time axis. The resulting filtered *range–time* data emphasise dynamic reflections corresponding to motion within the scene.

- **Time–Frequency Analysis (*STFT*)**: The *filtered* data are subsequently analysed in the *time–frequency* domain using the *STFT*. This transformation estimates the *Doppler* content of the *radar* signal over time, allowing the extraction of *velocity* information related to target motion. The resulting representation provides a compact description of the *dynamic* behaviour of the observed scene, capturing how frequency components evolve with time.

Overall, this processing chain establishes the foundation for subsequent radar data analysis, providing well-structured signal representations from which motion-related features can be derived.

c) Radar Movement Analysis and Deployment

The thesis of Shelly Vishwakarma [183] further emphasised the analysis of individual limb movements using *Micro-Doppler signatures*, providing a more detailed understanding of activity dynamics.

The study reported in [31, 168] demonstrated that *MoCap* can be used to obtain the *ground truth* for data recorded with *radar*. This represents a significant advantage for ensuring the reliability of both the collected data and its subsequent analysis. *MoCap* data can be superimposed on the *radar* data, allowing verification of the *Doppler* information obtained for each body part of interest in the study [31].

MoCap systems can be implemented using two main types of devices, the most common being *passive* sensors attached to the person being observed. Small reflective markers, often micro glass beads, are placed on specific points of the body, enabling *real-time reconstruction* of the person's *skeleton* [149]. This allows their movements and gestures to be visualised as they occur. The system achieves very high accuracy, as the cameras operate with an extremely short refresh rate of approximately *1000 frames per second*.

The studies [175, 164] have demonstrated that *machine learning* algorithms can be applied to *radar* data to recognise *human activities*. They present the initial results obtained using *classification* models for *human activity recognition*, with most studies relying on generated or simulated data.

Several *radar*-based *fall detection* solutions are already available on the market. Some remain at the prototype stage, while others are already deployed in users' homes. One example is the *Vayyar Walabot Home* [77], a *radar*-based *fall detection* system. This solution employs 3D imaging technology based on point clouds to monitor individuals within a defined area.

Falls are detected by comparing changes in movement and position within the point clouds against patterns of *normal activity*. Alerts are then sent in real time to carers, family members, or friends via a mobile application.

In the context of *activity recognition*, *classification models* require a sufficient amount of data to learn effectively. Moreover, these data must be consistent with the field of study, taking into account the objective of monitoring *daily activities* for clinical prevention purposes.

A variety of technologies have been explored for *human activity recognition*, ranging from *vision-based systems*, such as *cameras* and *MoCap*, to *wearable* inertial sensors, and *radar*-based solutions. *Vision*-based systems offer high *spatial resolution* and detailed *motion analysis* but can be affected by *occlusions* and raise *privacy* concerns. *Wearable* sensors provide accurate *motion tracking* and *portability*, yet they require user *compliance* and may be *intrusive* for continuous monitoring.

In contrast, *radar*-based systems enable *non-intrusive* monitoring under diverse environmental conditions, while *preserving individual privacy*. They also provide both *spatial information* and *micro-Doppler signatures* that are particularly well-suited for capturing the *dynamics* of *human activities*. Taking these factors into account, *radar* technology has been selected for this study. The subsequent section focuses on the data *processing* and *learning* techniques applied to *radar* data to extract meaningful *features* and support accurate *activity recognition*, as reported in the literature.

3.2 Data Processing and Learning Techniques

Human activity recognition can be divided into several areas which, when combined, enable the achievement of various objectives. Section 3.1 highlighted the need to first collect data before processing it for use as input to *classification* models. In the context of *falls*, *radar* appears to be the most viable solution. Once a sensor has been selected, it is essential to understand how to process the acquired data in order to interpret it correctly.

Radar data are pre-processed to generate *spectrogram* images, as illustrated in Figure 3.7.

These images can then be interpreted in a variety of ways. *Image processing* methods may need to be applied to emphasise previously unknown characteristics. The relevance of these features can subsequently be assessed by healthcare professionals.

The next objective is to select lightweight *classification* models in order to propose a complete solution. The aim is to develop a portable system capable of acquiring new data, processing it, and interpreting the results. This section will discuss the potential methods for processing and interpreting *spectrograms*, before evaluating the current state of the art in *classification* approaches within this field. Finally, the examination of more advanced methods will allow an assessment of whether the final solution requires additional data *acquisition* and *classification* strategies.

3.2.1 Data Processing and Evaluation

In the context of *radar*-based *human activity recognition*, a *spectrogram* represents a visual depiction of the frequency content of a *radar* signal over time. It illustrates how the signal's *energy* is distributed across different frequencies as the activity progresses, enabling the analysis and interpretation of distinctive patterns and features associated with specific *human activities*.

A notable advantage of *radar*-based *human activity recognition* is the ability to extract a person's *velocity* from the *spectrogram*. This feature allows for the estimation of the *speed* at which individuals are moving, based on the *spectrogram*'s characteristics.

a) Signal Processing and Spectrogram Generation

The *pre-processing chain* commonly employed in studies follows the procedure described in Section 3.1.3, enabling extraction of two types of information: *distance* over time (*range–time*) and *speed* over time (*spectrograms*).

Range–time data can be used to analyse distances travelled and the number of trips to and from the *radar*, whereas *spectrograms* allow movements of individual body parts to be examined via their *speed*. Studying early *warning signs* in *falls*, such as changes in activity *speed*, is particularly relevant. Consequently, most studies rely on *spectrograms*, which are the focus of our work.

A *spectrogram* can be represented as a matrix or an image. While matrices preserve detailed *speed* and *time* information, *feature extraction* is challenging, especially as *radar*-based *activity recognition* is still developing. Therefore, *spectrograms* are generally used as

images, encoding movement *energy*. For example, considering the walking phases discussed in Section 2.3.1, step details for each limb, particularly the feet, can be observed in Figure 3.8.

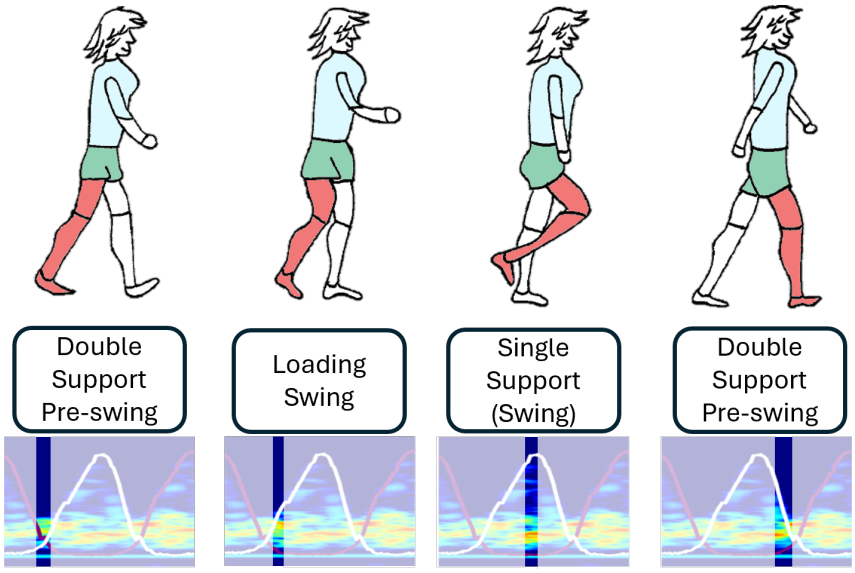

Figure 3.8 – Spectrogram Correspondance with Walk Phases

Spectrograms differ from conventional data such as photographs, video, or inertial measurements. As shown in [183], they allow the movements of individual limbs to be distinguished, enabling estimation of a person's *speed* and *distance* during an activity. Each activity produces a characteristic *signature*, which *activity recognition* algorithms can use to differentiate between *spectrograms*. Distinguishing these *signatures* is not straightforward and requires *machine learning* to handle variability in activity *speed* and *limb micro-movements*.

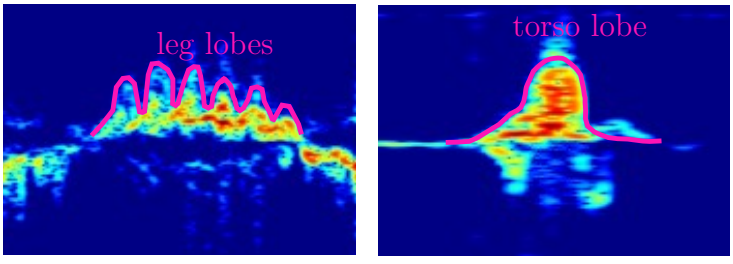

Walk Fall

Figure 3.9 – Comparison of Walk and Fall Signatures

Figure 3.9 shows two *spectrograms* (one "*Walk*", one "*Fall*"), with lines highlighting the

lobes relevant for our discussion. These examples are drawn from a FMCW radar dataset [57], illustrating the differences between "*Walk*" and "*Fall*" patterns. For "*Walk*", *leg lobes* repeat periodically, reflecting higher *velocity* during locomotion. In contrast, a "*Fall*" is dominated by *energy* from the *torso* lobe, followed by a gradual decrease in *speed* until movement ceases. The *torso*, having a *larger surface area*, reflects more *energy*, so the *spectrogram* predominantly captures the *torso lobe*.

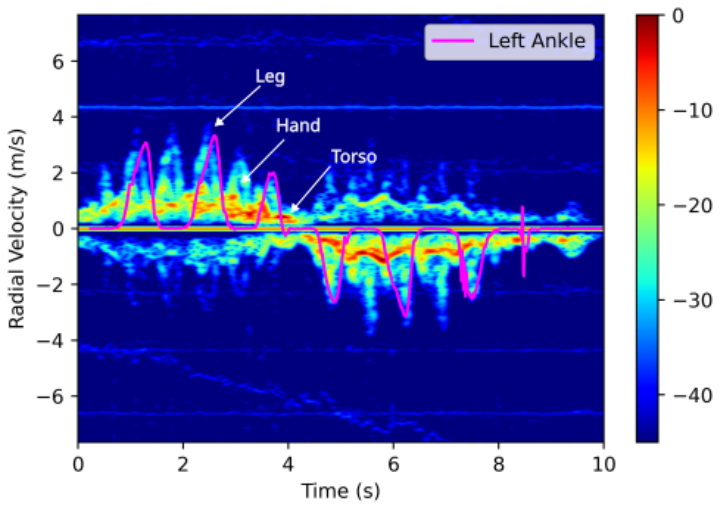

Figure 3.10 – Micro-Doppler Spectrogram Representation of Limbs, Obtained with our Experimental Setup

Figure 3.10 shows one of our *radar* acquisitions, with *MoCap* data superimposed to highlight the observed lobes in practice. In our study, we focus on *FMCW radar*. Range resolution depends on bandwidth, while *Doppler* resolution depends on signal duration and the number of *FFT* points in the *STFT*. Alexandre Bordat [23] use a range resolution of 37.5cm and a Doppler resolution of 1.25Hz (approximately 0.03m/s), allowing precise *velocity* analysis and *activity classification*. Using these parameters, a *spectrogram* of the captured activity can be generated in Figure 3.10.

b) Image Processing Methods

Among the commonly used image processing techniques are the following, all of which can be applied to *spectrograms* [96, 102, 173, 195, 181, 113]:

- **Binarisation** [140, 134]: Produces an image containing only two classes of pixels (0 or 1 for a binary image; 0 or 255 for a greyscale image), using a threshold applied

to the original image. For colour images, the R, G, and B channels can be converted to *greyscale* before thresholding, as illustrated in Figure 3.11.

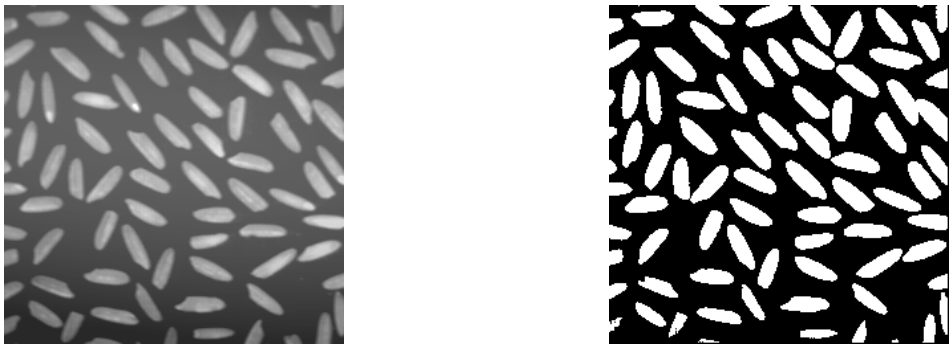

(a) Original Grayscale Rice Image (b) Binarised Rice Image

Figure 3.11 – Binarisation of Rice Image [2]

- **Filtering and noise reduction** [56, 154, 192]: Includes smoothing, enhancing image *quality*, *removing noise*, or adding noise intentionally. In *spectrograms*, conditional filtering (or "*heuristic thresholding*") can deactivate pixels in regions with no significant information, effectively reducing noise.
- **Contour detection** [94, 147]: Identifies points in an image where there is a sudden change in light intensity, highlighting *structural* elements such as edges, depth discontinuities, or surface changes. Depending on the algorithm and parameters, *contours* can be detected at different levels of detail. Additional processing, such as *segmentation*, is typically required to group pixels into meaningful regions. An example is illustrated in Figure 3.12.

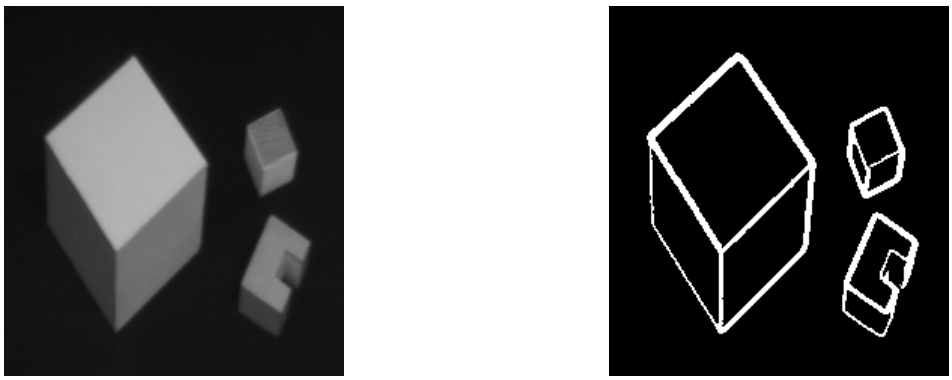

(a) Original Grayscale Cube Image (b) Cube with Contour Detection

Figure 3.12 – Contour Detection of Cube Image [2]

- **Channel analysis**: Analysing individual R, G, and B channels can provide additional insights, especially in *spectrograms* where certain components (e.g., blue) contain most of the movement-related information.

A summary of recent radar *spectrogram denoising* approaches is presented in Table 3.2. It lists the radar types, application contexts, *denoising* methods, and reported improvements in signal *quality* or *classification accuracy*.

Table 3.2 – Overview of Denoising Methods Applied to Radar Spectrograms and their Reported Benefits

Reference	Radar Type	Application / Target	Denoising Method	Classification / Evaluation	Reported Gain
Y. Lang et al.[96] (2025)	FMCW radar	Human micro-motions	Deep CNN-based denoising (DnCNN-like)	CNN classifier for activity recognition	Noise reduction ≈ 20 dB (PSNR) improvement
B. Li et al. [102] (2024)	mmWave radar	Automotive and surveillance scenes	Residual U-Net + wavelet filtering	Object detection	Enhanced clarity and target separation
C. Tang [173] (2021)	CW radar	Human activity (micro-Doppler)	GAN-based natural noise learning + DnCNN	Feature restoration for HAR	+3–5 dB SNR; clearer Doppler signatures
R. Yu et al. [195] (2024)	FMCW radar	Human activity recognition	Dual-path denoising convolutional net (DPDCNet)	CNN classification for HAR	+4 % accuracy vs. raw spectrograms
F. Tupin [181] (2024)	SAR (satellite radar)	Earth observation imagery	AI-based despeckling (self-supervised network)	Visual / PSNR evaluation	Improved texture preservation and resolution
B. Ma et al. [113] (2025)	Pulsed-Doppler radar	Target recognition (vehicles, humans)	1D DnCNN + AlexNet (cascade)	End-to-end classification chain	+6 % accuracy; better feature contrast

c) Ground Truth and Data Evaluation

A common approach to obtain the *ground truth* is to apply a verification method using corresponding *MoCap* data [184, 130, 168, 160]. This approach allows each piece of data detected from individual limbs to be highlighted in the resulting *spectrogram*. For example, the movement of the left ankle is represented in the *spectrogram* shown in Figure 3.10.

The MoCap system typically involves the use of multiple cameras to optimise camera *trilateration* [184, 168, 149], a process by which the *3D position* of each marker is determined using its measured *distances* from several *reference points*. This approach enables accurate *reconstruction* of the observed object or *human skeleton*, allowing precise tracking of movements and gestures. The use of multiple cameras ensures that each marker can be detected from different viewpoints, improving the *precision* of the position estimates and reducing the likelihood of occlusions affecting the *reconstruction*.

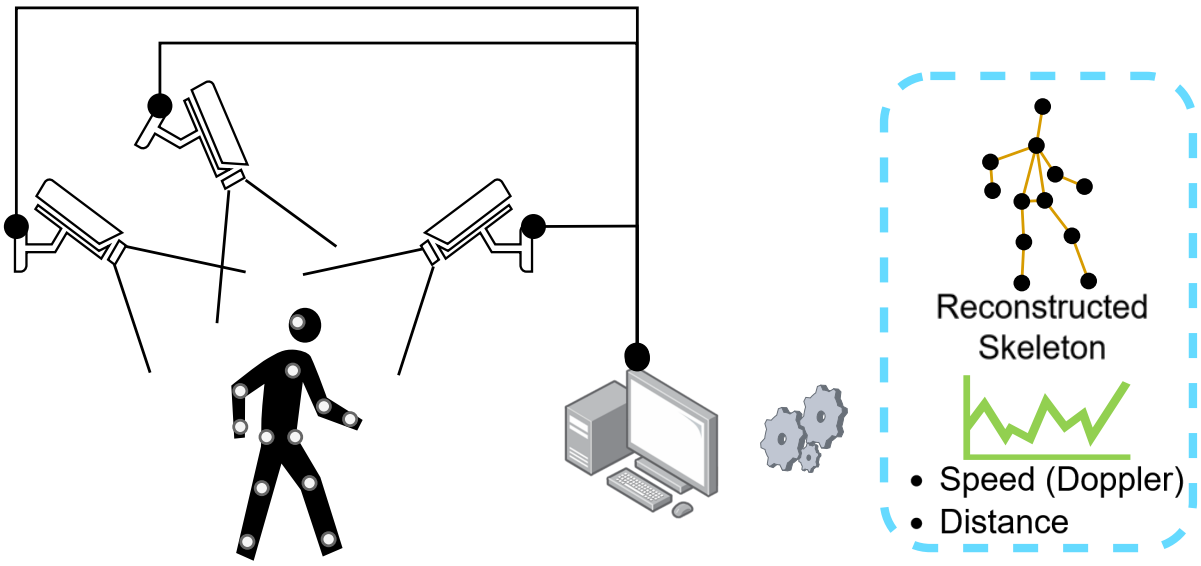

Figure 3.13 – Data Extraction from MoCap for Skeletal Reconstruction and Motion Analysis

Skeletal reconstruction enables each of the person's limbs to be labelled, allowing the tracking of each point in both *time* and *space*, as illustrated in Figure 3.13. This technology provides crucial information for the study of movement, such as the *speed* of a point or the *distance* travelled. These parameters correspond closely to the data that can be obtained using *radar*, which is why *MoCap* is often employed as *ground truth* to validate and calibrate *radar*-based solutions [168], which are more complex to implement.

Once data acquisition is *synchronised* between the *radar* and *MoCap* systems, meaning that the measurements from both systems are *aligned in time* so that each *radar* sample corresponds to the *same instant* as the *MoCap* data, the collected *data* can be *superimposed*. The error rate can then be calculated, allowing the reliability of the *radar* data to be assessed and enabling the development of solutions that no longer require *MoCap* for *validation*.

3.2.2 Learning Approaches for Human Activity Recognition

In recent years, *activity recognition* using *spectrograms* has received increasing attention [135, 88, 104, 99, 39, 114, 64, 118, 105, 186], due in part to the *non-intrusiveness* of *radar* sensors, their *respect for privacy*, and their ability to capture relevant motion information such as *distance* and *velocity*. Research has focused on the application of

classification models typically used for conventional image recognition tasks.

Spectrograms are images that each exhibit a pattern of activity, which varies according to the movements performed. *Image classification* models may offer a viable approach for *learning Doppler signatures*. *Spectrograms*, as described in Section 3.2.1, are images that each exhibit a pattern of activity, which varies according to the movements performed. *Image classification* models may offer a viable approach for *learning Doppler signatures*.

a) General Approaches to Human Activity Recognition Learning

In *human activity recognition*, several lightweight classification approaches are commonly employed, particularly for portable or embedded applications such as smartphones or *radar*. Among the most widely used are:

- **Random Forest (RF)**: A robust, straightforward, and computationally efficient model, *RF* achieves good accuracy in recognising activities such as "*Walk*", "*Run*", and "*Cycle*" [93].
- **K-Nearest Neighbours (KNN)**: A non-parametric model that classifies activities based on the proximity of feature vectors. Well suited for small *datasets* and platforms with limited resources, although its complexity grows with dataset size, often requiring dimensionality reduction techniques such as *Principal Component Analysis (PCA)* [150].
- **Principal Component Analysis (PCA)**: A multivariate statistical technique for dimensionality reduction by extracting the most significant variables. Applied to *spectrograms*, *PCA* selects the most relevant parameters before they are used as inputs to a classification model. This approach enhances model performance while reducing computational costs, particularly for lightweight systems [203].
- **Lightweight Neural Networks**: Architectures such as *Multi-Layer Perceptrons (MLPs)* and one-dimensional *Convolutional Neural Networks (1D-CNNs)* provide a balance between accuracy and efficiency, making them suitable for mobile and IoT applications. In particular, *1D-CNNs* exploit the temporal structure of sensor signals while reducing computational load compared to traditional *2D-CNNs* [133].
- **Hybrid Models**: Combinations such as *KNN+MLP* can improve classification robustness while maintaining low resource requirements, which is advantageous in *real-world* environments with noisy or partially missing signals [14].

b) Learning Approaches for Human Activity Recognition by Radar

Several studies [135, 88, 104, 198, 103, 89, 65, 98, 39, 55, 138] have been conducted using *radar* data from diverse applications, including *drones*, *animal activity*, and *human activity*. The most commonly investigated classification models include *SVM*, *KNN*, *GoogLeNet*, *AlexNet*, and *ResNet-18* [175, 164, 165, 130, 198, 20, 145, 186]. Employing these algorithms requires the acquisition of a sufficiently large *dataset* to train the model effectively and to validate the results. These models remain extremely popular, having demonstrated their effectiveness across a wide range of image *classification* tasks and challenges.

The majority of existing studies have been conducted using the *Radar Signatures of Human Activities* dataset [57]. These studies have demonstrated the feasibility of learning from *spectrograms* in the context of *human activity recognition*. Table 3.3 presents a comparison of *classification* models applied to this dataset, both with and without input enhancement methods.

Table 3.3 – Classification Models Using the Dataset from [57] (*Radar Signatures of Human Activities*)

Classification Models	**Year**	**Data Processing**	**# Parameters**	**Accuracy**
GoogLeNet [175, 164]	2021	Inception	6.7 M	74.70%
AlexNet + KNN [175, 164]	2021	Transfer learning	60 M	77.15%
AlexNet + SVM [175, 164]	2021	Transfer learning	60 M	78.25%
ResNet-18 [130]	2024	Threshold with CNN	11.1 M	92.70%
Q-SVM [104]	2024	Binarised spectrogram, use threshold & mask methods	NA	**93.10%**

GoogLeNet is a learning method built on a deeper architecture, comprising 22 layers compared with 8 in *AlexNet*. This convolutional neural network is based on the *Inception* architecture [171], which enables the network to select from multiple convolutional filter sizes within each block. In an *Inception* network, these modules can be stacked sequentially, with occasional max-pooling layers using a stride of 2 to reduce the grid resolution by half. In the study by [164], an average accuracy of 74.70% was achieved.

In the studies by [175, 164], *AlexNet* was employed to perform *Transfer Learning* on *SVM* and *KNN* models. *Transfer Learning* involves leveraging knowledge acquired from

a *pre-trained classification model* and applying it to a different, yet sufficiently similar, problem. Following feature extraction with *AlexNet*, the authors applied *Transfer Learning* to an *SVM*, achieving an accuracy of 78.25%. They subsequently applied *Transfer Learning* to a *KNN*, attaining an average accuracy of 77.15%. By employing *Transfer Learning* in conjunction with *SVM* or *KNN*, they demonstrated the feasibility of *learning* from *spectrograms*.

Finally, *ResNet-18* is a model comprising 18 layers, with reported accuracies ranging from 74% to 94%. This model has attracted attention due to its ease of use and its *compact* size in terms of the number of parameters. It is designed to have a small footprint and is suitable for deployment on *embedded systems*, eliminating the need for a computationally heavy model.

c) Adapted Learning Approaches

While deep learning models such as *ResNet-18* offer automatic feature extraction and are suitable for embedded applications due to their compact size, classical *learning models* like *SVMs* remain widely used in data *classification* and are particularly prominent in *human activity recognition* tasks. One of the main advantages of *SVMs* over other classification models is their ability to construct an optimal hyperplane that maximises the separation between feature classes. The use of support vectors is crucial, as they determine both the *position* and *orientation* of the separating hyperplane, thereby enhancing classification performance.

Other studies [104] have shown that applying an *SVM* directly to *radar* data, such as *CW radar*, can yield very good results. For instance, achieving an accuracy of 92.8%, and up to 95.4% with *FMCW radar* [104]. However, the *pre-processing methods* vary considerably and can be more or less complex and time-consuming to implement. The thesis [104] also demonstrates that *SVM* outperforms *KNN*, although it achieves only around 80% accuracy when applied to micro-Doppler data. The performance of such models is highly dependent on the choice of *kernel*, *hyperparameters*, and the *pre-processing* applied to the input data.

SVMs appear to offer a relevant solution, both in addressing hardware constraints and in providing effective classification performance. Following feature extraction, this approach has the potential to achieve results that are comparable to, or even surpass, those obtained with alternative methods.

KNN is frequently employed for *classifying human activities* based on micro-Doppler

radar data, as it effectively leverages the proximity of feature vectors extracted from *radar* signatures. Micro-Doppler provides detailed information on joint movements, enabling the distinction between *everyday activities* such as *walking*, *running*, *standing up*, *sitting down*, or even *arm swinging*. *KNN* is particularly valued in this context for its simplicity and computational efficiency, as it does not require complex models or deep learning. It performs well when the features are discriminative and the number of samples is manageable. However, its performance can be sensitive to *radar* noise and individual variations, often necessitating pre-processing or robust feature extraction [9].

3.2.3 Advanced Data Processing Techniques

The fields of *image processing* and *machine learning* are both distinct and complementary in the context of *activity recognition*. Section 3.2.2 emphasise the importance of the various processing steps that occur prior to the application of a *learning model*. They also highlight the diversity of existing models and the differences in their processing requirements. Consequently, the nature of the input data affects the quality of learning, while the models impose constraints on the type of data that can be effectively utilised.

The emergence of these two fields, and in particular their increasingly frequent combination, has given rise to new research aimed at further improving learning and data processing. In the *clinical domain*, there is a notable *lack of data* to address all requirements [95]. Consequently, methods have been developed specifically to overcome these *limitations*.

Firstly, *new data acquisition* methods combining multiple sensors have been proposed. These approaches open up new perspectives for *human activity recognition* and provide a larger and more robust data. The resulting data are both more reliable and contribute to the validation of the results obtained. However, the acquisition process can be lengthy and costly. To address this, acquisition simulation programmes have been developed [184, 86], enabling rapid expansion of the *dataset* within a short time and with minimal administrative complexity.

a) Optimisation of the Signal and Image Processing Chain

As noted in Section 3.1.1, there is a considerable need for *datasets* in the field of *human activity recognition* using *radar* data [170, 31, 129, 6]. The *Radar Signatures of Human Activities* dataset [57] remains the most widely used resource for *daily human activity*

recognition, with over 5,000 downloads worldwide. Table 3.3 summarises the results obtained using this *dataset*. Comprising 1,753 images, it is currently the most comprehensive *dataset* available. A related challenge [191] has been organised to showcase the various possibilities for utilising these data in *classification* tasks.

Research has also examined the importance of data and the processing techniques that may, or may not, be worthwhile for optimising *classification* performance. For instance, the studies [13, 40] modify the pre-processing chain to handle *radar* data. Both papers conduct a pre-processing analysis that emphasises the trade-off between pre-processing complexity and recognition accuracy. Specifically, [13] employs *FFT*, *STFT*, and *Smoothed Pseudo-Wigner–Ville Distribution* (SPWVD) as processing techniques, whereas [40] uses *spectrogram* binarisation as input.

Other studies [104, 106] have investigated the components of the data processing chain. In particular, they examine how the application of different types of masks and thresholds to raw data, along with normalisation, can enhance classification performance. According to the authors, the term "*mask*" refers to the *binarised image* employed for feature extraction.

Finally, the study [130] also utilises the *Radar Signatures of Human Activities* dataset [57] to examine data processing challenges, such as data cleaning and parameter selection, prior to considering the use of simulated data. The work [129] further investigates the effect of adjusting the threshold applied to raw data. Additionally, the recording angle can influence the resulting activity *signature*.

b) Data Augmentation Strategies

All of these studies can enhance *learning* performance, but they are quickly constrained by the limited amount of data, as simple image processing alone is no longer sufficient. Some works, such as [129, 204], have achieved around 90% accuracy. In certain cases, the use of *Data Augmentation* has allowed performance to approach 99%.

Among all these methods and studies, it becomes challenging to make proper comparisons between learning approaches due to the diversity of solutions applied to the same dataset. To ensure that the clinical approach is grounded in real data, *Data Augmentation* is not currently considered a viable solution. Similarly, with limited data, focusing on methods to enhance data processing or recognition accuracy may not be the most effective strategy.

Table 3.3 summarises *classification* models applied to the same *spectrogram* dataset [57]

and outlines the different methods employed. It highlights the use of various approaches discussed previously, such as *Transfer Learning* or image processing prior to training. These methodological variations and their associated limitations underscore the need for additional data in this field.

Furthermore, as most studies do not report the *F1-score*, comparisons are limited to accuracy. Consequently, only one study [130] provides this information, which is why it is included in Table 3.3.

c) Methods Combination and Dimension Reduction

In the field of *human activity recognition*, data and *model fusion* is frequently used to improve the robustness and accuracy of systems, particularly when using multiple sensors or heterogeneous signals such as micro-Doppler *radar* and accelerometers. Data-level fusion involves directly combining measurements from multiple sensors before feature extraction, while feature-level *fusion* aggregates feature vectors to enrich the representation of activities [87].

At the decision level, decision-level fusion, such as *majority voting* or weighted aggregation techniques, combines the outputs of several lightweight classifiers, such as *KNN*, *Random Forest* or *SVM*, to improve *classification* reliability [81, 12, 61]. These *fusion* approaches are particularly relevant for *daily activities*, where signals may be noisy or partially missing, and have shown significant gains in accuracy without drastically increasing computational complexity.

Studies employing various "*classical*" processing methods have demonstrated only modest improvements. Building on these observations, other works [41, 196, 203, 110] have shown that the use of *PCA* can also be beneficial.

3.3 Radar Human Activity Datasets: Review and Study Dataset

Radar is one of the most useful sensors for studying fall risks [74]. Its capabilities, particularly when using *FMCW radar* technology, allow the extraction of data relevant to analysing the movement of different body parts during various activities. The resulting *spectrograms* can be regarded as images composed of *signatures* that differ according to the activity performed. These micro-Doppler *signatures* are therefore essential for distin-

guishing between activities and for identifying potential warning signs of a *fall*.

A wide range of *classification models* has been proposed. However, for lightweight and portable solutions, simpler models are generally preferred, such as *CNNs*, *ResNet*, *SVM*, and *KNN*. Research has focused on how these models can be optimised and improved to ensure that deployed solutions are as reliable as possible. Data *fusion*, whether through the combination of models or their outputs, represents a promising avenue for further exploration.

To complete the *activity recognition* chain, it is necessary to identify the data that will be used in our study. While numerous *datasets* exist in the field of *activity recognition*, those employing *FMCW radar* for *human activity recognition* are comparatively scarce. This section aims to examine the most relevant datasets for our purposes and to highlight their limitations, thereby informing the choice of *dataset* that will be employed for our experiments and analyses.

3.3.1 Available Radar-Based Human Activity Datasets

Section 3.1.2 highlighted the wide range of sensors available for *activity recognition*. Each of these *sensors* has been studied and evaluated in various contexts and for different requirements. Over time, *datasets* have been developed to support and facilitate further research in this field.

For *clinical contexts*, particularly in relation to *fall risk*, *radar* appears to be the most suitable technology for this research. However, the application of *radar* in *human activity recognition* remains a relatively recent area of study, and existing *radar* datasets do not fully correspond to *human activities*.

In recent years, *radar* has attracted increasing interest as a means of capturing information on limb movements. Consequently, a few datasets have begun to emerge [100, 136, 7], each considering different types of activity and exhibiting a high degree of specificity. This again underscores the challenges of establishing a consistent *taxonomy* in the field of *activity recognition*. Since bandwidth information is generally unavailable, a proper comparison of the datasets is not feasible. All of the datasets discussed are summarised in Table 3.4.

Table 3.4 – Spectrogram Datasets

Datasets	Number & Type of Activities	Number of Participants[1]	Duration & Number of Repetitions	Number of Samples	Data Source	Radar Details	Format[2]	Open
Human Aquatic Activities, 2016 [136]	5 Aquatic Human Activities	5	Not specified	625	Simulation & Real	FMCW Radar Vivaldi Antennas with DAQ-6001	NA	N
Radar-Based Swimming Activities, 2024 [204]	5 Aquatic Human Activities	2	20s to 40s & Multiple repetitions	13,304	Real	FMCW Radar TI AWR1243	NA	N
DIAT MRADHAR, 2022 [36]	6 Military Human Activities	30	3s to 4s & 7 repetitions	3,780	Real	CW Radar AD1674	.jpeg & .mat	Y
FMCW radar dataset, 2024 [130]	11 Diverse Human Activities	Not specified	8s to 15s & 60 repetitions	19,800	Simulation from MoCap data	No radar, Simulation only	.mat	Y
Multi-Intensity Continuous Human Activities, 2024 [161]	19 Diverse Human Activities	7 (3 F, 4 M)	2 repetitions	25,000	Real	FMCW Radar TI IWR1642BOOST	NA	N
Dataset for Human Activity Recognition with FMCW, 2023 [166]	8 Diverse Human Activities	7	Not specified	2,520	Real	FMCW Radar	image	Y
PARrad, 2022 [19]	14 Hospital Room Human Activities	29 (8 F, 21 M)	Various (22h in total) & Not specified	21,569	Real	FMCW Radar TI MmWave	.h5 & .csv	Y
Vid2Doppler, 2021 [7]	12 Daily Human Activities	10 (2 F, 8 M)	2min 50s & 6 repetitions	720	Simulation from RGB Camera & Real	FMCW Radar TI AWR1642	numpy	Y
Radar signatures of human activities, 2019 [57]	6 Daily Human Activities	72 (25 F, 47 M)	5s to 10s & 3 repetitions	1,753	Real	FMCW Radar Ancortek	.dat	Y

[1] F: Female, M: Male; [2] NA: Not Available

a) Human Aquatic Activities

Among the *datasets* relating to *human activity* captured using *radar* technology, two focus specifically on *aquatic activities* [136, 204]. In both cases, the data were acquired using *FMCW radar*.

The *Human Aquatic Activities* dataset [136], as shown in Figure 3.14a, is based on the simulation of *aquatic activities* using video recordings, followed by Doppler *radar* measurements at the edge of a swimming pool. This research focuses on the ability to recognise activities in a complex environment with high levels of noise generated by water and splashes. However, this type of setup appears to be too remote from an environment whose constraints differ substantially from those of the *aquatic* domain.

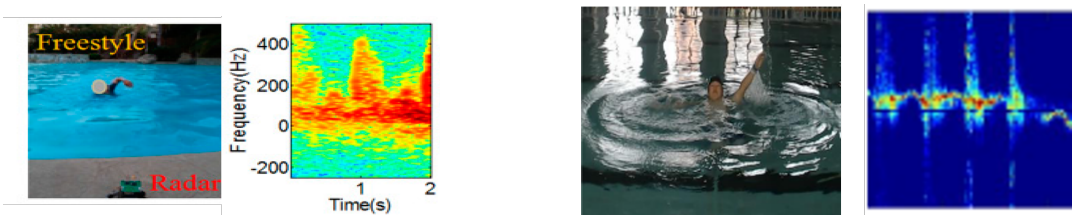

(a) *Human Aquatic Activities* [136] (b) *Radar-Based Swimming Activities* [204]

Figure 3.14 – Set of Human Aquatic Data

Radar-Based Swimming Activities [204], as illustrated in Figure 3.14b, introduces a new *dataset* collected in an indoor swimming pool. The study involved two participants who were asked to perform various swimming movements, with recording times varying according to the exercise. Five activities are proposed, including crawl and breaststroke. The objectives are twofold: first, to provide a new *radar*-based *dataset* in the field of swimming, and second, to enhance *classification* methods, which are currently ineffective. The study also proposes a *Data Augmentation* strategy, the use of a pre-trained model, and the addition of a new convolutional module.

Both studies are noteworthy for their data acquisition methodologies, and their approaches to processing *radar* data are also relevant. Nevertheless, the environmental constraints in these studies differ significantly from those encountered in clinical settings, particularly with respect to *human activity recognition* outside the aquatic domain. Consequently, while the *datasets* themselves are not directly applicable, some of the processing methods may be relevant to our context.

b) Military and Diverse Human Activities

Section 3.1.1 has shown that the *taxonomy* for defining an activity is highly complex. This complexity increases further when multiple activities are considered. Decomposing an activity to define it based on a set of characteristic movements is one of the major challenges in this field. These movements must clearly distinguish it from other activities.

In addition, the same activity can be defined in multiple ways, as illustrated by the *datasets* [36, 130, 161, 166]. For instance, "*Walk*" may be slow, fast, or rhythmic, and can be observed from different angles. As discussed in Chapter 2, from a *biomechanical* perspective, walking always consists of a *cycle* between *single support*, *double support*, and so on. However, variations in the *signature* or execution of this *cycle* can pose challenges when analysing the activity from an algorithmic standpoint.

The *DIAT MRADHAR* dataset [36], as presented in Figure 3.15a, was created to implement a *classification* model for recognising suspicious activities in a *military context*. Military marching differs significantly from everyday walking, as the *timing* and movement *cycles* are not the same. This can impact the observational model, highlighting the importance of the *acquisition* context. Despite the *dataset*'s relevance for *human activity recognition*, the proposed labels are too far removed from so-called *daily activities*. Furthermore, this *dataset* was acquired using *CW radar*, which, as discussed in Section 3.1.3, is less suitable for our study.

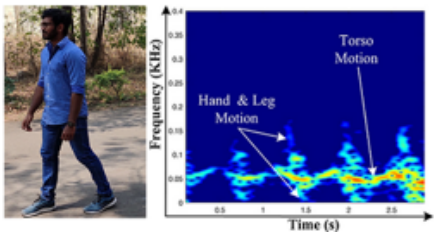

(a) *DIAT MRADHAR* [36]

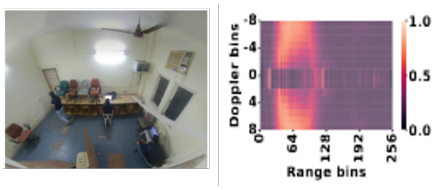

(b) *Multi-Intensity Continuous Human Activities* [161]

Figure 3.15 – Set of Military and Diverse Human Activities Data

The *FMCW Radar dataset* [130] comprises simulated data derived from real *MoCap* recordings, with different hardware-level bias configurations simulated. Specifically, using *MoCap* data and the *SimHumalator* tool provided by [184], the authors of [130] generated a version of this dataset in *spectrogram* form. Then, they conducted various tests to evaluate the performance of learning models on the generated data and to assess the impact of noise levels in the images. By adjusting noise level thresholds, the study demonstrated a

significant effect on data quality, highlighting the critical importance of parameter selection. *Masks* were subsequently applied to the data to identify regions of the *spectrogram* containing meaningful information. According to the authors, a mask is a *binary* function applied to the *STFT* that selects coefficients exceeding a given threshold, thereby isolating the signal's most significant components.

Moreover, the issue of *multi-person activity recognition* is also significant. The *Multi-Intensity Continuous Human Activities* dataset [161], as illsutrated in 3.15b, introduces multi-level Doppler analysis for *activity recognition* while allowing simultaneous recording of multiple individuals in a room. However, this *dataset* is too complex to be initially tested for *daily activity recognition*. In addition, it contains an excessive number of activity labels, with an unbalanced number of samples per activity.

The *Dataset for Human Activity Recognition with FMCW* [166] comprises a total of 2,520 samples, including both *daily* and *sports* activities. Very little information is provided regarding the acquisition method, making the dataset difficult to reuse. The recording conditions appear highly varied, and the constraints associated with the different acquisition environments are unknown.

These four datasets are relevant for recognising *human activity* in terrestrial conditions. However, the activities included are closely associated with specific contexts, such as *sport*, *military*, or *office environments*. Nevertheless, they provide valuable insights and underscore the importance of defining activities with respect to their context.

c) Hospital and Daily Human Activities

The focus of this thesis is on *falls* and the *risk of falls* among elderly or hospitalised individuals. This implies that the activities under consideration must be closely related to the *daily lives* of these individuals, enabling the rapid and accurate identification of potential behavioural changes. The objective is to develop a model capable of providing alerts, akin to the vigilance of a care team member who detects a potential *anomaly*. Ultimately, the aim is to support carers and facilitate continuous monitoring of the patient or individual.

In the context of *everyday human activities*, the *Vid2Doppler* dataset [7], as displayed in Figure 3.16a, is generated by converting videos of various activities into *spectrogram* images. This approach offers an interesting *simulation* method for a field that currently lacks data. However, it is not the most suitable solution for developing an end-to-end system that accounts for *real-world* constraints.

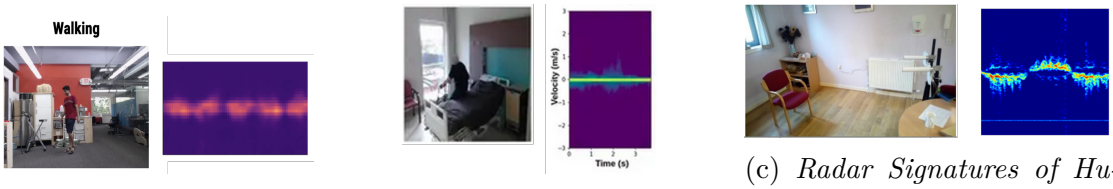

(a) *Vid2Doppler* [7] (b) *PARrad* [19] (c) *Radar Signatures of Human Activities* [57]

Figure 3.16 – Set of Hospital and Daily Human Activities Data

The *PARrad* dataset [19], as shown in Figure 3.16b, appears highly relevant to our field of study. It consists of real *radar* data acquired from patients in hospital rooms, with the aim of monitoring activities across different patient profiles. The *dataset* includes a substantial volume of data, encompassing the use or non-use of mobility aids and multiple age categories. However, limited information is provided regarding potential sources of noise, the presence of other individuals during continuous recording, and similar factors. Furthermore, the data cannot be used as provided, due to the absence of pre-processing code required to interpret it.

The *Radar Signatures of Human Activities* dataset [57], as depicted in Figure 3.16c, is the most widely used resource in *radar*-based *human activity recognition* research, with over 5,000 downloads worldwide. Consequently, it has been extensively evaluated across various aspects, including *pre-processing chain* parameters, *spectrogram* data manipulation, and the application of classification models. The *dataset* is generally balanced, with the exception of *fall* events, although some noisy and occasionally erroneous data have been identified. Notably, several studies have concentrated on image processing techniques, particularly *binarisation* and mask application, to extract relevant information from activity *signatures*.

Given these *limitations*, particularly regarding the distribution of activities, it is clearly necessary to identify a more suitable solution. Nevertheless, this *dataset* provides a valuable starting point for research and serves as a useful benchmark due to the range of activities it includes.

3.3.2 Data Challenges and Limitations in Radar Human Activity Recognition

Few fully accessible datasets exist in the field of *radar*-based recognition of *daily activities*, which constitutes a significant *limitation* in this area. Consequently, most available

data are simulated.

The most comprehensive datasets typically include only a small number of participants or focus on activities that do not necessarily correspond to *clinical* concerns. Among the *datasets* reviewed and summarised in Table 3.4, the most widely used is the *Radar Signatures of Human Activities* dataset [57].

a) Lack of Data: Augmentation Methods

Firstly, some studies transform data from *RGB cameras* or *MoCap* systems to generate *spectrograms*. For example, the *Vid2Doppler* [7] and *Human Aquatic Activities* [136] studies are based on *converting video images into spectrograms*. In addition, the *Radar-Based Swimming Activities* dataset [204] employs *Data Augmentation* techniques to maximise the available data.

Recently, some data have been generated through simulation based on *MoCap* recordings [184, 130]. Using *MoCap* data and the *SimHumalator* tool provided by [184], the authors of [130] produced a version of this *dataset* in *spectrogram* form. They subsequently conducted a series of tests to evaluate the performance of *learning models* on the generated data and to assess the *impact of noise levels* in the images.

b) Human Activity Data: Diversity and Challenges

Some *FMCW* datasets have been proposed with a variety of activities. However, these activities do not always correspond to the context of *daily activity recognition*. For instance, *Human Aquatic Activities* [136] and *Radar-Based Swimming Activities* [204] focus on *aquatic* activities, which are not representative of activities performed in a home environment. *DIAT MRADHAR* [36] is centred on *suspicious activities* in a military context, while the Dataset for *Human Activity Recognition with FMCW* [166] is a small dataset including both daily and sports activities. The latter two datasets are therefore not suitable for representing *daily activities* of ageing individuals.

The *Multi-Intensity Continuous Human Activities* dataset [161] aims to extend the domain by addressing multi-person scenarios. However, this also introduces greater hardware requirements. In a clinical context, this feature is relevant if the aim is to monitor a room with multiple residents, for example, or a corridor, which can also be a challenging area. Given that *radar*-based activity recognition is still a relatively new field, demonstrating the technology's effectiveness for a single individual before extending it to multiple people appears to be a more appropriate approach. Furthermore, this *dataset* contains an

excessive number of *activity labels*, many of which are not directly relevant to our study domain.

Finally, the *Radar Signatures of Human Activities* database [57] and *PARrad* [19] contain real *radar* data. *PARrad* is particularly interesting because it provides data collected in a *clinical* setting, with environmental conditions reflecting *real-world* scenarios. However, as explained in Section 3.3.1, we were unable to use these data, as the pre-processing code required to interpret them is not available.

Most research has relied on the *Radar Signatures of Human Activities* database [57], owing to its comparatively large size. The *dataset* is generally *balanced*, with the *exception* of the *fall* category, although some *noisy* and occasionally *erroneous* data have been identified. Despite these *limitations*, particularly regarding the distribution of activities, it remains essential to seek a more suitable solution. Nevertheless, this dataset provides a valuable starting point for research and serves as a useful benchmark due to its inclusion of *daily human activities*.

c) Data Quality and Reliability for Clinical Domain

Providing an effective solution for the *clinical prevention* of falls requires coherent and *balanced* data. In particular, *daily activity* data should include a balanced representation of each activity. During pre-processing, the type of filter applied, if any, can influence the resulting data. Furthermore, at the *pre-processing chain level*, the sensitivity of the threshold is also critical, as noted in [129].

Similarly, the data in the aforementioned *datasets* consist of either simulated data or raw *radar* recordings. For simulated data, it is challenging to assess consistency with the *ground truth*. In the case of raw *radar* data, however, if *ground truth* measurements were not collected concurrently, evaluating and verifying the accuracy of the data becomes difficult.

Nevertheless, in the clinical field, it is essential to be able to demonstrate and verify the accuracy of the collected data. The authors of [183, 168] used *MoCap* to confirm which elements could be observed in the *spectrograms*. Therefore, it would be highly beneficial to develop a solution that addresses this issue as well.

Finally, the data type, acquisition method, and *pre-processing chain* can all influence the inputs to a *classification* model. Data that are insufficiently accurate or not representative of the various activities can significantly affect the results obtained in *activity recognition*. These issues are highlighted in [129].

To enable the effective use of *classification* models based on *spectrograms*, particular attention must be paid to image quality. *Spectrograms* are time–frequency matrices converted into energy representations to facilitate interpretation. In this context, "*image quality*" refers to factors such as resolution and *noise level*, which can affect the model's ability to learn effectively and generalise to new data.

This is particularly important when *classification* models are deployed in *real-world* environments, such as hospitals. Models trained predominantly on clean or simulated data are at high risk of exhibiting poor recognition performance in clinical settings. Furthermore, inconsistencies may arise if the activities studied differ substantially from those of interest, or if the recording protocol lacks clarity, potentially affecting the accuracy of labelling.

3.3.3 Dataset Description and Pre-Processing

Sections 3.3.1 and 3.3.2 have outlined the *radar* datasets currently available. Each presents specific *advantages* and *limitations* that must be considered in relation to the working context, the required equipment, and the acquisition environment.

These observations suggest that the *Radar Signatures of Human Activities* dataset [57] best meets our criteria. It provides data acquired with *FMCW radar* in both *laboratory* and *clinical* environments. In addition, the proposed *pre-processing chain* enables the extraction of *distance–time* and *velocity–time* representations.

The *spectrograms* obtained from this *dataset* form the basis of our work. A more detailed analysis of the *dataset* makes it possible to highlight its main characteristics. This section therefore provides a concise description of the equipment used, the *acquisition* protocol, and the *pre-processing chain*. Finally, we present the generated data that will serve as the foundation for the remainder of this manuscript.

a) Material and Protocol of *Radar Signatures of Human Activities* [57]

Between *March 2017* and *March 2019*, the *University of Glasgow* carried out a data acquisition campaign. The objective of F. Fioranelli et al. was to develop the *Radar Signatures of Human Activities* dataset [57], as part of their study on *human activity recognition* using *FMCW radar* data.

To build this *dataset*, recordings were conducted in different rooms at the *University of Glasgow*, as well as in facilities providing care for elderly individuals. In total, 99

participants, aged between 20 and 100 years, were recorded while performing a range of daily activities.

In most cases, participants were asked to perform six different activities:

- **Walk:** Walking back and forth between the *radar* and a predefined point;
- **Sit:** Sitting down on a chair, with the position relative to the *radar* not specified;
- **Stand:** Standing up from a chair, with the position relative to the *radar* not specified;
- **Pick:** Picking up an object placed on the floor;
- **Drink:** Picking up a glass from a table, simulating drinking, and placing the glass back on the table;
- **Fall:** Falling from standing height onto a thin mat placed on the floor.

For older participants, the *fall* activity was not requested. This omission naturally introduces an *imbalance* within the *dataset*. Each activity performed was repeated two to three times by every participant, with the aim of capturing both *inter*- and *intra*-subject variability. In addition, recordings were carried out in nine different rooms to introduce variation in furniture and layout, thereby enhancing the robustness of the dataset.

The capture equipment consisted of a *C-band FMCW radar* operating at $5.8GHz$, with a pulse repetition period of 1 ms, a bandwidth of $400MHz$, and 128 complex samples per scan. Two *Yagi antennas*, each providing a gain of approximately $+17dBi$, were employed for signal transmission and reception. This configuration enabled the *acquisition* of detailed micro-Doppler *signatures*, capturing the fine movements of different body parts.

b) Pre-Processing Chain

As explained in Section 3.3.1, raw *radar* data cannot be used directly. It must undergo *pre-processing* to extract meaningful information about human movements. The *pre-processing chain* proposed by [57] consists of the following steps:

- **Range compression:** Each beat frame is transformed using a *FFT* to obtain the *distance* profile of the targets relative to the *radar*;
- **Stationary target suppression:** Echoes from static objects are filtered using a *Moving Target Indication* (MTI) technique, ensuring that only *dynamic* movements are retained;

- **Micro-Doppler signature extraction:** Frequency variations induced by fine movements, such as arm swings or leg bends, are extracted to capture detailed motion characteristics;
- **Spectrogram generation:** The processed signals are converted into *time–frequency* representations, producing *spectrogram* images suitable for activity analysis and *classification*.

This *pre-processing chain*, summarised in Figure 3.17, standardises the data and facilitates its use with *learning models* by emphasising the most discriminative information. The resulting files are named according to a specific *nomenclature*, enabling *classification* based on activity labels. The naming convention encodes the participant's ID, the activity label, and the number of repetitions performed.

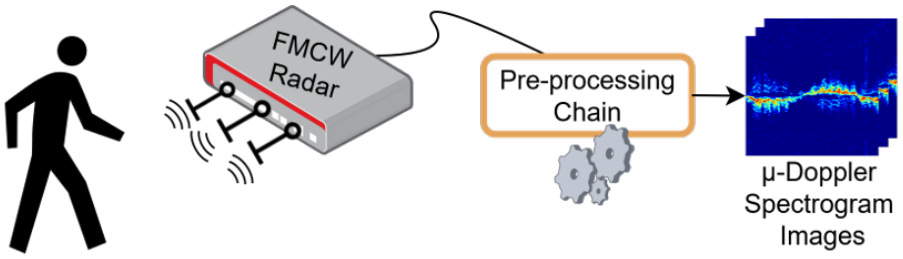

Figure 3.17 – Pre-Processing Chain Method

During data processing, we observed the presence of duplicate identifiers. However, these files were recorded on different days and do not belong to the same *acquisition* location. We assume that these duplicates result from errors in participant referencing across different *acquisition* sessions. Depending on the situation, we either ignored or renamed the affected files.

c) Spectrogram Data Specifications

Spectrograms provide a visual representation of the frequency variations of a signal over time, with amplitude encoded in the intensity of the pixels. In the context of *human activity recognition*, these representations allow the identification of *temporal* and *frequency* patterns characteristic of each movement, such as the swinging of the arms during walking or the rapid raising of a leg when standing up.

These *spectrograms* are particularly well suited as inputs for *Convolutional Neural Networks* (CNNs) and other deep learning models, which can automatically extract discriminative features. Their rich temporal and frequency information enables fine distinctions between similar activities and provides increased robustness against inter-subject and environmental variations.

The *dataset* consists of raw *radar* data, to which we apply the associated *pre-processing chain* using *Matlab* code. This processing produces the corresponding *spectrogram* data. In total, the *dataset* comprises 1,753 items. We propose to split the data, allocating 80% for *training* classification models and 20% for *testing*. The resulting *distribution* is summarised in Table 3.5.

Table 3.5 – Training Set Distribution Details

Activity Category	**Training Samples**	**Testing Samples**
Walk	264	47
Sit	265	47
Stand	264	47
Pick	264	47
Drink	263	47
Fall	168	30

This *distribution* highlights the imbalance within the dataset. Overall, the most common activities, such as *Walk*, *Sit*, or *Stand*, have a comparable amount of data, whereas falls are significantly underrepresented. Despite the explanation provided by [57], this *imbalance* may have implications for our results and subsequent analyses.

A more detailed analysis of the data enables us to assess the relevance of the *spectrograms* for classifying *human activities*. After examining the data and grouping the activities by label, we obtain the *dataset* representation matrix shown in Figure 3.18.

Figure 3.18 allows us to hypothesise that each activity performed by different participants exhibits distinct *signatures*. Some activities, such as "*Stand up*" and "*Sit down*", appear visually very similar, suggesting that the model may achieve lower accuracy for these classes. In contrast, "*Walk*" and "*Fall*" are clearly distinguishable. For "*Pick up*" and "*Drink*", however, it is difficult at this stage to discern a meaningful difference. This similarity may result from the highly standardised *protocol* for these two activities, and it will need to be verified in subsequent analyses.

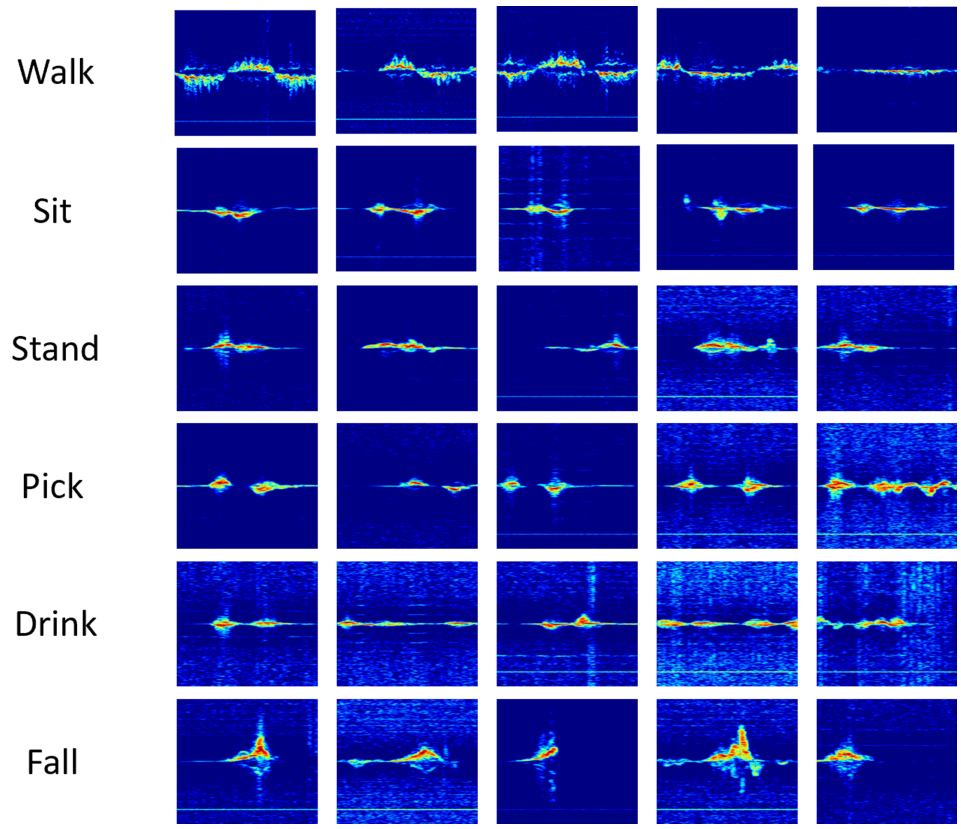

Figure 3.18 – *Radar Signatures of Human Activities* [57] Spectrogram Matrix

3.4 Conclusion

Human activity recognition is a vast and *multidisciplinary* field. This chapter provided an overview of the *foundations*, *technologies*, and *data processing techniques* used to analyse and classify *human activities* from *radar* signals. In particular, it highlighted *machine learning and deep learning* approaches, signal processing methods, and various *feature extraction techniques* designed to identify key characteristics of an *activity*. As shown in Figure 3.19, the main aspects of this context are summarised to provide a comprehensive overview of the discussed issues.

From a *clinical* perspective, these technologies offer opportunities to complement healthcare professionals' observations by providing *objective* and *non-intrusive* measurements. Among the studied approaches, *radar*-based methods are especially promising due to their *non-intrusiveness*, *privacy-preserving* nature, *spatial and temporal resolution*, and *robustness*. However, the field still faces a shortage of *datasets* tailored to the analysis of *daily activities* in clinical contexts, which limits the generalisation and validation of the

methods.

For this work, we selected a *dataset* whose protocol closely matches our intended application. Its processing chain and the resulting *velocity-time spectrograms* provide a suitable foundation for our *learning* approach. Chapter 4 will present experimental tests conducted with simple models on the *spectrogram* images, followed by the evaluation of different *feature extraction* methods and their integration into various *classification* models. Finally, the chapter will describe the *optimisation* of all parameters to achieve a robust and reliable solution for *radar-based human activity recognition*.

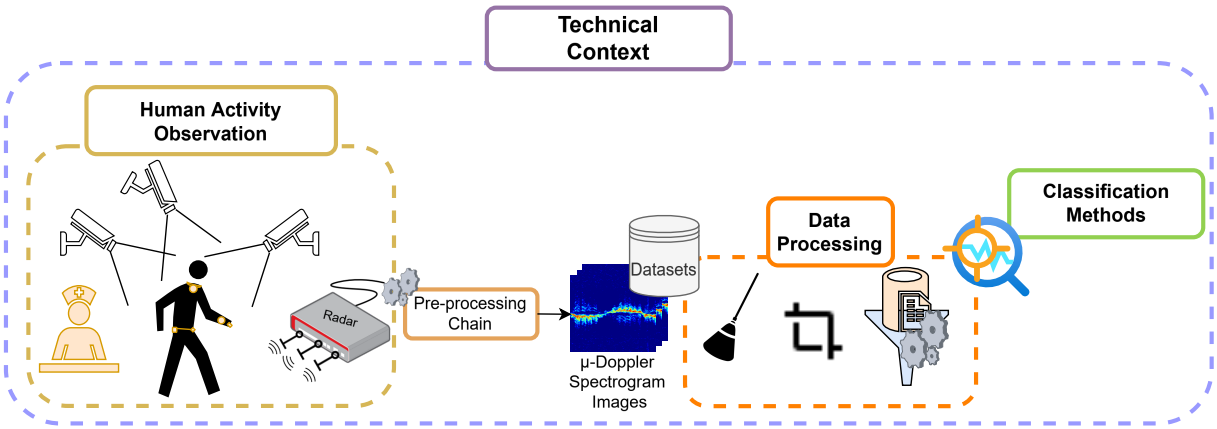

Figure 3.19 – Summary of Technical Context, Chapter 3

Chapter 4

Performance Evaluation of Feature Extraction and Recognition Methods

The aim of this work is to propose new solutions tailored to individuals at risk of falling, as well as to the carers responsible for supporting them. Chapter 2 highlighted the importance of monitoring the changing health status of anyone who may be vulnerable. Such risks tend to increase with age, and certain warning signs can often be detected in the individual's behaviour. These signs may be observed both at home and in specialised facilities, provided that the relevant criteria are assessed on a daily basis.

Technological developments, together with advances in clinical research and computer vision, have facilitated the emergence of innovative solutions. Chapter 3 examines the range of sensors and classification methods employed in the field of activity recognition. To assess potential risk indicators, caregivers typically observe the movements of different limbs during specific activities. Our aim is to propose a solution that, in a manner comparable to a clinician, can provide continuous monitoring and alert the appropriate staff if a potential risk is detected. This solution should be applicable both in domestic environments and in clinical settings.

In this chapter, the *Radar Signatures of Human Activities* dataset [57] is employed for all experiments. The reasons for its selection are outlined in Section 3.3.3. The aim here is to perform preliminary tests to evaluate the feasibility and potential of spectrogram classification. Our work begins with the use of a *ResNet-18* model [71]. A method for resizing inputs is implemented, and its impact on the results is assessed.

Building on our initial results, we propose methods to enhance the input images by applying techniques from the field of image processing. The purpose is to investigate whether the quality and content of the data have a significant influence on the accuracy of classification models.

Moreover, particular attention is given to the *optimisation* of the *classification* models themselves. This includes *strategies* such as careful parameter tuning, selection of appropriate pre-processing techniques, and the combination of complementary features. Classical machine learning models, including *SVM* and *KNN*, are optimised and evaluated on specific aspects of the data, such as *geometric* and *textural* features, while deep learning models benefit from optimised input representations and architecture choices. The goal of these optimisation strategies is to improve both the accuracy and the robustness of the models, ensuring reliable performance across different *datasets* and experimental conditions.

Finally, we explore the use of different classification models depending on the data under consideration. As each model may yield a distinct classification output, the objective is to introduce an element of uncertainty in order to enhance the reliability of the final decision. The results are ultimately combined through majority voting. Figure 4.1 presents the structure of Chapter 4, highlighting the sections dedicated to *classification*, *feature extraction*, and *model optimisation*.

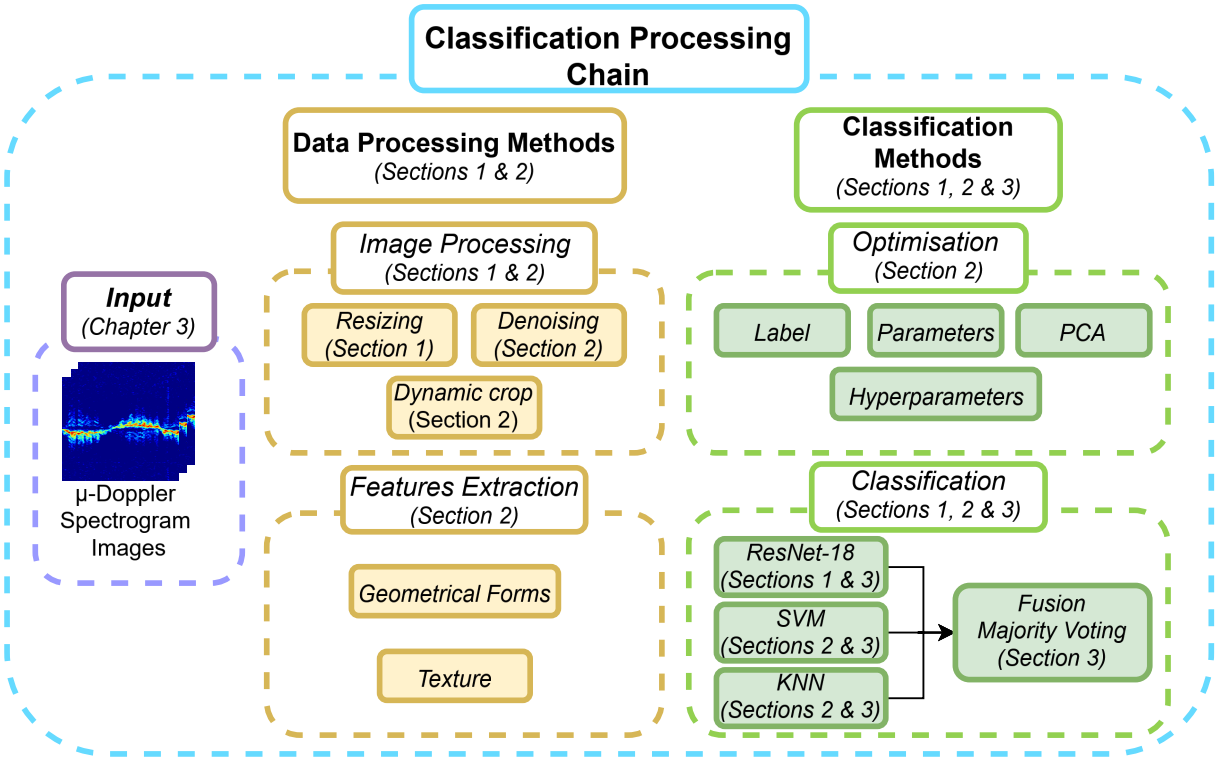

Figure 4.1 – Overview of Chapter 4

4.1 Experimental Evaluation of CNN for Radar-Based Human Activity Recognition

Chapter 3 introduces the *Radar Signatures of Human Activities* dataset [57]. This dataset offers the advantages of employing *FMCW radar* equipment and providing data collected in real-world environments. Although the data are imbalanced, it nevertheless constitutes a solid foundation for conducting initial classification experiments.

After being processed through a pre-processing chain comprising filters and transformations, the *radar* data are represented as *spectrogram* images. Section 3.3.3 emphasises that each recorded activity is characterised by a distinct *Doppler signature*.

The variation in signatures motivates the application of *classification* models to this type of data, *spectrogram*, with the aim of recognising and characterising human activities. This section presents preliminary findings in the field of *activity recognition* based on such data.

4.1.1 Preliminary Evaluation of Classification Architectures

Before defining the baseline model for the *radar*-based classification framework, several *Convolutional Neural Network* (CNN) architectures were reviewed and compared. The selection focused on representative models from different generations of *CNN* design, AlexNet [175, 164], GoogLeNet [175, 164], and *ResNet-18* [135], each offering specific architectural advantages in terms of depth, connectivity, and computational requirements. These differences directly influence the models' ability to extract meaningful features from *radar spectrograms* while maintaining acceptable processing times for practical applications.

Table 4.1 presents a synthesis of their main characteristics and reported performances on *radar*-based *human activity recognition* and *micro-Doppler dataset*s. In order to enable a fair comparison between models of different sizes, a *merit factor* was introduced, defined as the ratio between the achieved *classification* accuracy and the logarithm (base 2) of the number of trainable parameters. This indicator highlights architectures that achieve good performance while remaining computationally efficient, which is particularly relevant for *embedded or real-time radar systems*.

From this comparison, it can be observed that AlexNet-based [175, 164] methods achieve relatively high accuracy (about 78%), but at the expense of a very large num-

Table 4.1 – Classification Models for Radar-Based System

Classification Models	**Year**	**Dataset**	**# Param.**	**Merit Factor** $\left(\frac{\text{Acc}}{\log_2(\#\text{Parameters})}\right)$	**Acc.**
GoogLeNet [175, 164]	2021	*Radar Signatures of Human Activities* [57]	6.70 M	3.29	74.70%
AlexNet + KNN [175, 164]	2021	*Radar Signatures of Human Activities* [57]	60 M	2.99	77.15%
AlexNet + SVM [175, 164]	2021	*Radar Signatures of Human Activities* [57]	60 M	3.03	78.25%
ResNet-18 [135]	2020	Internal Dataset micro-Doppler UAV & HAR [135]	11.70 M	3.71	79.88%

ber of parameters (around 60 million), leading to high computational cost and slower convergence. GoogLeNet [175, 164] significantly reduces the number of parameters (∼6.7 million) but provides no significant accuracy improvement. In contrast, *ResNet-18* [135] exhibits the highest *merit factor*, combining strong *classification* performance (79.88%) with moderate model complexity (∼11.7 million parameters). Moreover, its residual *learning* structure helps prevent gradient degradation during training, improving *optimisation* stability and generalisation.

Based on these quantitative and qualitative considerations, *ResNet-18* was selected as the *reference* architecture for the present study. Section 4.1.2 details its structure, implementation settings, and specific adaptations made for *radar*-based classification, followed by a presentation and discussion of the obtained results.

4.1.2 Implementation and Performance Evaluation of ResNet-18

All subsequent experiments were conducted using the *ResNet-18* architecture [71, 92], which was identified as the most suitable compromise between performance and computational efficiency, as discussed in the previous section. For all classification models, it is essential to employ a dataset that is properly divided into distinct subsets. Typically, one subset is used for *training*, providing labelled examples from which the model learns representative features, while another independent subset is reserved for *testing*, allowing for an unbiased evaluation of the model's generalisation and robustness on previously unseen data.

As an initial validation step, the *ResNet-18* model was trained from scratch and tested using the *CIFAR-10* dataset [91]. This dataset was chosen because it provides a well-established benchmark for image classification and enables verification that both the model implementation and the computational hardware produce reliable and reproducible results. Once this verification was complete, the architecture was retrained from scratch and adapted using the open-source *radar* dataset [57], which was reformatted and pre-processed to serve as input for the *radar*-based classification task described in this study.

a) Methodology and Model Parameterisation

The aim of this study is ultimately to provide a lightweight, portable tool. The classification models under investigation must be compatible with these constraints. The solution intended for clinicians and patients will be embedded, and therefore requires models that are not overly complex, execute efficiently, and are straightforward to maintain.

ResNet-18 represents our initial approach. It is a deep *CNN* introduced in [71]. Its distinguishing feature is the use of residual blocks, which facilitate the training of deep networks. The architecture comprises 18 layers, including convolutional and fully connected layers, is illustrated in Figure 4.2. *ResNet-18* is considered one of the classical *CNN* architectures widely used in the field of deep learning.

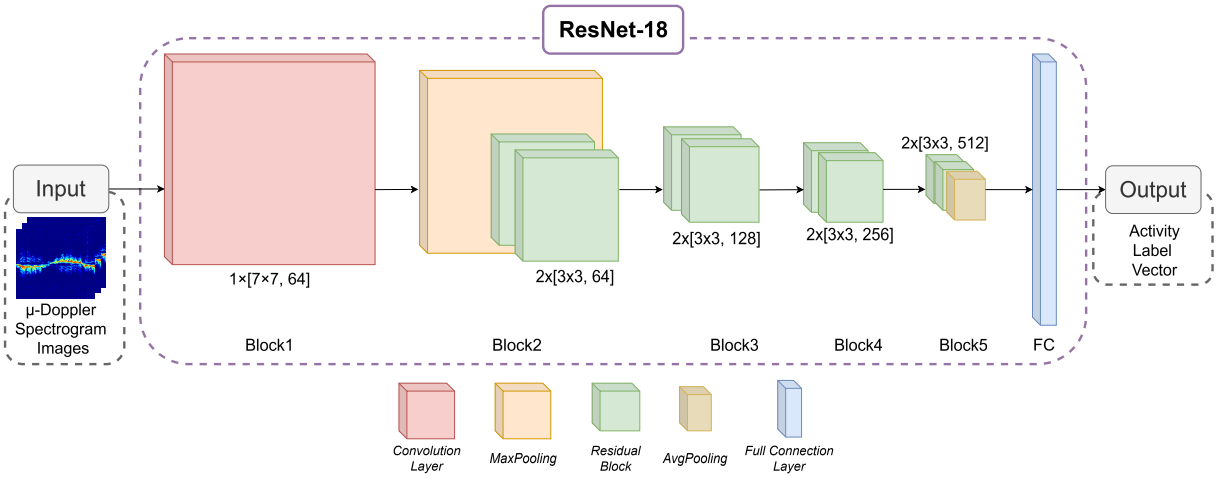

Figure 4.2 – Simplified Diagram of the ResNet-18 Architecture with Residual Blocks. The notation $[k \times k, n]$ indicates a convolutional layer with a kernel size of $k \times k$ and n filters, while the prefix $m\times$ denotes that the block is repeated m times at that stage.

The ResNet-18 model is trained on the target dataset, with its performance monitored

during training and subsequently tested. For the *Radar Signatures of Human Activities* dataset [57], the model is trained from scratch, without transferring weights from CIFAR-10 [91]. The protocols used with the reference *CIFAR-10* dataset [91] and for classifying spectrogram image data are compared in Table 4.2.

Table 4.2 – Comparison of Training and Testing Protocols for CIFAR-10 [91] and *Radar Signatures of Human Activities* [57]. "Train from scratch" indicates that no transfer learning was applied from CIFAR-10 [91].

Step	**CIFAR-10 [91]**	***Radar Signatures of Human Activities*** [57]
1	Implement ResNet-18	Adapt model
2	Train model	Train from scratch model
3	Test model	Test model (trained from scratch)
4	Verify process	Cross-validation
5	-	Test on spectrograms

The model takes RGB images as input, in the form of $(x, y, 3)$ arrays, where the third dimension corresponds to the three colour channels. The model parameters are set as follows: a batch size of 64, a learning rate of 0.01, and training over 20 epochs. The classification model was implemented using the Python library *PyTorch*[1], specifically the *torch*[2] and *torchvision*[3] packages. A 4-fold cross-validation was performed, with a random distribution of data at each iteration, ensuring that 85% of the data were used for training and 15% for validation, with each class proportionally represented.

All training and testing were carried out on two separate computer systems to compare execution times and performance. The first machine used was an Alienware laptop with an AMD Ryzen 9 5900HX processor with Radeon Graphics (3.30 GHz) and 32 GB of RAM. The hardware used on the server comprised 16 CPU cores, 64 GB of RAM, and four GTX 1080 GPUs, each with 11 GB of RAM and 3584 cores.

b) Comparative Results on CIFAR-10

The *CIFAR-10* dataset [91] consists of 50,000 training images and 10,000 test images, for a total of 60,000 images. It contains RGB images of size 32×32 across 10 classes. The classes included in this dataset are: Airplane, Automobile, Bird, Cat, Deer, Dog, Frog, Horse, Ship, and Truck, as illustrated in Figure 4.3.

1. `https://pytorch.org/`
2. `https://docs.pytorch.org/docs/stable/torch.html`
3. `https://docs.pytorch.org/vision/stable/index.html`

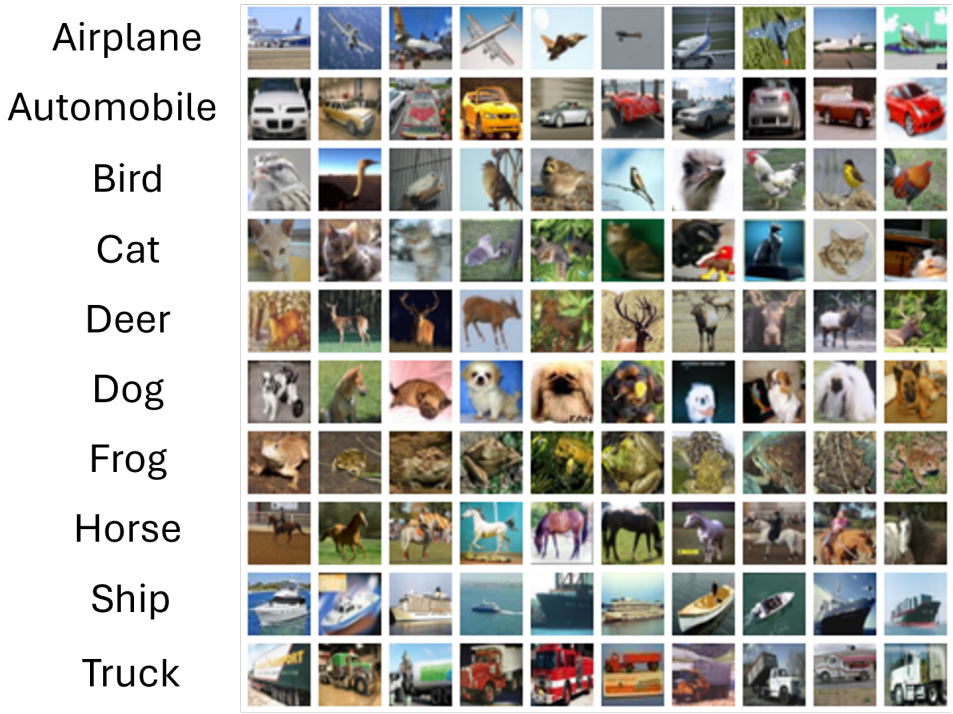

Figure 4.3 – Sample Images from the *CIFAR-10* Dataset [91]

The *CIFAR-10* dataset [91] is used solely for comparison with the original model results. Following adaptation of the model to this type of data, *Radar Signatures of Human Activities* [57] are employed for testing.

Table 4.3 – Accuracy Comparison between State-of-the-Art and our Implmentation on CIFAR-10 [91]

Learning Step	**Reference Accuracy**	**Our Accuracy**
Training	97.76%	97.64%
Testing	51.95%	56.99%

We obtained a training accuracy of 97.64% and a validation accuracy of 56.99%, as shown in Table 4.3. The reference results report a training accuracy of 97.76% and a validation accuracy of 51.95%, where training accuracy reflects the model's ability to recognise data used during training, and validation accuracy indicates performance on previously unseen data, providing a measure of model robustness. The slight differences between the reference *ResNet-18* values and those obtained with our implementation confirm that our model is valid.

While F1-scores provide useful per-class performance information, they are not com-

puted by the original implementation [148], and thus no direct comparison is possible at this stage. The current evaluation therefore focuses on accuracy metrics, which are sufficient to confirm the model's behaviour and allow us to proceed with adapting the micro-Doppler inputs and training on spectrogram images.

c) Model Adaptation and Initial Analyses

Since the *CIFAR-10* [91] images are 32×32 thumbnails, the 550×490 spectrogram images were resized to the same format. Consequently, our results are based on images that have been compressed in terms of size.

Using the data pre-processing programme provided with the *Radar Signatures of Human Activities* [57] dataset, all corresponding spectrogram images were generated. A sample of the resulting images is shown in Figure 3.18, with each class labelled according to the corresponding activity.

A programme was implemented to separate each class into 85% training and 15% testing, enabling the data distribution to be randomly regenerated according to this division criterion. This programme is also useful for cross-validation, allowing the reliability and performance of the model to be evaluated based on the *Radar Signatures of Human Activities* database [57].

Four data distribution schemes were generated. After inputting the spectrogram images and running the training process on each, the corresponding accuracies were recorded, as shown in Table 4.4.

Table 4.4 – Cross-Validation Results

Set number	**Train Accuracy**	**Validation Accuracy**
Set 1	100%	89.45%
Set 2	100%	87.50%
Set 3	100%	91.02%
Set 4	100%	87.11%

As expected, the relatively small size of the dataset allows the model to learn very quickly, achieving 100% training accuracy. However, the validation accuracy averages 88.70%. Given that the data differ substantially from *CIFAR-10* [91], these results are encouraging and satisfactory, considering the performance that can be achieved with directly acquired data.

For the remainder of the tests, we will use "Set 4" from the cross-validation. As it exhibits the lowest validation accuracy, it allows us to investigate the model's limitations and identify which activities are most difficult to recognise for a model of this type.

A test script was implemented following the training process. The validation dataset was input to the model, resulting in an accuracy of 87.11%. This confirms the validity of this implementation stage. Finally, the confusion matrix for the model trained on spectrogram images was generated.

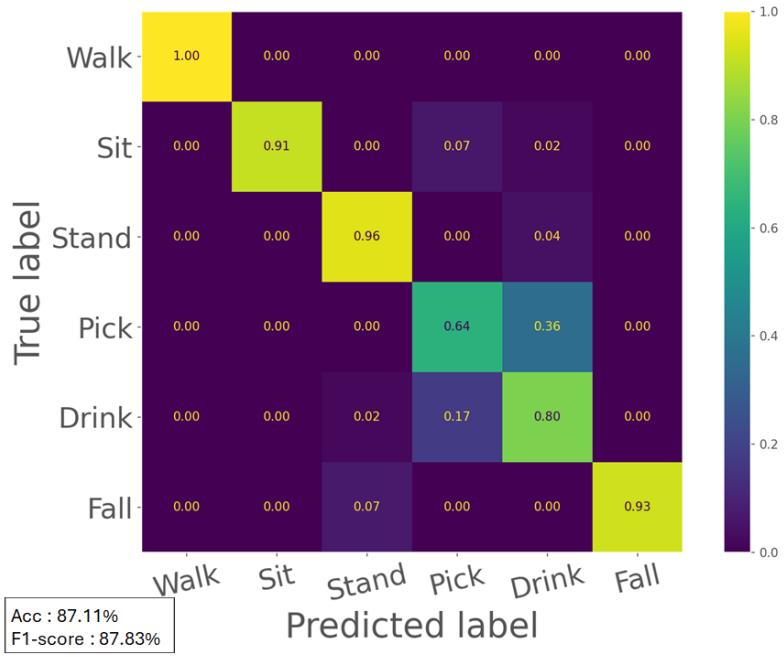

Figure 4.4 – ResNet-18 Normalised Confusion Matrix: 6 Activities

Figure 4.4 shows that, overall, the model has learned effectively and is able to recognise most of the data classes. The classes that present the greatest challenges are: "Sit", "Stand", "Pick", and "Drink".

This is not surprising, as the *spectrograms* for these activities are very similar, as discussed in Section 3.3.3 and illustrated in Figure 3.18. Indeed, the authors of *Radar Signatures of Human Activities* [57] noted that discrepancies sometimes occurred between the activities defined in the protocol and their actual execution. For instance, "Drink" activity could include grasping and placing down the glass, making it similar to "Pick", while "Walk" often involved back-and-forth motion without a clearly defined starting direction.

4.1.3 Impact of Image Resizing

Initial tests have yielded satisfactory results, enabling further study and use of spectrogram data for activity recognition. The intended approach for the classification model input is to use 32×32 thumbnails, as illustrated in Figure 4.5. Despite their small size, the model performs well. This raises the question of how the choice of thumbnail size affects the results. Consequently, the objective is to retrain the model on the spectrogram data from the *Radar Signatures of Human Activities* [57] dataset using varying thumbnail sizes and to analyse the impact of this parameter on model performance.

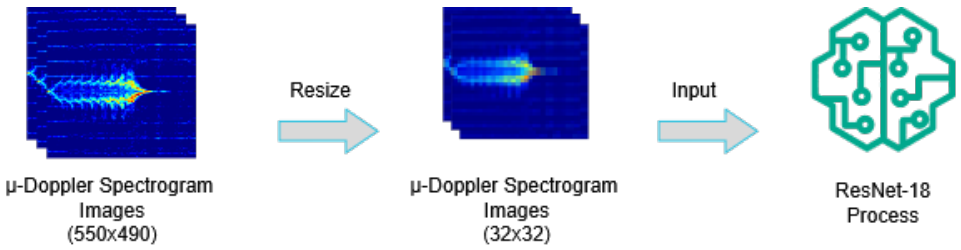

Figure 4.5 – Image Resizing Procedure for ResNet-18 Training

Table 4.5 – Comparison of Accuracy and F1-score Rates Achieved with Thumbnail Resizing

Image Size	**Accuracy**	**F1-score**
32x32	86.72%	87.50%
100x100	**92.97%**	**93.33%**
200x200	92.19%	92.50%
300x300	91.02%	91.17%
400x400	92.58%	93.00%
490x490	90.23%	90.17%
500x500	80.08%	76.67%
550x550	76.56%	70.00%
600x600	85.94%	85.50%

To begin with, the learning programme was duplicated to allow the *ResNet-18* model to be trained from scratch on images resized to different resolutions. The original images have a resolution of 490×550 pixels. For these experiments, the model was trained separately on thumbnails of various sizes, ranging from 32×32 to 600×600 pixels, including intermediate sizes such as 100×100, 200×200, 300×300, 400×400, and the original 490×490 to 550×550.

This procedure does not involve data augmentation: no new images are generated from the originals. Only the input resolution is varied to study the effect of image size on classification performance.

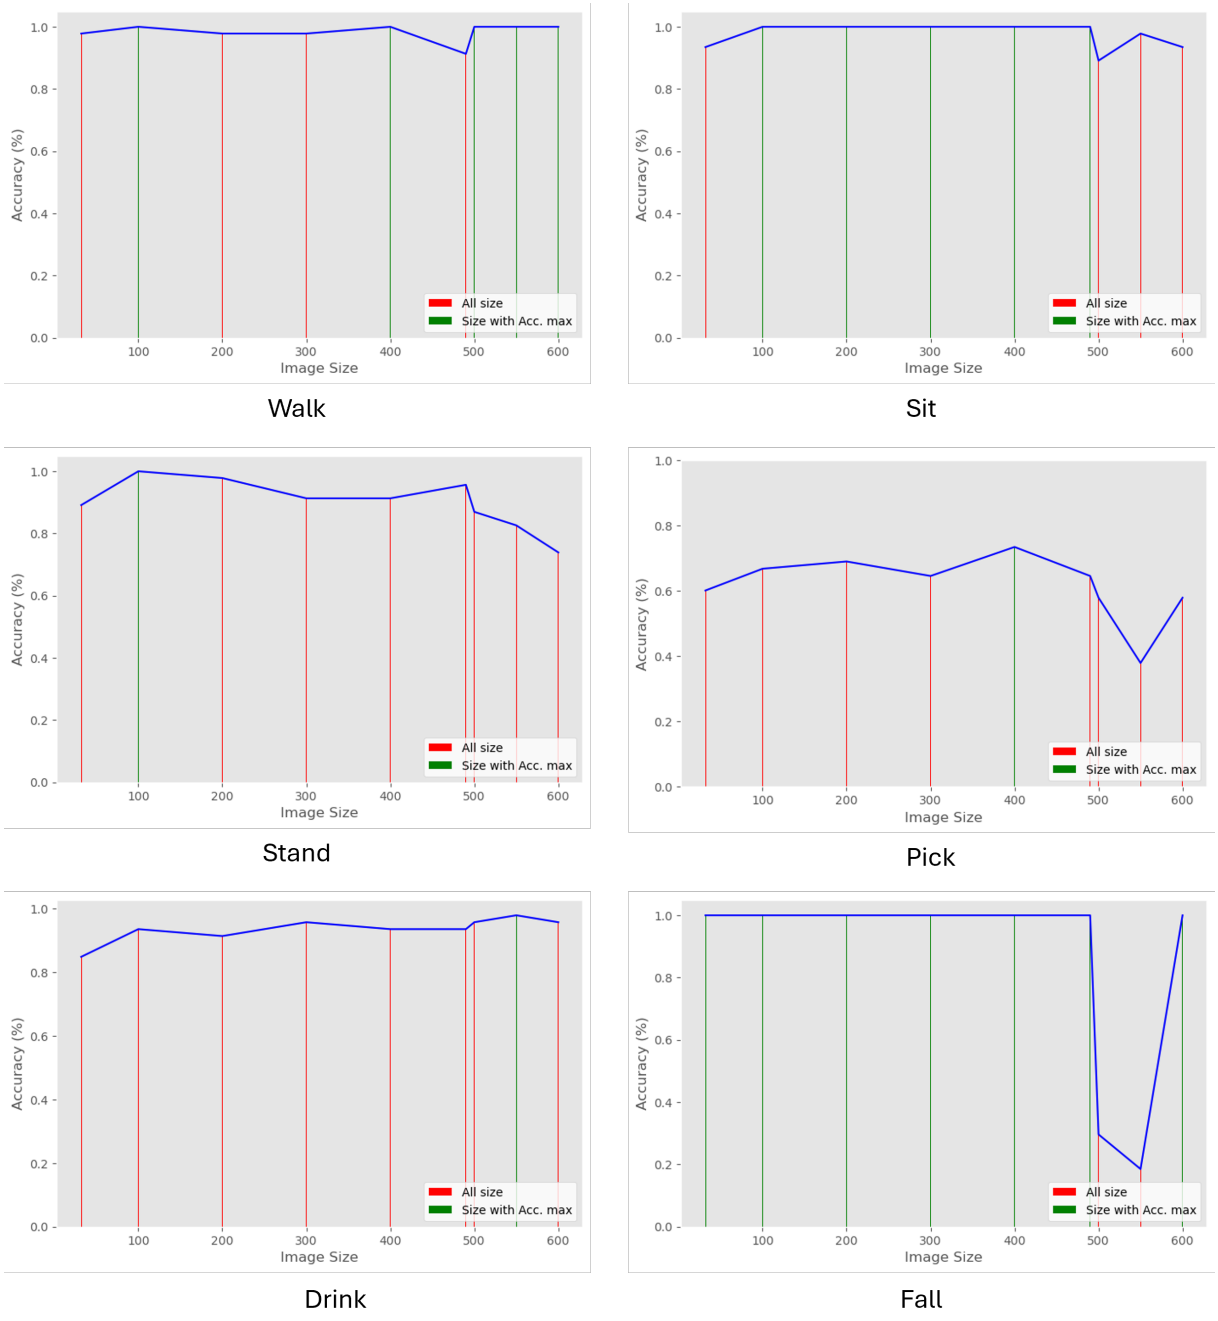

Figure 4.6 – Accuracy per Activity as a Function of Input Size

Table 4.5 summarises the accuracy rates obtained for each size tested. For each of these

selected sizes, a confusion matrix and the corresponding accuracy rates were generated. In most cases, the results remained satisfactory, with accuracies ranging from 76.50% to 93%. However, a slight improvement was observed for thumbnail sizes between 100×100 and 300×300. When the selected size was either too small or too large, the accuracy decreased.

When the accuracy rates are compared, recognition performance varies between activities,as illustrated in Figure 4.6. On average, the 100×100 thumbnail size appears to yield the best results, as shown in Figure 4.7. For "Sit" and "Stand", the confusion matrix indicates that, although these activities are visually similar both on spectrograms and in real life, they are relatively straightforward to distinguish, with only minor confusion observed.

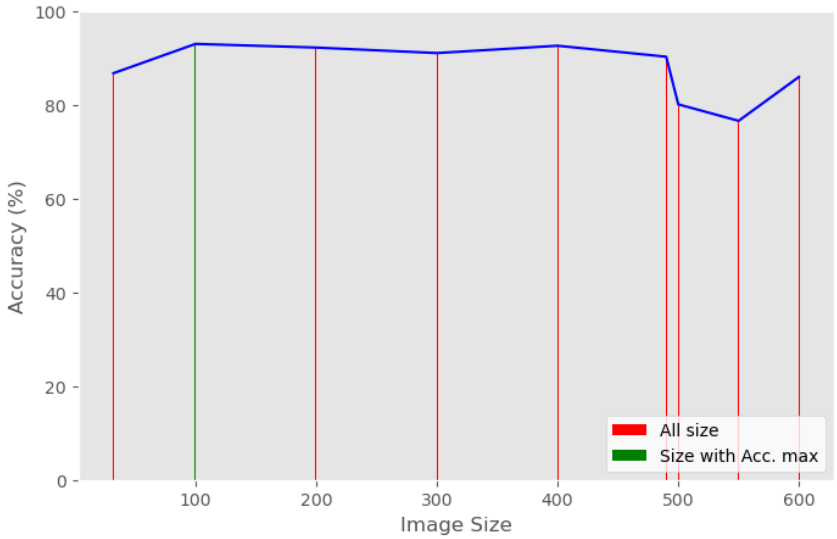

Figure 4.7 – Overall Accuracy as a Function of Input Size

By contrast, the activities "Pick" and "Drink", which are not particularly similar from a human perspective, appear to cause considerable confusion for the model. The authors of *Radar Signatures of Human Activities* [57] confirmed that the "Drink" activity does not consist solely of drinking, but also includes the motion of picking up the glass beforehand. Consequently, for all tests conducted following this Section 4.1, we treat the activities "Pick" and "Drink" as a single combined activity, labelled "Other".

4.1.4 Evaluation on Laboratory Data

The model is capable of classifying and recognising activities from spectrogram images acquired using an *FMCW radar*. To further evaluate its performance with variable data, we tested the model on data outside the *Radar Signatures of Human Activities* dataset [57].

A spectrogram image of a person walking in front of an Ancortek FMCW radar (SDR 980AD2) was acquired and fed into the adapted model, producing prediction scores for a single image, as illustrated in Figure 1.1. These tests were preliminary, using a small number of images (less than 25), and aimed to assess the model's behaviour with data external to the original dataset [57]. A spectrogram image of a person walking in front of the radar was acquired. This image was then fed into the adapted model, which produced prediction scores for a single image, as illustrated in Figure 1.2.

The model recognises the activity performed in each step acquisition test with an average accuracy of 88.77%. This demonstrates that it is possible to identify an activity from radar acquisitions using a classification model trained on spectrogram images.

The working hypothesis proposed earlier was that activities could be recognised following acquisition, after training on spectrogram images. To verify this, we carried out a series of acquisitions involving activities common to the training dataset. Variations of these activities were also introduced in order to assess the robustness of the model at this stage.

Table 4.6 – Summary of Acquired Test Activities

Activity Category	**Description**	**# Samples**
Walk	Conventional walking without imposed constraints	4
Walk	Walking with exaggerated and forced arm movements	4
Walk	Walking diagonally relative to the radar viewpoint	4
Sit	Sitting down on a chair facing the radar	4
Stand	Standing up from a chair facing the radar	4
Pick	Picking up an object from the ground facing the radar	4

Three types of gait were acquired, corresponding to conventional walking, walking with pronounced arm movements, and walking diagonally relative to the radar. These three variations were grouped into a single "Walk" class, in order to match the label used in the *Radar Signatures of Human Activities* dataset [57]. The other activities acquired were: "Sit", "Stand", and "Pick an object" facing the radar. Table 4.6 summarises the four resulting classes.

The confusion matrix shown in Figure 4.8 is based on these acquisitions but evaluated against the six classes used during the initial training.

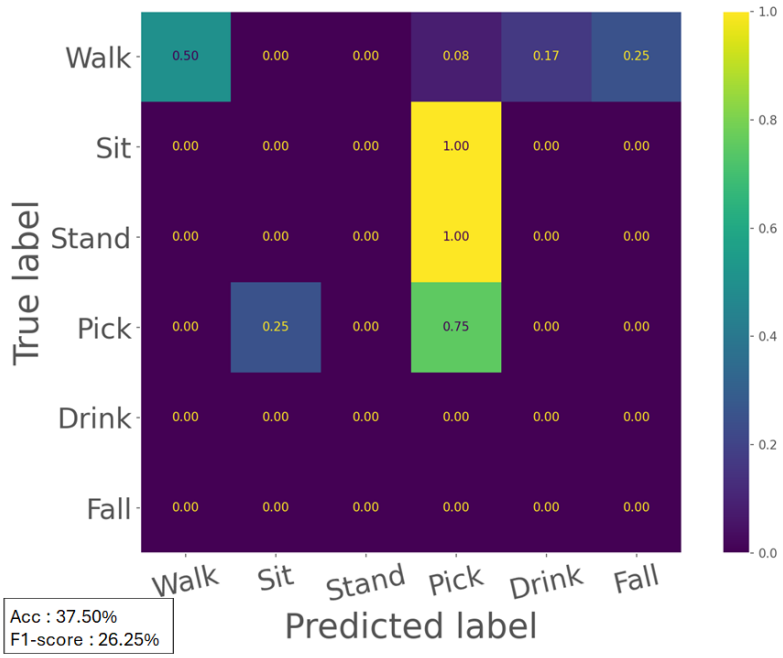

Figure 4.8 – Normalised Confusion Matrix on Four Micro-Doppler Classes with a Model Pre-Trained on Six Different Activities

It can be observed that, although the data acquired differ slightly from the referenced activities, the overall accuracy across all six categories is 37.50%. This relatively low value can be explained by the fact that our collected test data only cover four of the six categories, resulting in 0% accuracy for the two missing categories. Some confusions also occur between "Sit", "Stand", and "Pick" due to differences in the experimental protocols between our data and the original dataset, making it difficult to determine whether the exact actions were performed in the same manner.

Focusing on individual activities, the model achieves 50% accuracy for "Walk". For the "Sit" and "Stand" activities, the model tends to converge strongly towards the "Pick" activity, most likely because these actions were not performed in the same manner as in the open-source dataset protocol used for training. Finally, the "Pick" activity is recognised with an accuracy of 75%. These results therefore support our working hypothesis and reinforce the motivation to pursue our research, with the aim of developing a model that is robust across all types of daily activities acquired using *FMCW radar* data.

4.2 Feature Extraction and Recognition from Spectrograms

Section 4.1 provided initial results using the *ResNet-18* model. These tests highlighted that the data in the *Radar Signatures of Human Activities* dataset [57] were collected under a protocol that lacked sufficient precision and control. In particular, the activity "Drink" is defined very similarly to the activity "Pick".

The promising results obtained raise questions regarding data quality. The spectrograms depend on both the acquisition parameters and the pre-processing chain. Additionally, the quality of the radar itself can significantly influence the data acquired.

This section presents three methodologies that have been implemented to evaluate recognition performance based on radar data. Various image processing techniques are applied to the *Radar Signatures of Human Activities* dataset [57].

4.2.1 Processing and Extraction of Activity from Images

The spectrograms used thus far for classification and activity recognition are largely noisy. A closer analysis of the data revealed that approximately 13% of the images contain defects. In other words, the data either contains information that cannot be interpreted due to radar acquisition errors, or the images are heavily corrupted by noise.

a) Noise Processing in Images

To reduce the amount of noisy data, an approach was implemented to remove noise from the images. A threshold was applied to the green and blue components of each image. The red component corresponds to body parts with the largest surface area and energy density, while the blue component represents regions without useful information. The green component captures finer movements, such as hand or leg motions. Analysis of the per-pixel colour levels revealed that the noise is primarily concentrated in the blue and green channels. This method allows for the removal of a substantial portion of the noise present in most of the images in the *Radar Signatures of Human Activities* dataset [57].

The result can be seen in Figure 4.9. The denoised image still retains the essential information regarding the movements and micro-movements of the person's limbs. Some residual noise remains. However, increasing the threshold further carries a significant risk of losing information necessary for activity analysis. Overall, the chosen threshold appears

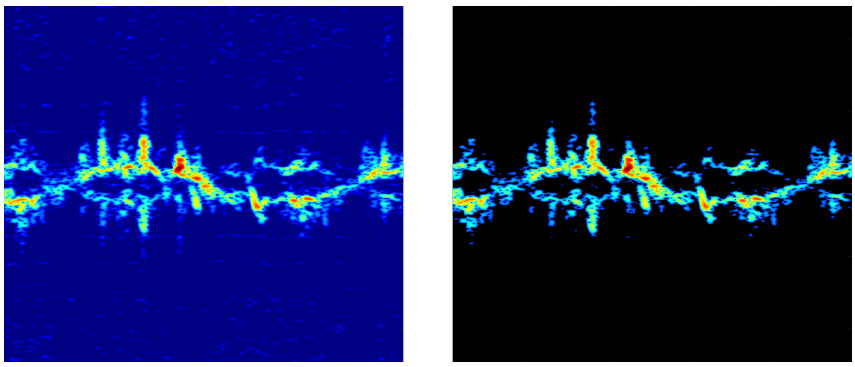

Figure 4.9 – Noise Removal in Spectrogram Data

to provide an adequate balance, sufficiently cleaning the data without compromising critical details.

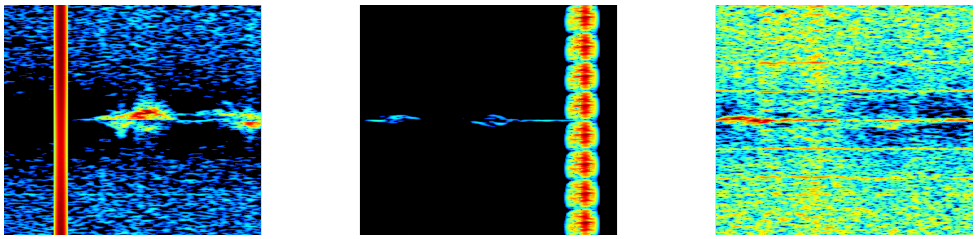

Figure 4.10 – Incorrect Spectrogram Data

Despite this denoising, the noisiest data, whose images contain information that is difficult to interpret, remain present. As illustrated in Figure 4.10, these data either exhibit anomalous signals in the micro-Doppler signature or are so noisy that the activity is almost invisible in the image. Consequently, such information cannot be removed using the simple noise-removal method employed.

b) Activity Zone Recuperation Through Energy Projection

To mitigate these issues, applying a cropping method appears to be a promising approach to improve the recovery of the most recent data, which contains the least noise. Some noise is difficult to remove at the vertical extremes of the image. In most spectrogram images, the activity of interest is concentrated near the centre, which corresponds to the informative region. An initial approach involves projecting the energy contained in the image onto the x-axis. This is achieved by calculating the sum of all pixel values in column j across each row i, using the formula for i ranging from 0 to n: $\sum_{i=0}^{m} p_{ij}$. This allows the identification of the region containing the activity to be preserved.

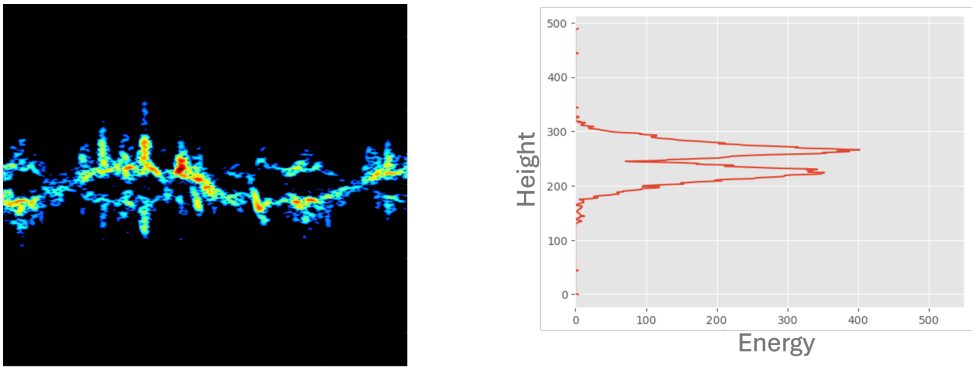

Figure 4.11 – Projection of the Energy of the Spectrogram Image

Using energy projection, it is possible to identify the region most affected by the activity of interest, as illustrated in Figure 4.11. Based on this projection, it appears beneficial to apply a cropping procedure informed by this information.

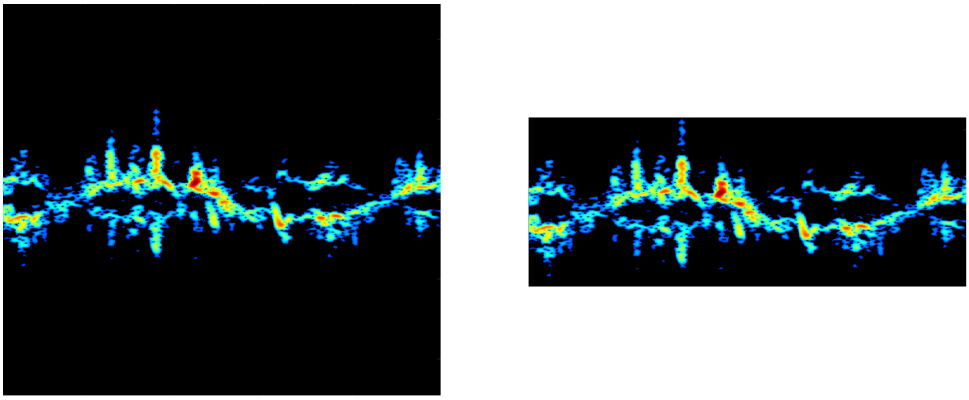

Figure 4.12 – Crop by Energy Projection

Following the energy projection, it is natural to aim to recover the region containing relevant information. To achieve this, the indices of elements with a positive pixel count are identified. In effect, this can be seen as a form of compression of all non-empty regions.

In other words, the maximum and minimum bounds are determined using the values obtained from the projection. If the projected area contains a certain number n of pixels, it is considered to represent an active region. The image can then be cropped to this region, as illustrated in Figure 4.12.

This method performs well with images that have been thoroughly denoised. However, when the data contains excessive noise, the cropping is less effective, as noisy regions are still included despite the initial search. As shown in Figure 4.13, the activity is successfully

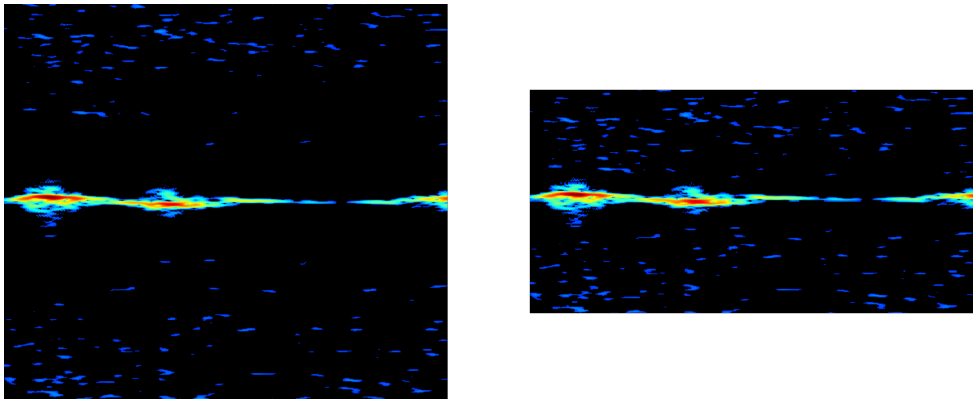

Figure 4.13 – Crop Problem Using Min and Max Method

recovered, but the image also retains noisy information. Consequently, the method requires improvement.

c) Activity Extraction Via Segmentation and Integration of Sub-curves

Simply identifying the minimum and maximum values is insufficient to handle all cases of noisy data. A method was therefore implemented to improve the initial approach. This method involves searching for a maximum in the projection list that is less than 500 pixels, thereby reducing the likelihood of including a noisy region. The index of this maximum is then used to identify the minima on either side, defining the bounds of the activity area to be cropped. Using this approach, the issue illustrated in Figure 4.13 is largely mitigated, resulting in the improved outcome shown in Figure 4.14.

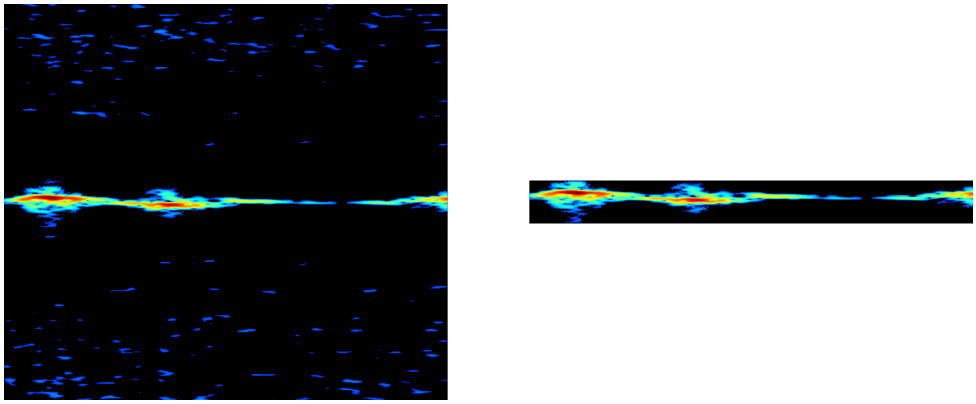

Figure 4.14 – Correction of the Crop Issue Using the Min and Max Method

However, this revised method is still insufficient to address all specific cases of noisy

images. When the image contains excessive scattered noise, the criteria for selecting the activity region are inadequate. As illustrated in Figure 4.15, the minimum number of pixels required to define an activity area is too high to allow extraction of meaningful information.

The final method is then considered to address the remaining complex cases as effectively as possible. Some data will inevitably be lost due to the high noise levels present, but this approach aims to recover as much information as possible from the final images containing scattered noise.

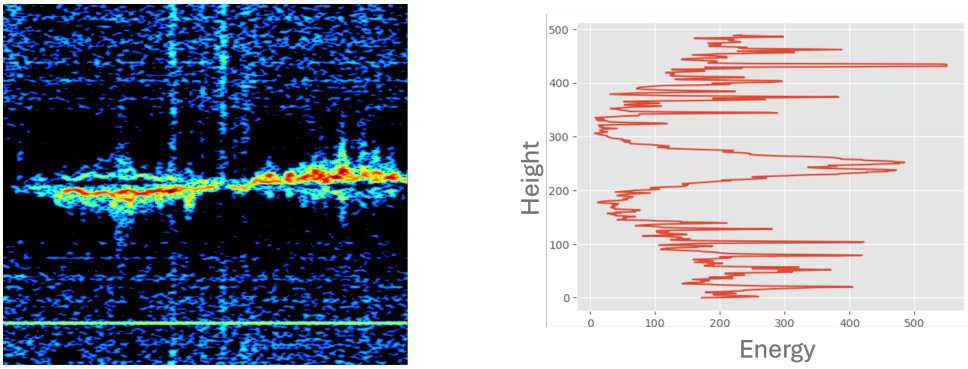

Figure 4.15 – Complex Cases with the Crop Method

The use of minima alone does not appear sufficient to achieve fully accurate results. However, it can be observed that this method generates sub-lists corresponding to distinct regions of information within the image. In other words, the image is effectively partitioned into multiple sub-regions. Regions containing activity are generally larger than those consisting solely of noise.

For this new method, the following procedure is proposed:

1. Two scenarios are considered:
   (a) If the noise level is low, the initial list is partitioned according to its minima to obtain n sub-lists.
   (b) If the partitioning produces only a few sub-lists due to excessive noise, the list is re-divided using a stronger minimum, defined by a ratio threshold. Specifically, the new minimum is calculated as $min + (percentage * min)$, and the list is partitioned accordingly.
2. The area under the curve of each sub-list is computed.
3. The sub-list with the largest area (largest integral) is selected.

4. The maximum of the selected sub-list is determined:
   (a) In very noisy cases, minima on either side are identified to recover information potentially lost due to the threshold adjustment.
   (b) Otherwise, the size of the selected sub-list is recorded.
5. The resulting area is saved.

More concretely, we calculate the value of the integral for each area, and we assume that the largest value corresponds to the area of activity, as illustrated in Figure 4.16.

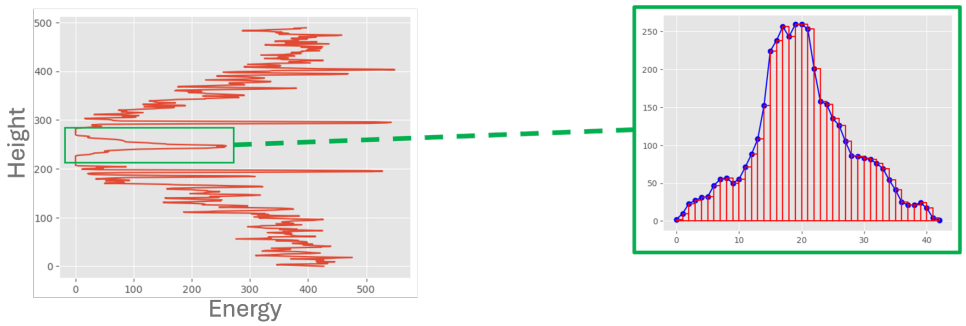

Figure 4.16 – Representation of Integral Calculation Under Curve

Our final method appears to best address the original problem of identifying the area where the activity is located in the image, while aiming to minimise the inclusion of noise. However, a few special cases remain, such as overcropping. Overcropping occurs when the activity area is partially cut, resulting in cases where the relevant information is present only in the positive velocity region, or vice versa. An improvement to this third method is proposed by introducing the following condition in step 3:

3. We select the sublist with the largest area (i.e., the largest integral) whose start index is closest to the centre of the image.

Adding this constraint corrects a large portion of the data, but it is not sufficient. Finally, it is proposed to incorporate a *symmetry method* relative to the centre of the image, with the addition of a margin of variable size. Thus, if only positive velocities relative to the centre are initially recovered, applying symmetry with respect to this selected area allows the corresponding negative velocities to be recovered as well. This results in more complete data. Verification via image subtraction confirms that the noise has been successfully removed from the image.

By applying this final method of dynamic cropping with symmetry and margin, we analysed the data in three distinct ways. In the first case, we consider the data from the

method without dynamic cropping, where all images contain as little noise as possible while preserving complete activity. In the second case, which incorporates dynamic cropping, any image containing activity is regarded as non-lost data. Finally, in the third case with dynamic cropping, we recover all images that are entirely free of noise. The results are summarised in Table 4.7.

Table 4.7 – Comparison of Loss Rates Between the Last Two Crop Methods

Data type	**Method**	**Loss rate**
Train	First crop	27.89%
Test	First crop	29.29%
Total	**First crop**	**28.09%**
Train	Dynamic crop if activity	0.67%
Test	Dynamic crop if activity	0.39%
Total	**Dynamic crop if activity**	**0.63%**
Train	Dynamic crop if no noise	13.37%
Test	Dynamic crop if no noise	9.37%
Total	**Dynamic crop if no noise**	**12.79%**

Finally, we note that the latest cropping method appears to allow for maximum data recovery, based on selection parameters that also account for data quality.

4.2.2 Spectrogram-Based Geometric Form Recognition

Dynamic crop allowed us to investigate methods for enhancing the quality of spectrogram images. Our focus was on the application of mathematical morphology and multimodal approaches. At present, however, these processes are too complex to provide a lightweight and easily deployable solution. Nonetheless, we consider that such processing methods could be valuable for extracting parameters that are specifically associated with the information-bearing region of the spectrogram.

ResNet-18 operates on the entire image, which we resized without observing any notable impact on performance. An alternative approach could involve extracting only the activity-related information in order to capture characteristics specific to this content. Our results indicate that each activity exhibits a distinct signature, thereby forming its own spectrogram image. This leads us to consider whether a method based on geometric shape extraction might provide a viable solution within the scope of our study.

a) Extraction of Geometrical Forms from Spectrograms

Once the spectrograms had been generated, only the activity signature information was retained by means of image binarisation. For this new approach, we propose the use of a lighter model, specifically a *Support Vector Machine* (SVM). To optimise information storage and reduce computational complexity, shape-based features are regarded as more appropriate.

The extracted features comprise *Area*, *Perimeter*, *Orientation*, *Major* axis length, *Minor* axis length, *Centroid* coordinates (X and Y), and *Eccentricity*. These were computed using mathematical methods and the Python library *skimage*[4]. All extracted data, illustrated in Figure 4.17b (derived from Figure 4.17a), were stored in a text file for subsequent processing.

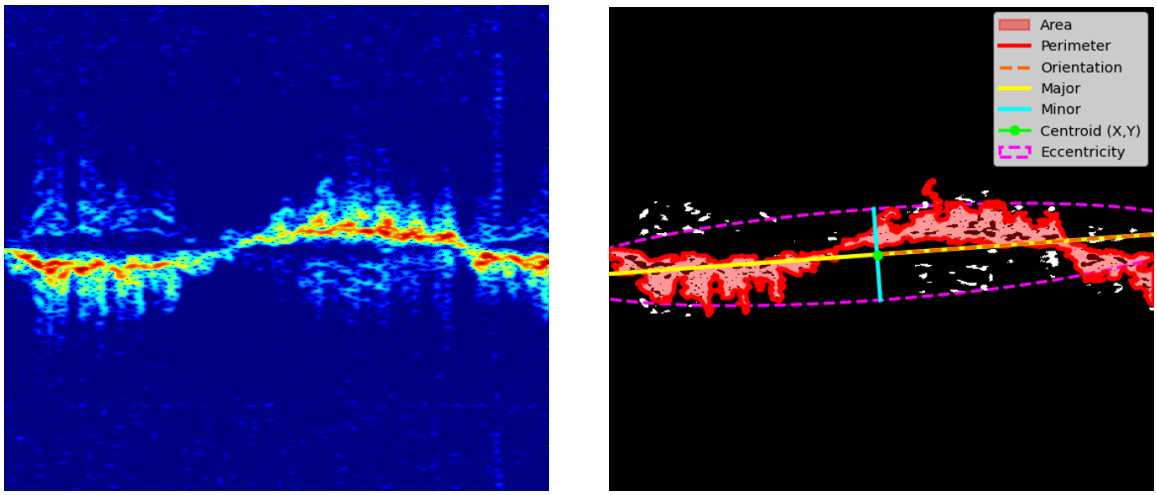

(a) Spectrogram of "Walk" [57] (b) Parameters Extraction

Figure 4.17 – Spectrogram Analysis, from the Original to the Parameters Extraction

To extract the eight previously mentioned shape features, the activity signatures in the binarised spectrograms were first isolated by identifying connected regions. Morphological and geometrical descriptors were subsequently computed using the *regionprops* function from the *skimage* library.

Contours were extracted using the $measure.find_contours$ function, with the threshold set to 0.9 and the $fully_connected = "high"$ option enabled to account for diagonal connectivity. To ensure geometric consistency, contours initially traced in a clockwise direction were reversed to adopt a counterclockwise orientation, in accordance with standard

4. `https://scikit-image.org/`

mathematical conventions.

The following parameters are extracted for each detected region:

- **Area**: Number of pixels contained within the region;
- **Perimeter**: Sum of Euclidean distances between consecutive contour points (x_j, y_j), where $j = 1 \ldots n$, and n is the total number of points on the contour, including the closing segment connecting the last point (x_n, y_n) and the first one (x_1, y_1) (see Equation 4.1);

$$P = \sum_{j=1}^{n-1} \sqrt{(x_{j+1} - x_j)^2 + (y_{j+1} - y_j)^2} + \sqrt{(x_1 - x_n)^2 + (y_1 - y_n)^2} \tag{4.1}$$

- **Orientation**: Angle (in radians) between the region's major axis and the horizontal axis. The orientation θ is converted to degrees and adjusted to conform to the digital image coordinate system, in which the y-axis increases downward. The transformation is described in Equation 4.2;

$$\theta = \begin{cases} \theta + 90, & \text{if } \theta < 0 \\ \theta - 90, & \text{else} \end{cases} \tag{4.2}$$

- **Major and Minor axis lengths**: Dimensions of the equivalent ellipse fitted to the region;
- **Centroid** (X, Y): Coordinates of the region's centre of mass;
- **Eccentricity**: Measure of the shape's elongation, ranging from 0 (perfect circle) to 1 (highly elongated ellipse).

For visualisation purposes (as shown in Figure 4.17b), the orientation of each region is represented by a directional line originating from the centroid and extending along the corrected angle. The endpoint of this line is calculated using half of the major axis length L and the corrected orientation θ, as defined in Equation 4.3.

$$(x_{\text{end}}, y_{\text{end}}) = (x + L \cdot \cos(\theta),\ y - L \cdot \sin(\theta)) \tag{4.3}$$

The extracted parameters are then stored for subsequent feature selection and the training of classification models.

b) SVM Method for Geometric Form Classification

The study of *ResNet-18* in Section 4.1 has demonstrated the feasibility of recognising activities using a simple and easily implementable method. We are now exploring an alternative approach for interpreting the activity signatures, while still aiming to maintain a lightweight and straightforward solution.

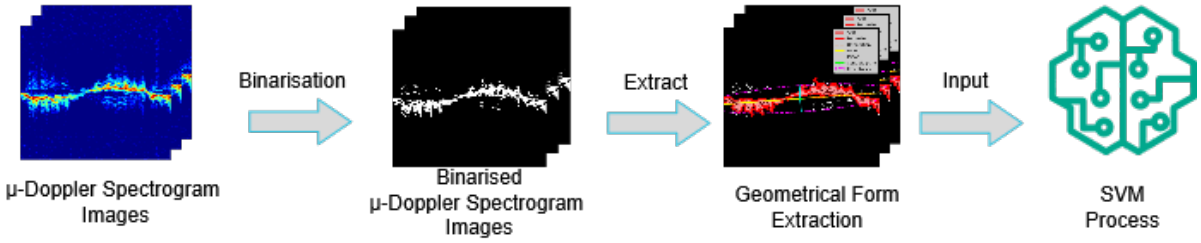

Figure 4.18 – Geometrical Form Extraction for SVM Training

A *SVM* is proposed for learning from spectrograms after the relevant features have been extracted, as illustrated in Figure 4.18. The SVM represents a lightweight, fast, and highly effective solution, particularly well suited to small datasets. In addition, it allows for efficient storage of minimal information. The *SVM model* from the Python library *scikit-learn*[5].

The training procedures were applied to the data from the *Radar Signatures of Human Activities* dataset [57]. The training and testing datasets are identical to those used in Section 4.1 to ensure consistency when comparing our results with this method. Consequently, the same train and test sets obtained through cross-validation for *ResNet-18* were employed.

A total of 1,753 data samples were used, distributed as described in Table 4.8.

For training, the model is provided with a set of parameters extracted from the spectrograms and stored in text files. The model hyperparameters, as well as the extracted features, can be optimised according to their relevance to the data and the specific requirements of this study. A more comprehensive analysis is necessary to select them appropriately.

5. `https://scikit-learn.org/stable/`

Table 4.8 – Training Set Distribution Details

Activity Category	**Training Samples**	**Testing Samples**
Walk	264	47
Sit	265	47
Stand	264	47
Other	527	94
Fall	168	30

c) Method for Selecting SVM Hyperparameters and Parameters

Several types of mathematical representations exist to best suit the specific problem posed by the data to be classified. In the context of human activity recognition from radar data, the observations are real-valued and the problem is inherently nonlinear. Among the various *SVM* variants, four kernel types are commonly used: *Linear*, *Radial Basis Function* (RBF), *Sigmoid*, and *Polynomial* [104, 110, 194]. *Linear* and *Sigmoid* kernels were excluded due to the nonlinear nature of the data. The *RBF* and *Polynomial* kernels, which performed well in initial tests, were compared to determine the most suitable kernel and hyperparameters for human activity recognition.

Preliminary tests indicated that both the *RBF* and *Polynomial* kernels can achieve satisfactory performance, depending on the selected parameters. Consequently, these two kernels were systematically evaluated to ensure effective and optimal results. The use of SVMs and their kernels also requires the specification of certain hyperparameters.

- **For a RBF Kernel:**
  - *Error Penalty C:* corresponds to the penalty associated with classification errors. A high C forces the model to correctly classify most training points, which can lead to overfitting, characterised by a narrow decision margin and lower tolerance to errors. Conversely, a low C widens the margin and increases tolerance to errors.
  - *Range of Influence of the Point γ (gamma):* controls the influence of a single training point on the decision boundary. A high value of γ focuses learning on local details, increasing the risk of overfitting, while a low value has a more global effect, potentially resulting in underfitting.
- **For a Polynomial Kernel:**

* *Error Penalty C*
* *Model Complexity degree:* A higher degree allows modelling more complex decision boundaries, but also increases the risk of overfitting.
* *Importance of lower degree terms, coef0:* adjusts the relative importance of lower-degree terms compared to higher-degree terms, playing a crucial role when $degree > 1$.

Finally, *Principal Component Analysis* (PCA) is applied to reduce data dimensionality, capture the key variations, and compare results with and without this technique. The aim is to optimise the model parameters and hyperparameters by selecting the most relevant features from the spectrogram signatures.

4.2.3 Spectrogram-Based Textural Feature Recognition

The extraction of geometric forms opens up the possibility of exploring methods that are more applicable to the features that can be derived from spectrograms. In image processing, a range of techniques is available for extracting information from an image. Contour detection, element labelling, and filtering are among the methods that can be employed in different contexts.

Among these techniques is the study of image texture [90, 17, 68]. This concept is very complex and poses difficulties in the field of texture classification. There are many different types of texture, which makes it difficult to categorise them. They are generally divided into three main categories:

- **Structural textures (macro-textures):** characterised by the repetition of a pattern, known as a *texton*, following specific rules of direction and arrangement. Examples include brick walls, certain fabrics, and wire mesh.
- **Random textures (micro-textures):** lacking any regular or localisable pattern. Visually, these textures tend to appear homogeneous and often represent microscopic structures. Examples include the surface of water, fields of grass, wood, and similar materials.
- **Directional textures:** lacking a distinctive structural pattern but not entirely random, these textures exhibit an oriented structure.

Numerous methods [90, 17] have been proposed for extracting and identifying elements from image textures. These include spatio-frequency approaches, which rely on filters such

as *Fourier*, *Gabor*, and *discrete wavelets* to perform global texture analysis inspired by mechanisms in the brain. Model-based methods, such as those using *Markov* models, characterise image textures by estimating parameters in order to construct a texture-generating model. Finally, *statistical methods* do not attempt to identify a specific texture pattern within the image; instead, they evaluate the statistical distribution of grey levels to derive texture properties.

For texture analysis, the second-order approach based on the *Gray Level Co-occurrence Matrix* (GLCM) is among the best-known methods. It remains highly popular due to its strong performance. The purpose of this Section is to present a method for extracting texture features from spectrograms for subsequent application in machine learning.

a) Extraction of Textural Feature from Spectrograms

As with geometric form extraction, a *texture feature extraction* method is applied once the spectrogram images have been obtained. To this end, the RGB images are first converted into greyscale, as shown in Figure 4.19, after which the co-occurrence matrix is computed from the transformed image.

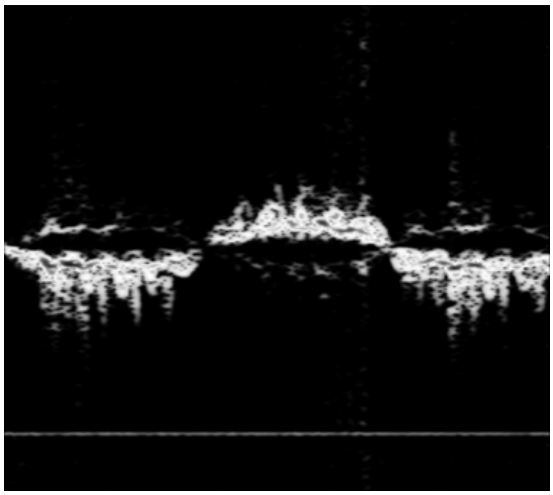

Figure 4.19 – Grayscale Spectrogram

For this approach, we propose the use of a lightweight classification model, *KNN*. Given that the *Radar Signatures of Human Activities* database [57] is relatively small, KNN is particularly well suited to texture descriptors, as it makes no strict assumptions about data distribution and directly exploits the distances between feature vectors.

We propose the use of Haralick's parameters [68]. Haralick introduced the use of co-occurrence matrices to extract image texture features. For a greyscale image with 256 intensity levels, the co-occurrence matrix is a 256×256 matrix in which the element at row i and column j represents the number of times a pixel of intensity i occurs at a distance D and direction θ from a pixel of intensity j.

As the size of the matrix can become very large, leading to long computation times, quantisation is generally applied before calculating the *GLCM*. This matrix encodes a substantial amount of information regarding the orientation and distance of elements and is often symmetrical. However, such data are difficult to exploit directly. Haralick [68] therefore proposed extracting 14 attributes that describe the physical characteristics of texture, which are summarised in Table 4.9. In this table, $P(i, j)$ denotes the (i, j)-th entry of the normalized *GLCM*, representing the probability that a pixel with intensity i occurs adjacent to a pixel with intensity j. $P_x(i)$ and $P_y(j)$ are the corresponding marginal probabilities along the rows and columns.

Texture features were initially extracted in MATLAB and later migrated to Python to ensure full integration within the image processing pipeline. The implementation relies on the *graycomatrix*[6] and *graycoprops*[7] functions from the *skimage.feature*[8] library. During validation, minor discrepancies were observed between MATLAB and Python outputs, primarily due to differences in quantisation, normalisation, and matrix indexing. Although these variations do not affect overall feature trends, a detailed comparative analysis and optimisation of the Python implementation are planned as part of future work.

Attributes f_6 to f_{14} provide complementary information on homogeneity, image complexity, and correlation. The features most relevant for classification are *Angular Second Moment* (ASM), *Dissimilarity*, *Correlation*, *Contrast*, *Homogeneity*, and *Energy*. All extracted data are stored in lists, and the resulting distributions are illustrated in Figure 4.20.

Finally, the data extraction process involves first converting the spectrogram image to greyscale. The co-occurrence matrix is then computed and normalised using *graycomatrix*. Subsequently, the various texture parameters are extracted with *graycoprops*, enabling the retrieval of information from the *GLCM*.

6. `https://scikit-image.org/docs/0.25.x/api/skimage.feature.html#skimage.feature.graycomatrix`
7. `https://scikit-image.org/docs/0.25.x/api/skimage.feature.html#skimage.feature.graycoprops`
8. `https://scikit-image.org/docs/0.25.x/api/skimage.feature.html`

Table 4.9 – Description of Haralick's 14 Parameters [17, 68]

Features	Description	Formula
Angular Second Moment (ASM)	Energy quantifies the homogeneity of an image. Lower values indicate less uniformity and greater colour variation	$f_1 = \sum_i \sum_j P(i,j)^2$
Contrast	Inertia measures local colour variations. When these variations are pronounced, the contrast is high. This parameter is largely uncorrelated with energy	$f_2 = \sum_{i=1}^{n} \sum_{j=1}^{m} (i-j)^2 P(i,j)$
Correlation	Where μ_x, μ_y, σ_x and σ_y are, respectively, the means and standard deviations of px and py. This parameter enables us to determine whether certain columns of the matrix are identical, i.e., whether there are linear dependencies within the image	$f_3 = \frac{\sum_{i=1}^{N_g} \sum_{j=1}^{N_g} ij\, P(i,j) - \mu_x \mu_y}{\sigma_x \sigma_y}$
Sum of Square (Variance)	Variance measures the dispersion of values around the mean. Higher values indicate greater differences between the values and the mean, reflecting the heterogeneity of the texture	$f_4 = \sum_i \sum_j (i-\mu)^2 P(i,j)$
Homogeneity	This parameter measures the homogeneity of the image; the more homogeneous regions it contains, the higher the value	$f_5 = \sum_i \sum_j \frac{P(i,j)}{1+(i-j)^2}$
Sum Average	-	$f_6 = \sum_i i\, P_{x+y}(i)$
Sum Variance	-	$f_7 = \sum_i (i - f_8)^2 P_{x+y}(i)$
Sum Entropy	-	$f_8 = -\sum_{i=1}^{2N_g} P_{x+y}(i) \log(P_{x+y}(i))$
Entropy	This parameter measures the disorder in the image. Its value is high for completely random textures, lacking any apparent structure. Entropy is elevated when the values of the co-occurrence matrix are nearly all equal	$f_9 = -\sum_{i=1}^{N_g} \sum_{j=1}^{N_g} P(i,j) \log(P(i,j))$
Difference Variance	-	$f_{10} = \sum_i (\mu - \mu_{x-y})^2 P_{x-y}(i)$
Difference Entropy	-	$f_{11} = \sum_i P_{x-y}(i) \log P_{x-y}(i)$
Information Measures of Correlation 1	Let P_x, P_y be the probabilities according to x and y, H_X, H_Y their respective entropies $H_{XY} = -\sum_{i=1}^{N_g} \sum_{j=1}^{N_g} P(i,j) \log P(i,j)$ with $H_{XY1} = -\sum_{i=1}^{N_g} \sum_{j=1}^{N_g} P(i,j) \log\left(P_x(i)P_y(j)\right)$	$f_{12} = \frac{H_{XY} - H_{XY1}}{\max(H_X, H_Y)}$
Information Measures of Correlation 2	$H_{XY2} = -\sum_{i=1}^{N_g} \sum_{j=1}^{N_g} P_x(i)P_y(j) \log\left(P_x(i)P_y(j)\right)$	$f_{13} = \sqrt{1 - e^{-2(H_{XY2} - H_{XY})}}$
Maximal Correlation Coefficent	Let Q(i,j) be defined by: $Q(i,j) = \sum_k \frac{P(i,k)\,P(j,k)}{P_x(i)P_y(k)}$ and V be the second largest eigenvalue of Q	$f_{14} = \sqrt{V}$

b) KNN Method for Textural Feature

For classification, we propose using a *KNN* model, which is a non-parametric algorithm. Distances between Haralick [68] feature vectors are computed, making the method well suited to multidimensional data. Since the Haralick [68] parameters form a multidi-

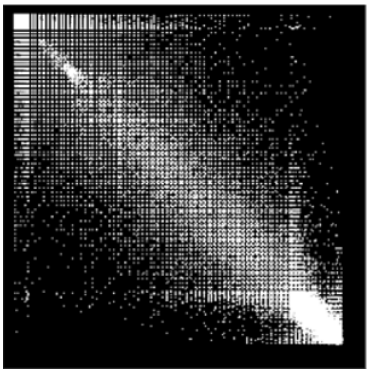

Figure 4.20 – Gray Level Coocurence Matrix of Spectrogram: Walk Activity

mensional feature vector, *KNN* can effectively exploit the relationships between features for classification.

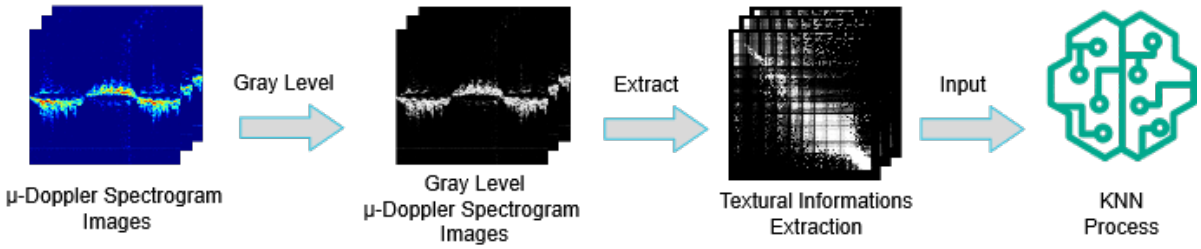

Figure 4.21 – Haralick Parameters Extraction for KNN Training

KNN naturally handles these vector spaces through a distance metric. This model is particularly useful because Haralick [68] parameters can exhibit complex and correlated distributions. KNN can create complex decision boundaries, which is advantageous since image textures often produce classes that are difficult to separate due to their non-linear nature. By adjusting the parameter K, the model can be easily tuned, and it is robust even with limited training data, as is the case with the dataset *Radar signatures of human activities* [57]. For implementation, we employ the *KNeighborsClassifier* from the Python library *sklearn.neighbors*[9].

The learning protocol follows the same procedure as for *SVM* and *ResNet-18*, as illustrated in Figure 4.21. Specifically, data from the *Radar Signatures of Human Activities* dataset [57] are used as input to the model, with the samples categorised into five activities, as summarised in Table 4.8.

As with *SVM*, we aim to optimise the parameters used as inputs to the model. The

9. `https://scikit-learn.org/stable/api/sklearn.neighbors.html`

hyperparameters of the classification model will also be examined in order to achieve the best possible optimisation for our specific requirements.

c) Method for Selecting KNN Hyperparameters and Parameters

The *KNN* model has three hyperparameters that must be optimised to enhance its learning performance [142, 75, 155, 5, 162]. The factors influencing *KNN* are the number of neighbours considered, the distance metric employed, and the method used to weight the neighbours. Optimising these three criteria allows for improved classification results. These hyperparameters can therefore be summarised as follows:

- **Number of neighbours K:**
  - *Small K (e.g. 1–3)*: the model becomes sensitive to noise, leading to high variance.
  - *Large K (e.g. 15–40)*: the model is more stable, but there is a risk of excessive smoothing of class boundaries, which can result in significant bias.
- **Distance metric:**
  - *Euclidean*: performs well if the features are normalised.
  - *Manhattan (L1)*: can be more robust when certain dimensions dominate.
  - *Minkowski*: generalises both Euclidean and Manhattan metrics.
  - *Cosine*: useful when the focus is on vector direction rather than magnitude.
- **Neighbour weighting:**
  - *Uniform*: each neighbour contributes equally.
  - *Distance-weighted*: closer neighbours have greater influence; this weighting is often more effective.

For the Haralick texture feature parameters [68], we use the same method as for *SVM* explained in Section 4.2.2. That is, we will use *PCA* to reduce and optimise the inputs considered for learning by the *KNN* model.

4.3 Optimisation and Evaluation of Learning with Diverse Features

Section 4.2 presents three classification methods, each based on a different spectrogram processing approach. The objective is to compare the accuracy achieved by these various

methods.

We first evaluate the effectiveness of our data denoising method with *ResNet-18*. Subsequently, we analyse the performance of *SVM* and *KNN* after optimising their respective hyperparameters and input features.

Finally, we propose a method for combining the predictions from each model. In particular, we investigate whether such an ensemble approach can improve the reliability of the final prediction results.

4.3.1 Impact of Data and Pre-Processing on Learning

Section 4.1 highlighted the potential for learning from spectrograms. While the results are already highly satisfactory, they exhibit a degree of uncertainty that could increase the risk of false alarms. We hypothesised that the definition of activities may influence the model's ability to discriminate effectively.

This observation prompted further analysis of the data. Section 4.2.1 demonstrated that the *Radar Signatures of Human Activities* dataset [57] contains erroneous and noisy data. Various denoising methods were proposed, and in this section, we evaluate them to determine whether they provide a tangible improvement in training accuracy.

a) Impact of Labels on Performance

The results and configurations obtained during our preliminary tests prompted a more detailed analysis of the data and its impact on the performance of the *ResNet-18* classification model. The objective here is first to compare the performance gain achieved by reducing the number of classes from six to five.

The data used consist of the spectrograms obtained from the pre-processing chain proposed by the *Radar Signatures of Human Activities* dataset [57], producing data as illustrated in Figure 4.17a.

After adjusting the number of classes input to the *ResNet-18* model described in Section 4.1, the model was trained from scratch.

An accuracy of 97.26% was achieved, representing an improvement of approximately 9–10% following the grouping of the "Pick" and "Drink" activities. The confusion matrix shown in Figure 4.22 confirms that this accuracy is consistent across activities rather than being isolated to specific cases. The diagonal remains prominent, with only the "Stand" activity exhibiting slight confusion. These results underscore the importance of activity

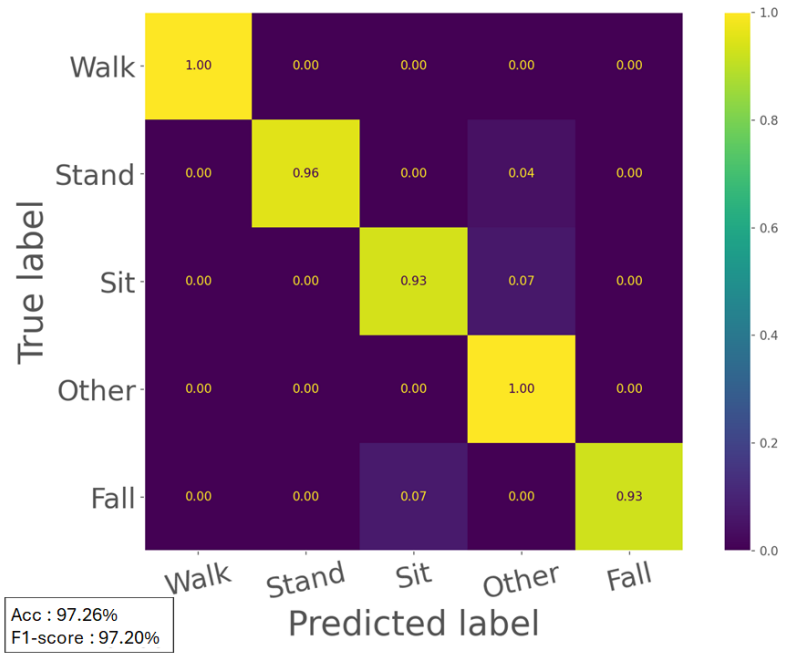

Figure 4.22 – ResNet-18 Normalised Confusion Matrix: 5 Activities

definitions for recognition based on spectrograms.

b) Denoising Impact on Learning

The first method involves denoising the images by directly manipulating the histogram of the three colour channels in the RGB images. The process is described in greater detail in Section 4.2.1.

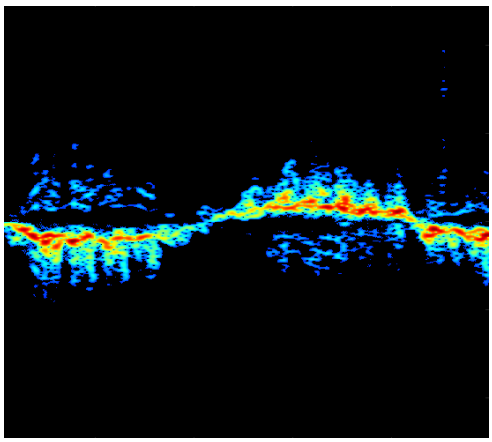

Figure 4.23 – Denoise Spectrogram Example

After denoising the spectrograms, the resulting images have a black background. Most of the artefacts and residual noise are no longer visible, as illustrated in Figure 4.23.

This processed data was then input into the *ResNet-18* classification model, yielding a recognition accuracy of 92.96%. This result is approximately 3% lower than the accuracy obtained with the unprocessed images.

The confusion matrix shows slightly increased confusion between the "Sit" and "Stand" activities, as evident along the diagonal in Figure 4.24. For the other activities, the model demonstrates greater confidence.

Denoising has therefore made the model more selective in recognising activities whose spectrograms differ markedly from those of other classes. However, for activities with similar signatures, this method does not lead to improved performance.

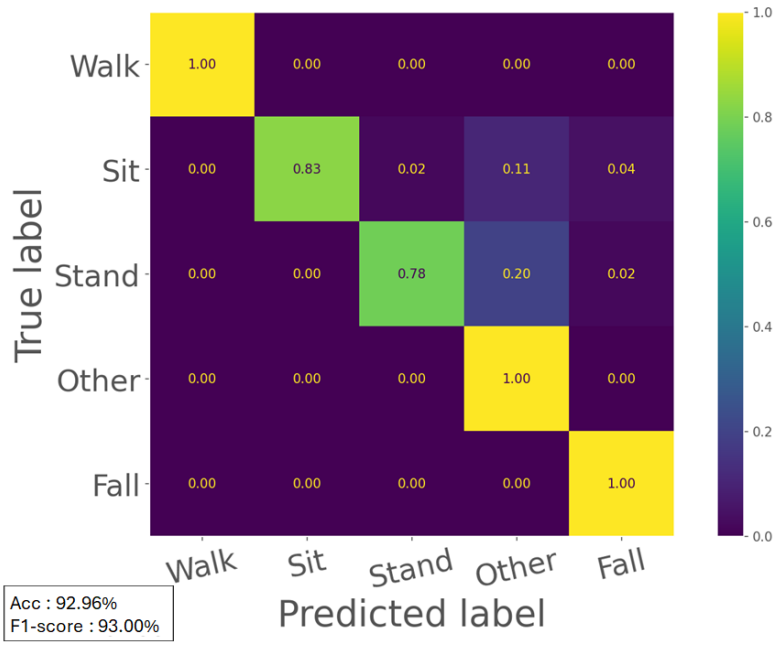

Figure 4.24 – ResNet-18 Normalised Confusion Matrix: 5 Activities with Denoise

This approach may be particularly relevant in one-vs-all scenarios, as it allows the model to concentrate on the energy profile of a single activity. For more general classification tasks, however, its benefit appears limited.

c) Cropping Impact on Learning

The *dynamic crop* method presented in Section 4.2.1 incorporates the energy projection technique. The objective in this case is to focus exclusively on the active data within the

spectrogram, thereby removing all information deemed irrelevant for analysis and training.

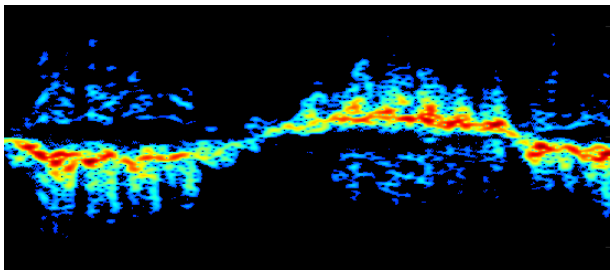

Figure 4.25 – Crop Spectrogram Example

Figure 4.25 provides an example of the data obtained following the application of dynamic cropping, showing that only the activity flow is retained. This data was then used as input to the *ResNet-18* classification model.

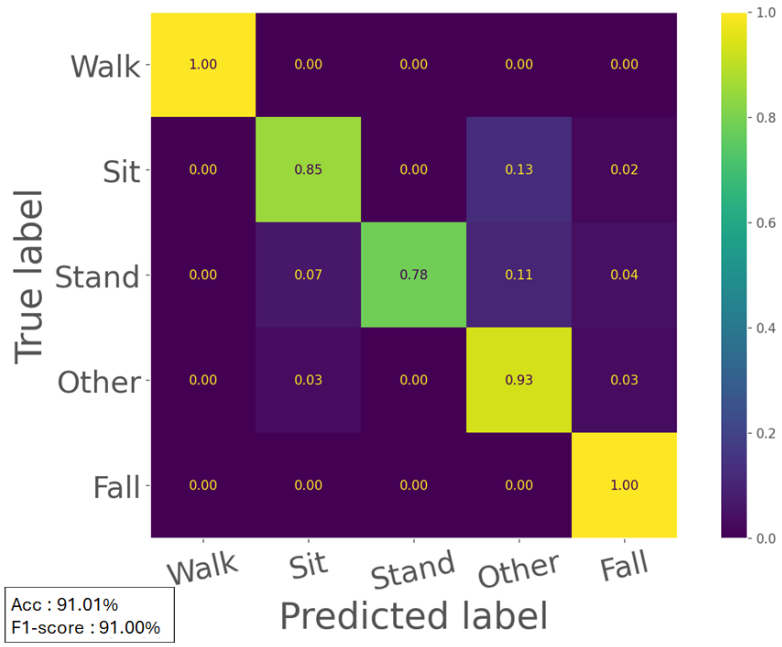

Figure 4.26 – ResNet-18 Normalised Confusion Matrix: 5 Activities with Crop

Figure 4.26 shows the confusion matrix obtained after training. As with the denoising method, uncertainty remains for the "Sit" and "Stand" activities. Overall, this approach appears to introduce even greater confusion. The accuracy achieved with this method is 91.01%.

Table 4.10 provides a summary comparison of the three results obtained. Overall, the results are an improvement over those achieved with the six activities presented in

Table 4.10 – Accuracy Comparison Between Original, Denoise and Crop Methods

Data Type	**Accuracy**	**F1-score**
Original	**97.26**%	97.20%
Denoise	92.96%	93.00%
Crop	91.01%	91.00%

Section 4.1. However, the methods for removing noise and elements deemed non-essential for discriminating activities do not appear to enhance performance. In fact, the model trained on unprocessed spectrograms retains the highest accuracy, achieving 97.26%.

4.3.2 SVM Optimisation and Application for Geometric Forms

Our methodology for improving image quality has shown that the level of noise in the images can affect activity recognition from spectrograms. However, the contribution of this approach is limited, as the differences in results are relatively small. These denoising processes may also render the solution more cumbersome. We consider that it would be preferable to identify more discriminative parameters within the activity information collected.

Section 4.2.2 outlines our methodology for extracting geometric shapes and optimising an *SVM* model. The aim of this section is to conduct tests to evaluate the relevance of a classification approach based on the geometric shape characteristics of a Doppler signature.

a) SVM Hyperparameters Optimisation

Two kernel types were selected for our *SVM*-based human activity classification method. To determine the optimal hyperparameters, we employed two search techniques: *GridSearchCV*, which exhaustively explores a predefined set of parameters, and *RandomizedSearchCV*, which samples a fixed number of combinations at random.

For the *RBF* kernel, we specifically searched for the hyperparameters C and γ. Both search methods converged on the same optimal combination: $C = 1000$ and $\gamma = 0.1$, yielding an accuracy of 87.45%. All the obtained performances are summarised in Table 4.11.

During these tests, we observed an improvement of over 10% in accuracy when comparing the original six activities defined in the dataset with our choice to merge two activities into a single class, termed "Other". Similarly, applying this method to images

Table 4.11 – Accuracy for RBF Kernel Hyperparameter Combinations

	C=100	**C=1000**
γ **= 0.1**	86.75%	**87.45%**
γ **= 1**	82.92%	80.83%

denoised using our dynamic cropping approach did not yield any performance gains. This finding supports our decision to omit denoising for the extraction of geometric shapes.

Similarly, for the *Polynomial* kernel, we searched for the hyperparameters C, *degree*, and $coef0$. Both search methods agreed on $coef0 = 1.0$. However, for C and *degree*, two values each appeared to be of interest: $C = 100$ or $C = 1000$, and $degree = 3$ or $degree = 4$.

We therefore tested the four possible combinations of these hyperparameters. A difference in accuracy of up to 10% was observed between $degree = 3$ with $C = 1000$ and $degree = 4$ with $C = 100$. The best performance was achieved with $degree = 4$ and $C = 100$, as summarised in Table 4.12.

Table 4.12 – Accuracy for Polynomial Kernel Hyperparameter Combinations

	C=100	**C=1000**
degree = 3	87.45%	87.45%
degree = 4	**87.80%**	86.41%

The hyperparameters for both kernel types have thus been selected. The next step is to extract and identify the most relevant features from the spectrogram data.

b) Geometrical Forms Parameters Optimisation

To extract the characteristic data from the spectrogram, the image is first binarised, as illustrated in Figure 4.17b. Following binarisation, the activity signature within the spectrogram is recovered through contour extraction. The eight parameters described in Section 4.2.2, including *Area*, *Perimeter*, and *Orientation*, can then be extracted.

These parameters serve as inputs to the *SVM*, but their relevance must first be assessed to determine whether the full set is necessary. *PCA* is then applied: the data are centred and standardised, and the explained variance of each principal component is computed.

After applying this method, we obtained the graph shown in Figure 4.27. It can be seen that retaining a minimum of five components preserves 90% of the explained variance, which is highly significant for the subsequent results. The higher the explained variance,

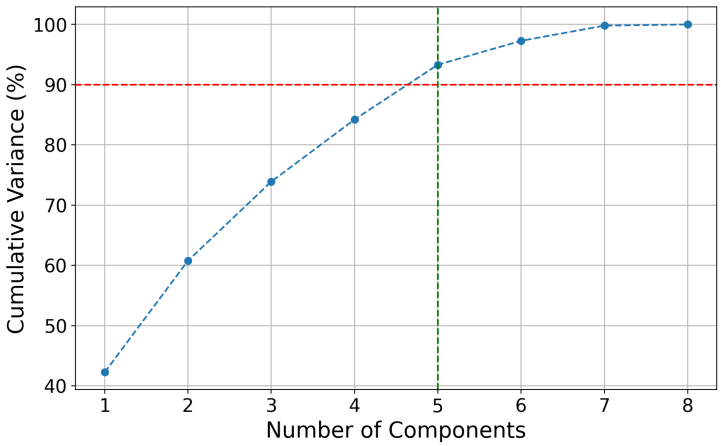

Figure 4.27 – Explained Variance per PCA for SVM Learning

the greater the confidence in the relevance of the outcomes obtained during the learning process.

We then evaluated all possible combinations of parameters. By comparing subsets ranging from two parameters up to the full set, it was found that a subset of five to six parameters is sufficient to achieve a recognition accuracy of approximately 80% or higher. After compiling the top five results, we selected the parameters that appeared most consistently.

For the *RBF* kernel, following the application of the chosen hyperparameters, the most relevant features identified were *Area*, *Perimeter*, *Orientation*, *Major*, *Minor*, *Centroid X*, and *Centroid Y*. In contrast, for the *Polynomial* kernel, with its selected hyperparameters, the necessary features were *Orientation*, *Major*, *Minor*, *Centroid X*, *Centroid Y*, and *Eccentricity*.

c) SVM Application to Geometrical Forms

A *SVM* is employed to learn from the parameters extracted from the spectrogram signatures. The features *Area*, *Perimeter*, *Orientation*, *Major*, *Minor*, *Centroid X*, *Centroid Y*, and *Eccentricity* can thus serve as inputs to the model.

The *SVM* was evaluated using two types of kernel (*RBF* and *Polynomial*) and two approaches: classical SVM and SVM with *PCA*-based dimensionality reduction, employing the *PCA* implementation from *sklearn.decomposition*. For the classical SVM, the best result achieved an accuracy of 88.50% with the *RBF* kernel, using the features *Area*, *Perimeter*, *Orientation*, *Major*, *Minor*, *Centroid X*, and *Centroid Y*. Figure 4.28 shows

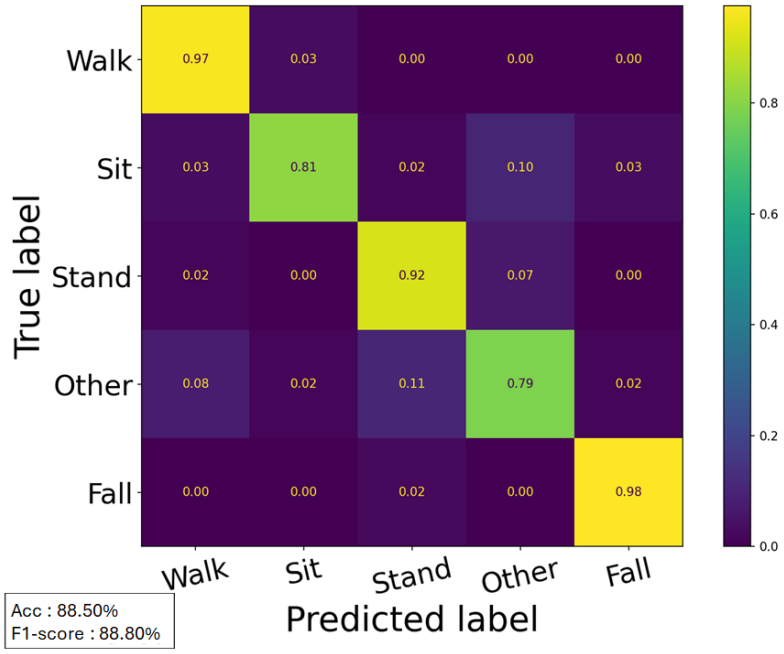

Figure 4.28 – SVM Model with RBF Kernel

the corresponding confusion matrix, which demonstrates a low level of confusion between classes.

The SVM with dimensionality reduction via *PCA*, using a *Polynomial* kernel, achieved an accuracy of 88.85%. This configuration was based on the features *Orientation*, *Major*, *Minor*, *Centroid X*, *Centroid Y*, and *Eccentricity*. The corresponding confusion matrix is presented in Figure 4.29. In both configurations, the matrices show promising results, with well-defined diagonals and minimal confusion between classes. However, the approach incorporating *PCA* achieves slightly higher accuracy.

The use of *SVM* has led to a slight improvement in activity recognition compared with *ResNet-18*. The combination of certain activities has already reduced some of the confusion. Furthermore, the learning approach based on parameters extracted via image processing has enabled further refinement, achieving an improvement of over 1%.

These results highlight the importance of clearly defining the signatures contained in the spectrograms. The more distinctive the data for a given activity, the more accurately activities can be recognised from this type of data. The signature encapsulates the essential information required to distinguish one activity from another.

The results presented here demonstrate that it is entirely feasible to adopt a lightweight solution that is fast to execute and requires minimal resources. Our methodology focuses

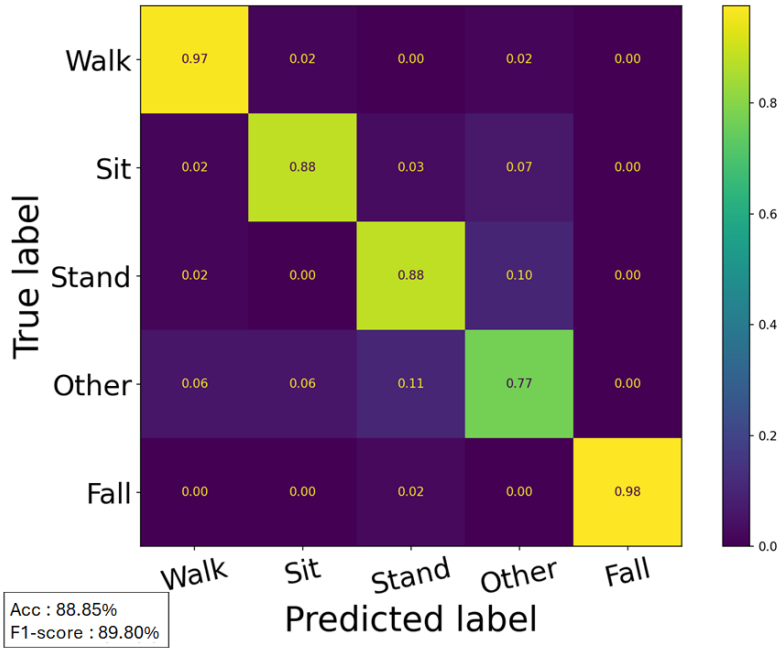

Figure 4.29 – SVM Model with Polynomial Kernel and PCA

on retaining only the most relevant data, thereby optimising the processing pipeline for recognition while maintaining performance.

4.3.3 KNN Optimisation and Application for Textural Features

We have observed that the parameter extraction method based on geometric shapes can be useful. Although the accuracy does not surpass the one achieved with *ResNet-18*, the confusion matrices indicate that the *SVM* model demonstrates a well-distributed overall recognition capability. Consequently, we considered exploring alternative data processing methods, in particular the extraction of features associated with image texture.

Section 4.2.3 describes the method employed to optimise hyperparameters and input features. We also present the Haralick parameters [68] used to extract image texture features. The objective of this section is to determine the optimal parameters and to evaluate their effectiveness for spectrogram classification.

a) KNN Hyperparameters Optimisation

Three hyperparameters must be optimised for our *KNN* model to adapt it to human activity recognition using radar data. To evaluate and select the optimal hyperparameters,

we again employ two search methods: *GridSearchCV* and *RandomisedSearchCV*.

Both methods converge towards the same distance and weighting hyperparameters, summarised in Table 4.13. To determine the most optimal value of K, we employ *cross_val_score* from the Python library *sklearn.model_selection* [10], using the previously identified distance and weighting hyperparameters.

Table 4.13 – Optimum Hyperparameters of KNN

Weights	**Metric**	**K Neighbours**
Distance	Manhattan	22

Once all hyperparameters have been selected, we proceed to search for the optimal texture parameters. Without this optimisation, the hyperparameter search methods suggest that an accuracy of only 58.60% can be achieved.

b) Textural Features Parameters Optimisation

First, the images are converted to greyscale, as illustrated in Figure 4.19. The *GLCM* co-occurrence matrices are then extracted using the *graycomatrix* and *graycoprops* methods. From these matrices, we obtain six texture parameters: *Dissimilarity*, *Correlation*, *Contrast*, *Homogeneity*, *ASM*, and *Energy*.

The Matlab and Python implementations each have limitations. In Matlab, 13 of the 14 Haralick features can be extracted with relative ease. In Python, these methods are still less commonly used, and some parameters are not as well defined. Since we opted to implement an end-to-end pipeline combining data processing and classification within a single codebase, we chose to use the Python implementation.

The six parameters extracted in this manner are used as input to the *KNN* model. However, we aim to optimise this set as much as possible. As with the SVM approach, we employ *PCA* to identify the most informative parameters.

This approach produces the graph shown in Figure 4.30. It can be observed that, using just two parameters, the model achieves 90% explained variance.

We evaluate all possible combinations of parameters. By comparing the different sets, we find that three parameters provide sufficient accuracy, exceeding 60%. The optimal parameters identified are *Correlation*, *Contrast*, and *Homogeneity*. Using this texture

10. `https://scikit-learn.org/stable/modules/generated/sklearn.model_selection.cross_val_score.html`

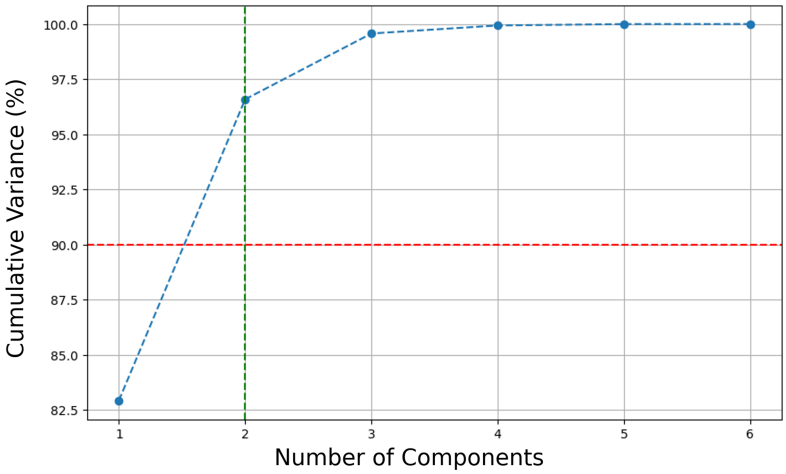

Figure 4.30 – Explained Variance per PCA for KNN Learning

extraction method, however, we observe that the classification models still struggle to differentiate between certain classes.

The complexity of the texture data and the definitions of the parameters appear to constrain the model during training. This may also be attributed to the presence of multiple textures within the same image. Overall, the model generally achieves accuracies above 50%. By optimising the parameters, we are able to increase the accuracy to above 60%.

c) KNN Application to Textural Features

Once the parameters and hyperparameters have been optimised, we proceed to train the *KNN* model. Initially, the model is trained using the optimal parameters directly. Subsequently, it is trained using the *PCA* function from the Python *sklearn.decomposition* library.

Since *KNN* is well suited to small datasets and is highly flexible, we evaluate the model on the dataset *Radar Signatures of Human Activities* [57]. The accuracies obtained are modest, not exceeding 70%. At first glance, this approach does not appear to be the most decisive for producing a reliable model for detecting and predicting human activities. However, this suggests that the *KNN* model could provide useful uncertainty estimates in cases where the first two models might be overly confident in recognising an activity.

We also tested these data using *SVM* and *RF* models, and the accuracy results were essentially identical to those obtained with the *KNN* model. For this reason, we present

only the *KNN* results here. The other tests did not provide additional insights that would allow us to draw new conclusions or explore alternative approaches.

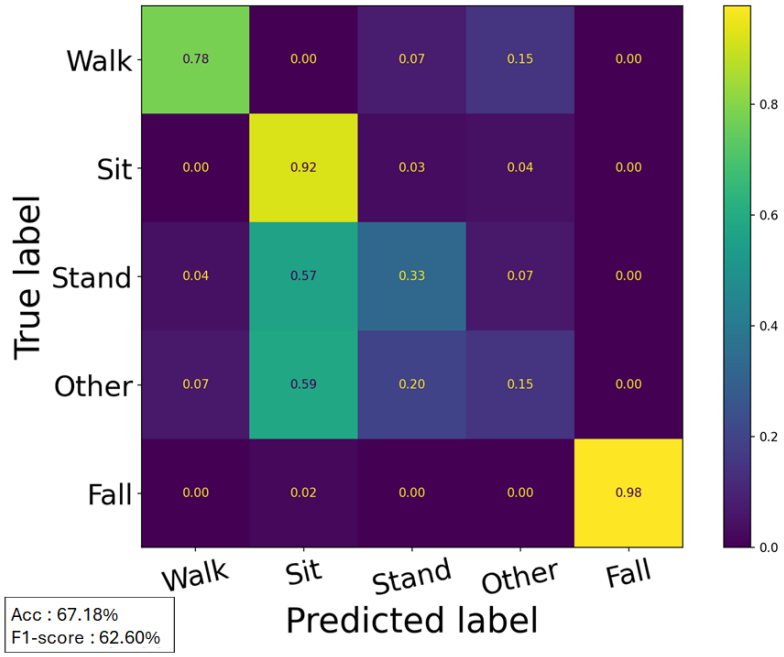

Figure 4.31 – KNN Normalised Confusion Matrix

The matrix shown in Figure 4.31 exhibits the beginnings of a diagonal, but there is considerable confusion between the activities "Sit", "Stand", and "Other". This may be attributed, on the one hand, to the similarity between "Sit" and "Stand" noted in previous sections, and, on the other hand, to the presence of multiple textures, which creates overlap with the "Other" category. Despite this confusion, the model achieves an accuracy of 67.18%.

Applying the *PCA* method does not lead to a significant improvement in accuracy. The confusion matrix in Figure 4.32 shows that the results are very similar across all classes. The accuracy achieved is 66.79%, indicating that the use of *PCA* is unnecessary in this case.

Although these results are less favourable than those achieved with the *ResNet-18* and *SVM* models, we believe they may still be useful in complementing the other models. Indeed, the predictions from each model could be leveraged to provide an indication of certainty or uncertainty in signature recognition.

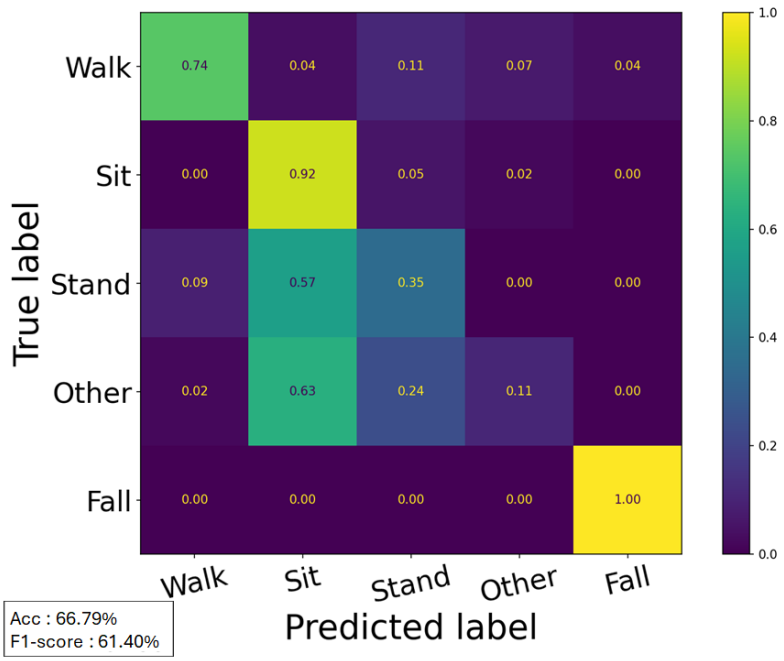

Figure 4.32 – KNN Normalised Confusion Matrix with PCA

4.3.4 Evaluation of the Fusion by Majority Vote

In the previous three sections, we evaluated three distinct classification models. *ResNet-18* and *SVM* produced convincing results with their respective input data processing methods, achieving accuracies around 90%. *KNN* exhibited more confusion overall but performed well for the "Walk" and "Fall" activities, which are the most frequently observed and clinically relevant in a medical context.

For this reason, we propose a method for merging the predictions from the three models. Specifically, we aim to investigate whether combining their results can enhance the reliability of the final predictions.

Figure 4.33 provides a schematic representation of the proposed processing workflow. First, the three learning processes are carried out using the *ResNet-18*, *SVM*, and *KNN* models. *ResNet-18* takes data from the *Radar Signatures of Human Activities* [57] dataset as input, with images resized to thumbnails of 32×32 pixels. The *SVM* model operates on the geometric shape features of the Doppler signatures, using optimised hyperparameters and input features. Finally, the *KNN* model utilises Haralick features [68], also with optimised hyperparameters and input parameters.

Once the three learning processes have been completed, their predictions can be com-

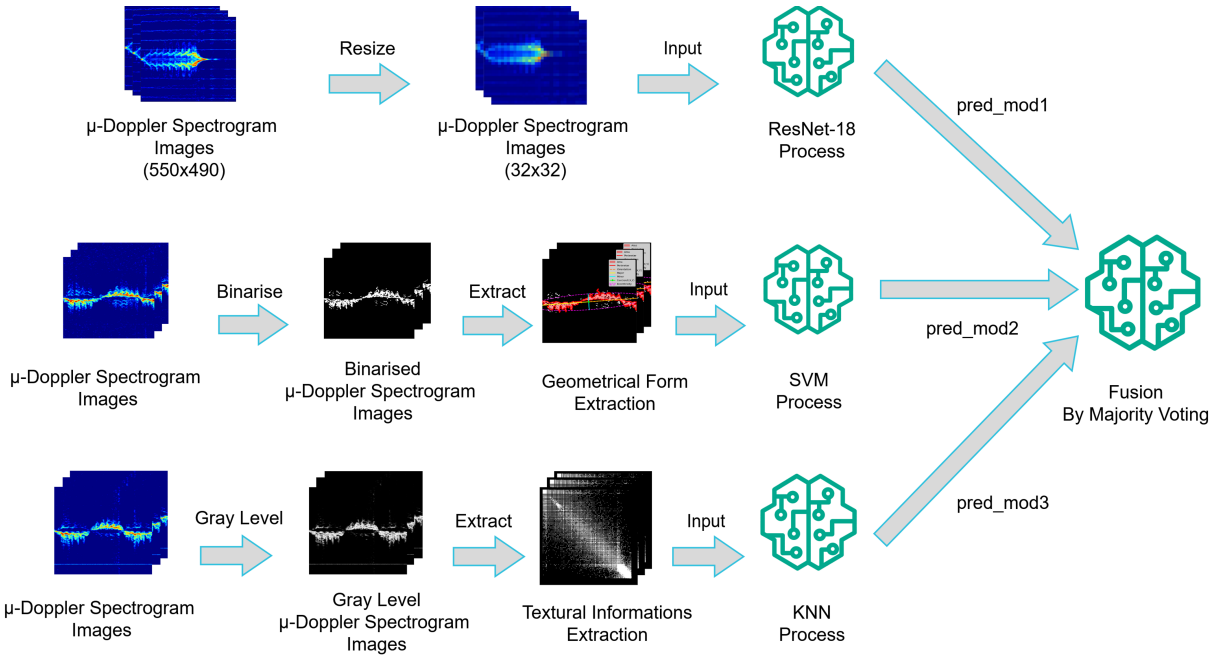

Figure 4.33 – Global Chain for Training with Majority Voting

bined using a majority voting rule. The underlying idea is that if two or more models predict the same activity, there is a higher probability that this prediction is correct. This approach introduces both certainty and controlled ambiguity into the final model, with the aim of mitigating overconfidence and avoiding overly subjective predictions.

Our method can be summarised in the following steps:

- Let $pred_mod1$, $pred_mod2$, and $pred_mod3$ denote the predictions from *ResNet-18*, *SVM*, and *KNN*, respectively. $pred_fusion$ represents the prediction resulting from the fusion of the three models.
- If $pred_mod1 = pred_mod2 = pred_mod3$,
  Then $pred_fusion = pred_mod1$
- Otherwise, if at least two of the models agree on the prediction,
  Then $pred_fusion = pred_majority$
- Otherwise, as *ResNet-18* is considered the most confident model,
  $pred_fusion = pred_mod1$ by default

We use the ResNet-18 results presented in Figure 4.22 in Section 4.3.1, which achieve an accuracy of 97.26%. For SVM, we focus on the results obtained using a *Polynomial* kernel combined with *PCA*, yielding an accuracy of 85.54%. Finally, KNN without *PCA*

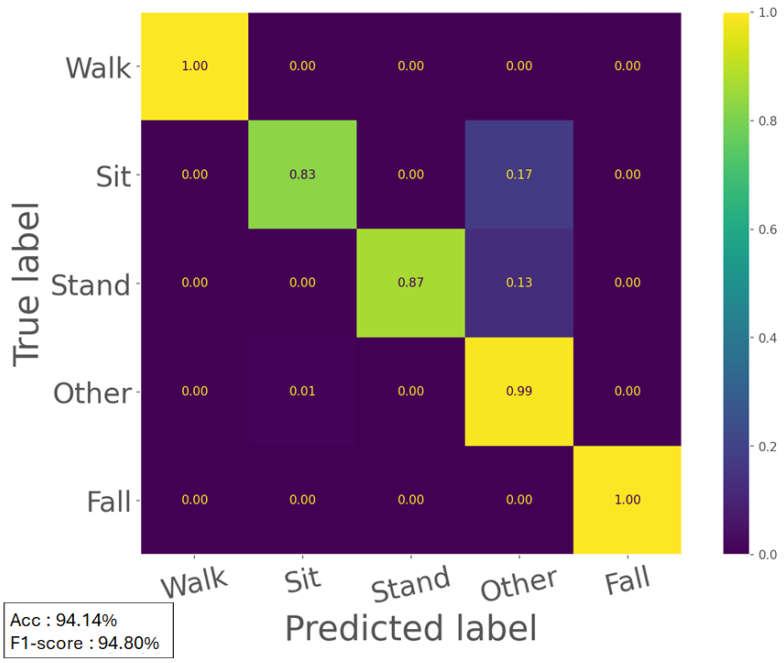

Figure 4.34 – Majority Voting Normalised Confusion Matrix

achieves an accuracy of 67.18%.

Table 4.14 – Summary of the Highest Accuracies Achieved on the *Radar Signatures of Human Activities* [57] Dataset

Model	**Accuracy**	**F1-score**
ResNet-18	97.26%	97.20%
SVM	85.54%	85.20%
KNN	67.18%	62.60%
Fusion	94.14%	94.80%

This majority voting fusion method produces the confusion matrix shown in Figure 4.34. The final accuracy achieved is 94.14%. Despite the high confusion observed in the KNN model, a closer examination of the votes indicates that the contributions from *SVM* and *KNN* were effective. Although the fusion accuracy is initially around 3% lower than that of *ResNet-18*, it is noteworthy that some errors made by ResNet-18 were corrected through the majority vote of the other two models. The detailed accuracies are summarised in Table 4.14.

This method could be further improved by considering the *Top-N* predictions rather than only the *Top-1*. The three methods complement each other effectively, and the ro-

bustness and reliability of the results could be enhanced by refining this fusion approach.

4.4 Conclusion

Radar spectrograms provide rich information on human movements, making them a valuable tool for human activity recognition. In this chapter, we conducted an *experimental evaluation of feature extraction and recognition methods*, focusing on both *deep learning architectures* (ResNet-18) and *classical machine learning models* (SVM, KNN) applied to spectrogram data.

Our experiments highlighted several factors influencing model performance. The *data collection protocol* can create similarities between classes, while *image resizing* significantly improved ResNet-18 accuracy by up to 10%. *Image quality*, affected by noise, was addressed using a *dynamic cropping method*. Additionally, *geometric* and *textural features* of activity signatures were extracted and optimised through *PCA* and feature combination search. SVM and KNN hyperparameters were carefully tuned to maximise performance, and a *majority-vote fusion* strategy was applied to combine the outputs of different models and balance their predictions.

The results demonstrate the effectiveness of these approaches: *ResNet-18* achieved 97.26% accuracy, *SVM* 88.85%, *KNN* 67.18%, and the *fused* model reached 94.14%. Overall, the combination of optimised models improved *accuracy, robustness, and generalisation* across the *Radar Signatures of Human Activities* dataset [57].

Despite these promising results, several challenges remain. *Dataset* variability, data quality, class imbalance, and limited sample sizes indicate the need for further validation. Moreover, the *scarcity of datasets* specifically aligned with the targeted *clinical* context limits the generalisation of these methods. Figure 4.35 provides an overview of the complete experimental chain developed and assessed in this chapter.

To address these *limitations*, Chapter 5 will present the creation of a *new dataset*, including the collection protocol, equipment, recorded activities, and associated statistics. The *classification chain* will be evaluated on this dataset to assess its *robustness and effectiveness*, thereby preparing for eventual deployment in real clinical environments.

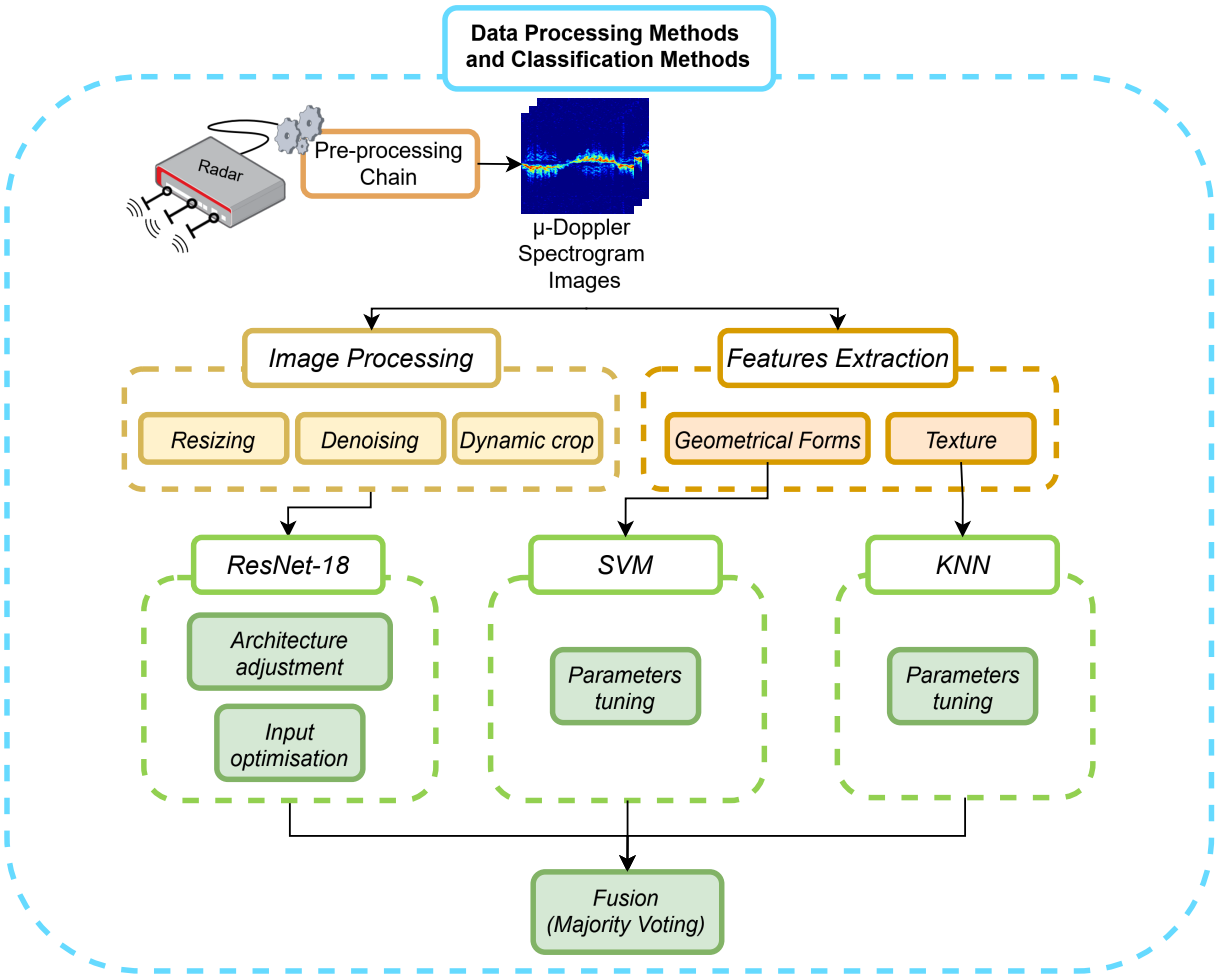

Figure 4.35 – Summary of Data Processing and Classification Methods, Chapter 4

Chapter 5

MAD: Creation of a Multimodal Dataset and Classification Analysis

The results obtained in Chapter 4 enable us to consider several possible solutions. We have explored different approaches to data processing, as well as the potential use of various types of learning models. However, the field of *Radar*-based activity recognition suffers from a lack of data. Both the quality of the acquisitions and the protocols adopted by the designers of existing datasets can significantly affect the learning performance of the models.

Indeed, as demonstrated by the *Radar Signatures of Human Activities* [57] dataset, defining one activity relative to another can rapidly introduce ambiguity. More generally, this observation highlights the persistent problem of *taxonomy* within the broader field of activity recognition, as discussed in Chapter 3.

There is therefore a notable scarcity of both acquired and publicly available data, further constrained by the specific contexts in which these studies were conducted. Consequently, this imposes *limitations* on the scope for implementation and processing. For instance, the *Radar Signatures of Human Activities* [57] dataset presents several shortcomings: it comprises a limited amount of data, exhibits significant imbalance, and is, in large part, affected by noise.

The reliability of our proposed learning approaches also depends on the data employed. Our context is highly specific, focusing on falls and fall risks among elderly and hospitalised individuals. This necessitates validation of our solution by healthcare professionals and authorised institutions. Furthermore, the acquisition of new data in environments closely resembling the target setting has become increasingly essential.

The primary objective is to strengthen collaboration with all stakeholders directly concerned by the proposed solution, including healthcare professionals, patients, and their

families or caregivers. This engagement also enables an assessment of the practical and material constraints involved in the project's implementation. Within this framework, we have compiled a new dataset, developed a dedicated set of equipment, and established a testing platform within the laboratory. In addition, a protocol has been designed to enable comparison with the *Radar Signatures of Human Activities* [57] dataset while introducing activities of particular relevance to our study. Finally, the project received approval from the CY Cergy Paris Université Ethics Committee (reference number: 202307-004), authorising the inclusion of members of the public, see Appendix A.6.

This chapter details all the equipment and materials utilised in the study, together with their respective functions. It subsequently describes the complete data acquisition process. Finally, it presents the results obtained using the same data processing and learning methodologies as those applied to the *Radar Signatures of Human Activities* [57] dataset. To provide a clear overview of the methodological workflow described in this chapter, Figure 5.1 summarises the design of the experimental framework, the acquisition of multimodal data, and the subsequent evaluation of learning models.

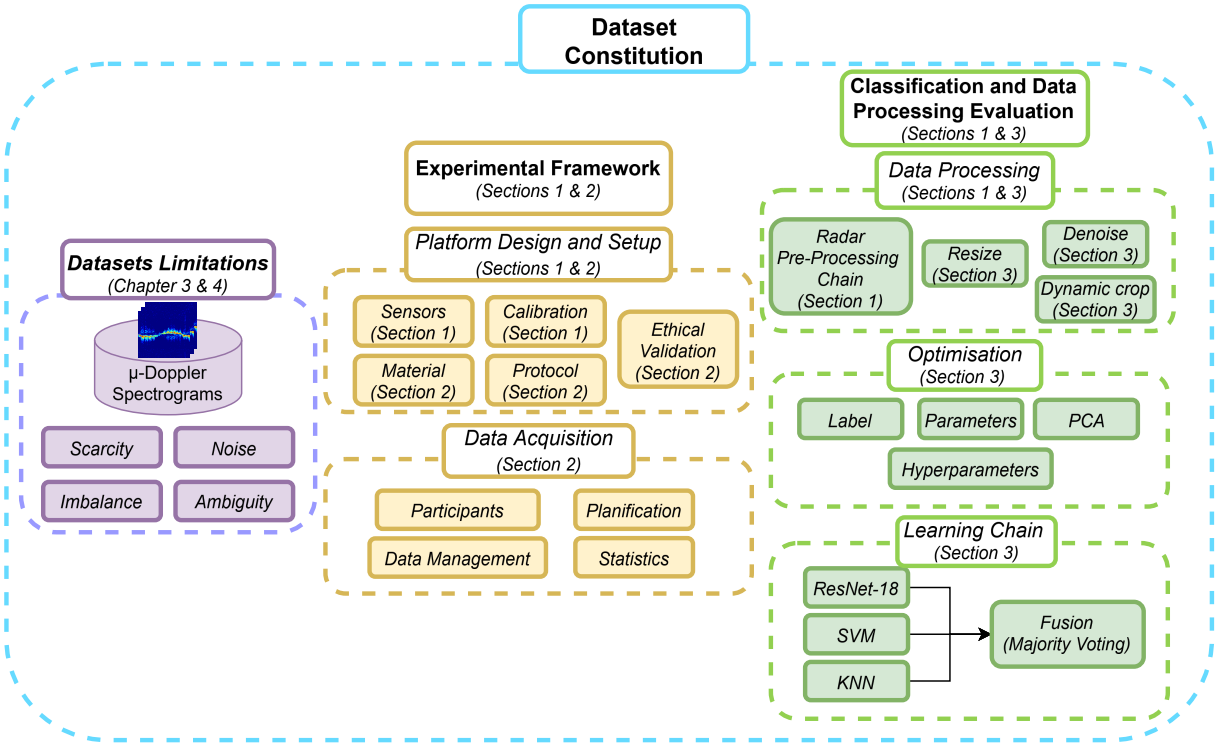

Figure 5.1 – Overview of Chapter 5

5.1 Movement Platform and Data Processing

To perform our various tests and data acquisitions, we utilise a dedicated facility at the ETIS laboratory known as the *Movement Platform*. This room is specifically designed for experimental testing and enables the acquisition of data from a variety of sensors for the study of human movement, whether involving single participants or groups.

The aim of this section is to describe the sensors installed in the room and their application to human activity recognition. Specifically, a *Radar* sensor is utilised, and a dedicated pre-processing chain has been developed to produce custom *spectrograms*. In addition, a validation system has been established to assess the accuracy of our configuration using *ground-truth* data.

The aim of this section is to describe the sensors installed in the room and their application to *human activity recognition*, building upon the review of sensor technologies in Section 3.1.2 and methods in Section 3.2.2. Specifically, a *Radar* sensor is utilised, and a dedicated pre-processing chain has been developed to produce custom *spectrograms*, following the approaches discussed in Section 3.1.3. In addition, a validation system has been established to assess the accuracy of our configuration using *ground-truth* data, as outlined in Section 3.2.1.

5.1.1 Movement Platform Configuration and Equipment

The *Movement Platform* is a facility within the ETIS laboratory that has been repurposed as an experimental space for the human activity analysis project. The room, rectangular in shape and measuring 12x4.5m, was designated for this project in 2021.

The room is equipped with five active sensors and is designed to accommodate additional devices as required. This flexibility enables the testing of various configurations and the combined operation of selected sensors for comparative analysis. Over the years, the layout and configuration of the room have evolved significantly, improved by experimental outcomes, specific research requirements, and the phenomena targeted for observation.

a) Material: Sensors

This thesis constitutes a component of a wider research project, as detailed in the Chapter 1. It extends the work carried out in Alexandre Bordat's thesis [23], which addresses the early detection of biomechanical deficiencies through the use of unconventional *Radar* imaging. Notably, his work utilises the *TUG* test, as outlined in Section 2.1.2.

In his thesis, Alexandre Bordat [23] established an initial *Radar* system and evaluated the application of a *Pressure Sensor Mat* within a medical context. Additionally, a *Motion Capture* (MoCap) system was deployed to facilitate comparison between the data acquired from the *Radar* and that from *MoCap*. This system includes a *conventional RGB camera* as part of its setup, as illustred in Figure 5.2.



Figure 5.2 – Set-Up of the *Movement Platform*

At the same time, the facility is utilised by another research team [3] engaged in a distinct project. Their focus is on assessing the application of an *Event-based camera* in various contexts, with the specific goal of devising an implementation that allows learning to be performed directly on complex event data, rather than on reconstructed images derived from the sensor's captured events.

The *Movement Platform* is equipped with the following sensors:

- **FMCW Radar:** is primarily used in the room to conduct various tests. We employ the SDR 980AD2 model from Ancortek, as shown in Figure 5.3, which is equipped with a transmitting antenna (Tx). Its function is to emit an electromagnetic wave towards a target, in this case a person. The wave is reflected upon striking the target, generating an echo that is captured by two receiving antennas (Rx).

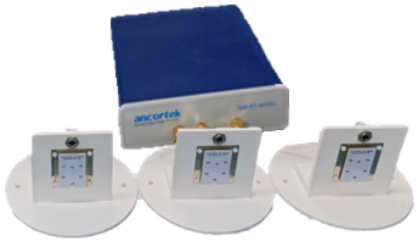

Figure 5.3 – SDR 980AD2 Ancortek FMCW Radar

This echo contains information on both the distance and radial velocity relative to the *Radar*. The *Radar* operates at a central frequency of $9.8GHz$ with a bandwidth of $4MHz$ and is controlled via a Jetson Xavier [25] using the C++ programming language.

- **Optical Motion Capture System:** validates the *Radar* data and ensures its accuracy. It uses passive retro-reflective markers tracked by infrared cameras, a highly flexible and widely adopted method. The *Movement Platform* is equipped with 10 Vicon Vero v2.2 cameras, as shown in Figure 5.4.



Figure 5.4 – Vicon Vero V2.2 Camera

These cameras can track fast-moving subjects, such as athletes or drones, and multiple participants or objects simultaneously with minimal latency. Marker positions are recorded at $100Hz$. As the room's walls are fully glazed, a light-filtering film was applied to the windows to minimise daylight reflections and interference.

- **Camera:** is a *conventional RGB* unit provided as part of the Vicon system, designed to complement the *MoCap* setup, as depicted in Figure 5.5. Synchronous calibration with the Vicon optical cameras ensures precise alignment between the optical and video perspectives, enabling the capture of detailed movements.



Figure 5.5 – Vicon Vue Camera

The camera offers a resolution of $2.1MP$ (1920x1080). Its primary function is to provide visual verification of the movements recorded by the remainder of the system, which tracks only the retro-reflective markers affixed to the subject being analysed.

- **GAITRite Pro Mat:** serves to complement the *Radar* and *MoCap* system by recording spatio-temporal gait parameters [123], including cadence, step symmetry, and step regularity, as shown in Figure 5.6. It is the sole sensor in the room capable of measuring the projection of the centre of gravity onto the support polygon. The mat is employed primarily for single-leg tests.

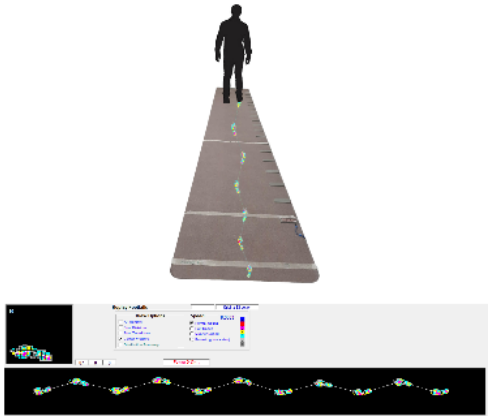

Figure 5.6 – GAITRite Pro Mat

However, the standard configuration does not permit synchronised recording with the other sensors, rendering it unsuitable for use in the *TUG* test.

- **Event-Based Camera:** is used by a project team independent from our own, operating in the same laboratory. It is the EVK1-VGA model provided by Prophesee, as illustrated in Figure 5.7. The sensor has a resolution of 640×480 pixels, a latency of 1μs, and an optical format of 3/4".



Figure 5.7 – EVK1-VGA from Prophesee Event-Based Camera

They use the original lens supplied with the camera, the VSTechnology SV-0813V, which provides a diagonal *Field of View* (FOV) of 70°. The camera requires only a USB 3.0 connection to a computer for power and data transfer. Data acquisition is performed using the Prophesee Metavision Python API [3].

This project employs the *Radar*, *MoCap*, and *GAITRite mat* systems. The thesis aims to achieve activity recognition based solely on *Radar* data. Although all laboratory sensors require calibration, only the *Radar* data are used for learning, with the *MoCap* system providing sufficient validation. The *GAITRite mat* is therefore unnecessary. To ensure calibration accuracy and data consistency, the *Radar* is synchronised with the *MoCap* system, following the temporal-domain approach proposed by Alexandre Bordat [23].

b) Calibration and Synchronisation Equipment: Radar and MoCap

A *Force Sensing Resistor* (FSR) was integrated into a chair identical to those used in nursing homes to detect the start of a *TUG* test or any movement from a seated position. When the chair is unoccupied, data recording can be manually triggered via remote control. The *FSR* varies its electrical resistance with applied pressure, and upon standing, a 433 MHz transmitter emits a signal to synchronise all measurement devices. The configuration developed by Alexandre Bordat [23], including the sensor's placement and integration, is shown in Figure 5.8.



(a) Instrumented chair found in care facilities

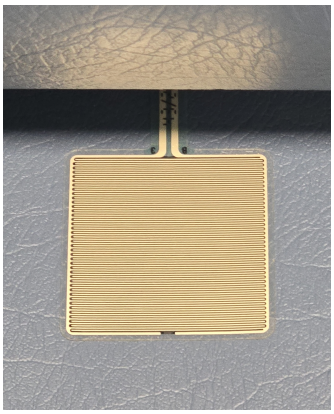

(b) Force sensor located on the chair seat

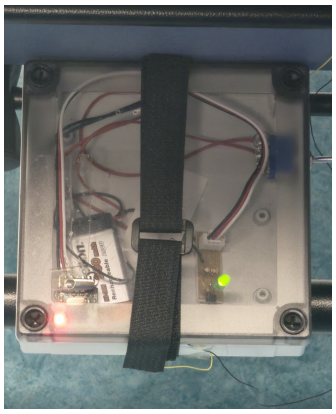

(c) FSR transmitter and resistance measurement unit

Figure 5.8 – Illustrations of the Components in the Instrumented Chair Setup [23]. The first image shows the instrumented chair, the second shows the force sensor, and the third shows the transmitter and resistance measurement unit.

Upon transmission of the signal by either the force sensor unit or the remote control, the receivers detect the message and initiate the recording process. In the *Radar* system, the receiver is interfaced with the *General-Purpose Input/Output* (GPIO) ports of the Jetson Xavier, which handles both data acquisition and processing for the *Radar*. The Jetson Xavier is connected to the Ancortek *Radar* via USB, as depicted in Figure 5.9. In the *MoCap* system, the receiver is connected to the Vicon Lock Lab, which triggers recording and ensures synchronization across all cameras in accordance with the parameters configured in Vicon *Nexus*.

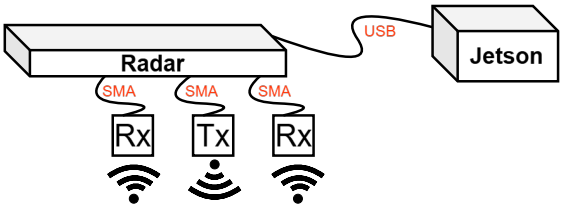

Figure 5.9 – Radar Setup Diagram

The spatial references of the Vicon cameras and the *Radar* are initially aligned through a calibration procedure. Additionally, at the commencement of each *Radar* data acquisition session, spatial calibration is performed using a Vicon calibration wand, the *Magic Wand V2*, as illustrated in Figure 5.10. Notably, the y-axis of the *MoCap* system is aligned with the *Radar* antennas, which are in turn aligned with the walking axis of the *TUG* test. This configuration provides optimal conditions for precise measurement of both distance and radial velocity [23].



Figure 5.10 – Magic Wand V2

The system delivers reliable and verifiable data. Moreover, it is highly adaptable, allowing for data acquisition from both sensors simultaneously or from a single sensor as required. The chair or remote control facilitates the selection of activities to be performed with full control over the experimental conditions.

c) Radar Acquisition

For the initial acquisition tests, which supplied the data employed in the *ResNet-18* analyses presented in Section 4.1, recordings were conducted with a duration of 10 seconds per session.

Our *Radar* operates in the X-band with a carrier frequency of $9.8GHz$, an instantaneous bandwidth of $400MHz$, and a chirp duration of $1ms$. This configuration provides a Doppler frequency range of $\pm 500Hz$, which is sufficient for detecting indoor human activities [146].

The system operates with a transmit power of $+19dBm$. The angle of arrival of the backscattered signal is determined using the two receiving antennas. These antennas are equipped with *SubMiniature version A* (SMA) connectors, which are semi-precision coaxial RF connectors designed as a standard interface for coaxial cables, featuring a screw-coupling mechanism.

The *Radar* system utilises a hardware dechirping, or stretch processing, approach, whereby the received signal is mixed with the transmitted signal to perform down-conversion. The resultant down-converted signal, known as the beat frequency signal, is then processed through low-pass filtering and amplification before being digitised.

The *Intermediate Frequency* (IF) signal consists of two channels, *In-phase* (I) and *Quadrature* (Q), allowing the capture of complex-valued samples. The *Analog-to-Digital Converter* (ADC) samples at $128kHz$, yielding 128 samples per $1ms$ sweep. The *Radar*'s Application Programming Interface (API) requires predefined recording lengths, with each session lasting 10 seconds. This duration is adequate to capture the daily activities.

5.1.2 Pre-Processing Chain

The *Radar* data acquired using our equipment in the *Movement Platform* are processed to derive multiple types of information. As highlighted in previous chapters, and particularly in the context of the reference *Radar Signatures of Human Activities* [57] dataset, it is essential to implement a processing chain to interpret the recorded signals. Furthermore, as part of his thesis, Alexandre Bordat [23] introduced a pre-processing chain designed to rapidly produce reliable data.

The pre-processing chain is structured into multiple sub-processing blocks, as depicted in Figure 5.11. The initial block handles the recording and formatting of the input signal. Since the signal is interleaved, it must first be deinterleaved to extract the I/Q data from

each antenna. The resulting data are then arranged into a $128 \times M$ matrix, where 128 denotes the number of time samples per chirp, and M denotes the total number of chirps corresponding to the activity recording.

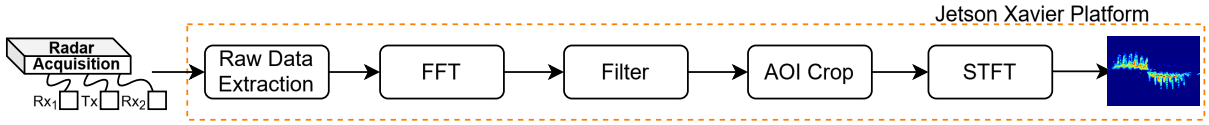

Figure 5.11 – Pre-Processing Chain with FMCW Radar (1 Transmit and 2 Receive Antennas)

Following this, a *Fast Fourier Transform* (FFT) is performed on each received chirp to estimate the radial distance between the receiving antenna and the participant, which is directly proportional to the beat frequency. The resulting data are filtered using a *5-pulse canceller* (5-PC) [42], applied along the slow-time dimension, representing the sequence of chirps over time. This filter serves as a moving target indicator, suppressing components near $0Hz$ generated by stationary objects in the environment, while maintaining the fidelity of the spectrogram.

A range bin band between 5 and 25, corresponding to $1.875m$ to $9.375m$, is selected to concentrate on the spatial region encompassing the target's movement relative to the *Radar*. This procedure, known as the *Area of Interest* (AOI) crop, serves to decrease processing time and suppress noise in the calculation of temporal velocity signatures.

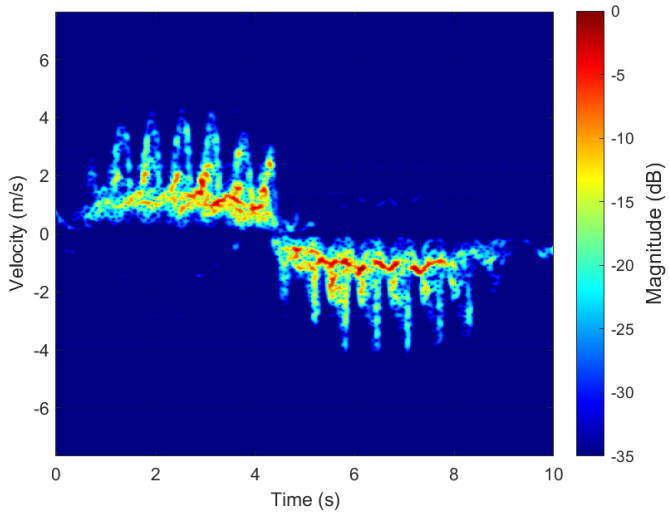

Figure 5.12 – Our Velocity Time Graph (Spectrogram)

Finally, a *Short-Time Fourier Transform* (STFT) is performed on the range-time data matrix employing a *Hamming* window of 0.2 seconds with a 95% overlap. This facilitates the extraction of time-varying micro-Doppler features. The *STFT* achieves a range resolution of $37.5cm$, a maximum measurable velocity of $7.65m/s$, and a Doppler resolution of $1.25Hz$ (corresponding to $0.03m/s$).

$$X(t,\omega) = \sum_{m=0}^{N-1} x[m] \cdot w[m-t] \cdot \exp\left(-i\omega m\right) \tag{5.1}$$

Equation 5.1 defines the *STFT*, where $X(t,\omega)$ denotes the time-frequency representation of the signal, $x[m]$ is the input signal, $w[m-t]$ is the *Hamming* window centred at time t, and $\exp\left(-i\omega m\right)$ represents the complex exponential term. In this notation, m is the sample index, N the total number of samples, ω the angular frequency, and i the imaginary unit. This transformation enables the examination of the temporal evolution of the signal's frequency content, yielding detailed micro-Doppler features suitable for velocity analysis and activity classification.

The output of this processing chain comprises data matrices that depict the radial velocity of the target over time, referred to as *spectrograms*, as shown in Figure 5.12.

The achieved resolution is adequate for precise velocity analysis and activity classification. This *Radar* configuration was chosen on the basis of its commercial availability and the need to validate the proposed methodology in realistic conditions. While the range resolution (∼38 cm) is lower than that attainable with higher-bandwidth or millimetre-wave *radars*, this compromise guarantees reliable performance in indoor settings.

For the experiments, the pre-processing chain was executed on the Jetson Xavier platform [1]. This system incorporates *One HMP* (Heterogeneous Multi-core Processing) CPU comprising 4 × 2 Carmel cores, alongside a Volta GPU with 512 CUDA cores. The CuFFT library [2] in CUDA is utilised for batch processing, employing the *PlanMany* function to perform *FFT* and *STFT* operations.

Custom CUDA kernels have been implemented to derive the spectrogram from the STFT. Currently, the filtering stage is executed on the CPU using coefficients precalculated in MATLAB. The configuration of the *Radar* system is depicted in Figure 5.9.

1. `https://www.nvidia.com/fr-fr/autonomous-machines/embedded-systems/jetson-agx-xavier/`
2. `https://docs.nvidia.com/cuda/cufft/index.html`

5.1.3 Motion Capture as a Ground Truth: Equipment, Acquisition and Processing

We use *MoCap* as a reference for data quality and, more importantly, to verify our interpretation of the acquired data. To ensure accurate field measurements, the participant must first be appropriately equipped. Our focus is on the lower body, as signs of balance impairment predominantly manifest in this region during the performance of activities.

a) MoCap: Equipement and Principles

The deployment of 10 cameras, relative to the recorded space, ensures accurate data tracking through *trilateration*. Markers, composed of glass microspheres, are placed on the participant to reflect the body segments whose movement data are to be captured. The system also includes an *RGB camera*, which visualises the scene and provides support for skeletal reconstruction if required.

The *Nexus* software is employed for system calibration and for both real-time and offline visualisation of the data. Based on the markers affixed to the participant, a skeletal model can be calibrated and automatically reconstructed, enabling more effective visualisation and interpretation of the captured motion data.



(a) Markers worn by the participant at the ankles, knees and around the waist (circled in red)

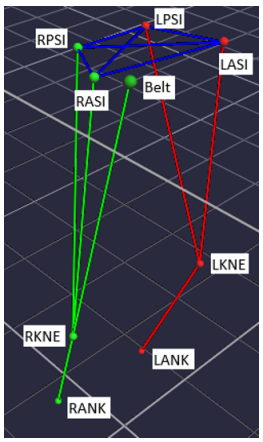

(b) Reconstructed skeleton (segments) in Vicon's *Nexus* Software, dots represents the markers

Marker ID	Description
Belt	Belt-mounted
LASI	Left Anterior Superior Iliac Spine
RASI	Right Anterior Superior Iliac Spine
RPSI	Right Posterior Superior Iliac Spine
LPSI	Left Posterior Superior Iliac Spine
RKNE	Right Knee
LKNE	Left Knee
LANK	Left Ankle
RANK	Right Ankle

(c) Correspondence table of marker acronyms and their descriptions

Figure 5.13 – Motion Capture Markers in Real and their *Nexus*'s Representation

The *Nexus* software facilitates the calibration of all cameras and the generation of a

skeletal template. Starting from a predefined model in *Nexus*, the skeleton was modified to produce the configuration illustrated in Figure 5.13.

A total of 9 markers are used to capture lower limb movement data. An additional marker is placed on the belt to capture *asymmetry*, enabling differentiation between the right and left sides of the participant. This belt marker is assumed to have the same velocity as the torso.

The ankles are specifically monitored to analyse leg swings during activities that place substantial demands on the lower limbs. The torso is similarly of interest for high-acceleration movements, such as falls. Data acquisition for both the *Radar* and *MoCap* systems is initiated simultaneously at time 0 to ensure synchronisation.

b) Acquisition Space and Relevant Data

The *MoCap* system acquires marker positions at 100 measurements per second within the recording volume, relative to a reference origin positioned at the centre of the volume, referred to as point B in Figure 5.14. The dedicated Vicon *Nexus* hardware processes the trajectories of these markers and extracts the corresponding data.

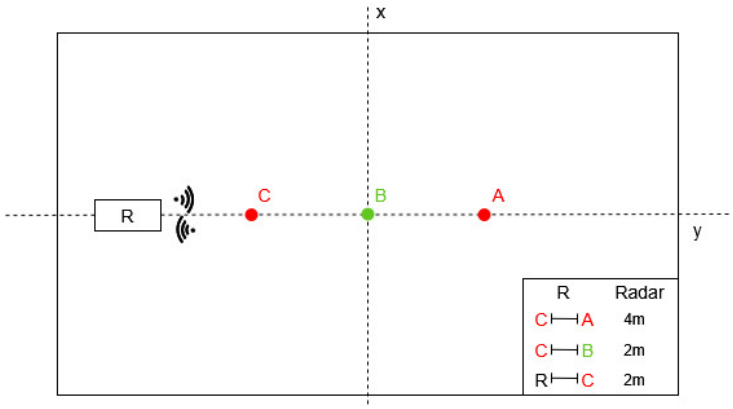

Figure 5.14 – Top-View Schematic of Activity Layout within the *Movement Platform*

The files include data on the distance, velocity, and acceleration of each marker, positioned as illustrated in Figure 5.13. The y-component aligns with the radial axis of the *Radar*, permitting direct comparison with *radar*-derived measurements. Velocities are originally recorded in *mm/s*, and are subsequently converted to *m/s*.

This enables the simultaneous plotting of velocity–time graphs for multiple markers, as illustrated in Figure 5.15, providing a comparative visualisation of movement dynamics across different body segments.

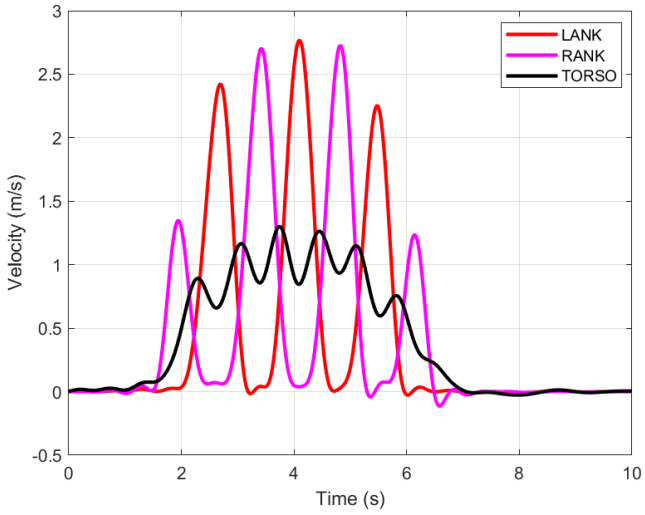

Figure 5.15 – *MoCap* Velocity-Time Graph of *Walk*

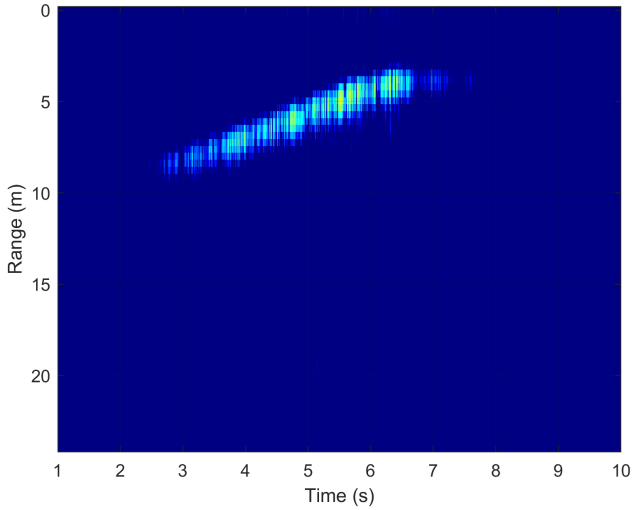

Figure 5.16 – Range-Time Representation of *Walk* Activity Detected by *Radar*

Following *Radar* signal acquisition, a *FFT* is performed along the fast-time axis to obtain distance information. Each chirp yields a beat frequency f_{beat} that is proportional to the target's distance, expressed as:

$$R = \frac{c \cdot f_{\text{beat}}}{2B}$$

where c denotes the speed of light and B the chirp bandwidth. The resulting range–time

matrix comprises columns corresponding to successive time instants and rows corresponding to distances. This matrix is illustrated in Figure 5.16, where the intensity encodes the amplitude of the reflected signal at various distances over time, enabling an intuitive visualisation of target motion.

To determine the person's location in space and time, a horizontal projection is performed on the range–time matrix, also referred to as a *histogram*, highlighting the range bins where the person is present. This selection enables computations for the *STFT* to be restricted to the relevant range bins, thereby optimising velocity extraction.

After performing a *STFT* along the slow-time dimension, a Doppler frequency–time matrix is obtained, in which each row represents a specific Doppler frequency and each column corresponds to a particular time instant. The Doppler frequency f_D is directly related to the target's radial velocity v_r by the equation:

$$v_r = \frac{\lambda f_D}{2}$$

where $\lambda = \frac{c}{f_c}$ represents the wavelength of the transmitted signal, with c as the speed of light and f_c the carrier frequency. By converting the Doppler frequency axis to velocity, the Doppler frequency–time matrix is converted into a velocity–time spectrogram, facilitating the temporal tracking of the target's motion. An example of this velocity–time representation is shown in Figure 5.12, offering a clear visualisation of velocity variations over time. Velocities with positive values indicate motion towards the radar, while negative values correspond to motion away from it.

c) Data Processing and Comparison

The synchronisation between the *MoCap* and *Radar* systems enables the straightforward superposition of the y-velocity components, sharing the same y-axis, as illustrated in Figure 5.14, of the markers onto the spectrogram, as shown in Figure 5.17.

Through analysis of the superimposed signals, the torso velocity can be identified as corresponding to the centroid of the spectrogram. The velocities of the feet, marginally greater than those of the ankles, are evident along the upper envelope during motion towards the *Radar* and along the lower envelope when moving away from it. This differentiation allows for accurate segmentation of individual body part velocities based on movement direction. It should be noted that the specific body part represented by the envelope may vary depending on the activity.

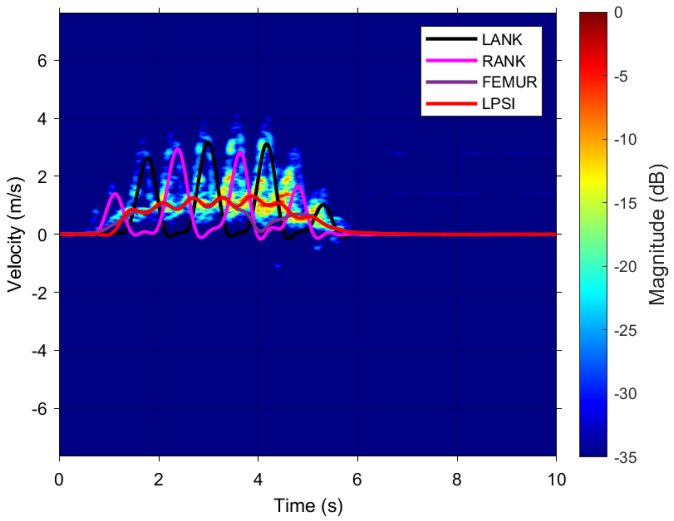

Figure 5.17 – Overlaid Spectrogram of *Walk* and *MoCap* Data (Velocity-Time Graph)

For each activity, velocities from markers placed on the left (LANK) and right (RANK) ankles, as well as on the front right (FEMUR) and rear left of the belt (LPSI), were superimposed. These markers, illustrated in Figure 5.13b, were employed to compare the ground-truth data from *MoCap* with *Radar*-derived measurements. This is attributable to the low radial velocities of such movements relative to the *Radar*, which hinders meaningful comparisons between *Radar*-based estimates and *MoCap* measurements. In future studies, transforming *MoCap* data from its Cartesian coordinate system (x, y) into the *Radar*'s polar coordinate system (r, θ) will allow for more relevant spatial alignment and comparison.

5.2 Preparation and Creation of a Multimodal Dataset: MAD

Chapter 4 demonstrated our ability to detect and recognise activities using *Radar*-based classification models, while highlighting the scarcity of existing data. Section 5.1 described the equipment and our capacity to extract project-relevant data, presenting a complete chain from acquisition to learning and activity recognition. To address these limitations, we propose creating a new dataset balanced across *Activity*, *Age*, and *Gender*, outlining procedures for public involvement, the *Ethics Committee* application, and the study protocol, followed by a description of the dataset's characteristics.

5.2.1 Ethical and Organisational Preparation of the Study

Given the limitations of existing datasets, we created *MAD* (*Multimodal Activity Dataset*) to provide a larger volume of high-quality data, a robust protocol, and flexibility for additional labels, ensuring comparability with existing datasets and adaptability to future research. Participants performed prescribed tasks under an *Ethics Committee* approved protocol.

The project builds on collaboration with Alexandre Bordat [23] on room preparation and protocol design, and with Laure Acin [3] on the *Event-based camera*. This joint work meets multiple sensor requirements and produces a multimodal dataset supporting several theses within the laboratory.

Details on ethical submission, team organisation, participant recruitment, selection criteria, and session planning are in Appendix A, with the main points summarised below:

- **Ethical Approval, Appendix A.1:** The dataset creation project was approved by the *Ethics Committee of CY Cergy Paris Université (CYU)* on 11th October 2023. The approval covered all procedures related to participant recruitment, data acquisition, storage, and processing, ensuring compliance with *General Data Protection Regulation* (GDPR) regulations and institutional standards. The formal letter of approval is shown in Appendix A.6.
- **Organisation, Appendix A.2:** The acquisition campaign involved three doctoral projects and a nine-person team. A rotation system was implemented to conduct two sessions per day, ensuring technical reliability and participant safety.
- **Communication, Appendix A.3:** A structured communication strategy was implemented to inform potential participants about the study. This included institutional mailing lists, posters, newsletters, and a dedicated website. The communication aimed to reach a diverse population and encourage balanced participation across *Age* and *Gender* groups.
- **Recruitment, Appendix A.4:** Participants registered via the *Evento* platform, providing demographic information for pre-selection. 100 participants were recruited via institutional communication channels, ensuring balanced representation across *Gender* (*Male*/*Female*) and *Age* groups (*18–40* and *41–60* years).
- **Planning, Appendix A.5:** Recording sessions took place from January to July 2024, with each session lasting approximately two hours for participants and two and a half hours for the acquisition team.

5.2.2 Experimental Planning and Protocol

Prior to welcoming participants, we developed a detailed protocol. This protocol outlined the equipment to be used, the list of human activities to be observed, and the procedure for each recording session. The complete protocol was approved by the *Ethics Committee*.

a) Activity Selection and Labeling

We defined the activities to generate data comparable with the state of the art, summarised in Table 3.4, while capturing context-specific insights. The *TUG* test was selected as most relevant for analysing weakness and fall risk in elderly or hospitalised individuals, comprising 4 fundamental activities: "*Walk*", "*Semi-Turn Around*", "*Stand Up*", and "*Sit Down*". Participant orientation relative to the *Radar* was considered, producing subcategories for each activity: "*Toward*" the *Radar*, "*Backward*", and "*Right to Left*".

Table 5.1 – Description of *MAD* Activity List

Activities	**Toward**	**Backward**	**Right to Left**	**Diagonal**	**TUG**	**Arbitrary**	**Material**	**Recording Activation**
Walk	011	012	013	014	-	101	-	Remote Control
Walk (Clinical test)	-	-	-	-	015	-	Chair	Pressure Sensor
Stand	021	022	023	-	-	-	Chair	Pressure Sensor
Sit	031	032	033	-	-	-	Chair	Remote Control
Up Stairs	041	042	043	-	-	-	Steps	Remote Control
Down Stairs	051	052	053	-	-	-	Steps	Remote Control
Pick	061	062	063	-	-	-	Basketball	Remote Control
Step Over	071	072	073	-	-	-	Hurdle	Remote Control
Semi-Turn Around	081	082	083	-	-	-	Hurdle	Remote Control
Fall	091	092	093	-	-	-	Mat	Remote Control

Comparison with existing datasets, such as [57], guided activity selection. Only certain activities ("*Pick*" and "*Fall*") were relevant for assessing fall risk and lower limb weakness, while others (e.g., "*Drink*") were excluded. Additional activities of interest included "*Up stairs*", "*Down stairs*", and "*Step Over*" obstacles, as they involve effort and are frequently associated with falls.

Each activity is performed in three repetitions per subcategory. The "*Walk*" activity includes two additional subcategories, "*Diagonal*" and "*Arbitrary*" walking, reflecting natural gait variations, and also incorporates the full *TUG* test. In total, *MAD* comprises 9 primary activities with three subcategories each, except "*Walk*" which has 6, all repeated 3 times to enhance robustness. Unique identifiers for each activity are summarised in Table 5.1.

b) Experimental Setup

One of the initial steps involved measuring the room to prepare the protocol and the *Ethics Committee* file, which allowed us to identify the active detection zone and usable space.

The 2024 configuration, including floor markings and final sensor positions, is illustrated in Figure 5.18. Floor markings indicate start and end points of activities, maximum and minimum room dimensions, and reference points for MoCap calibration and skeletal reconstruction.

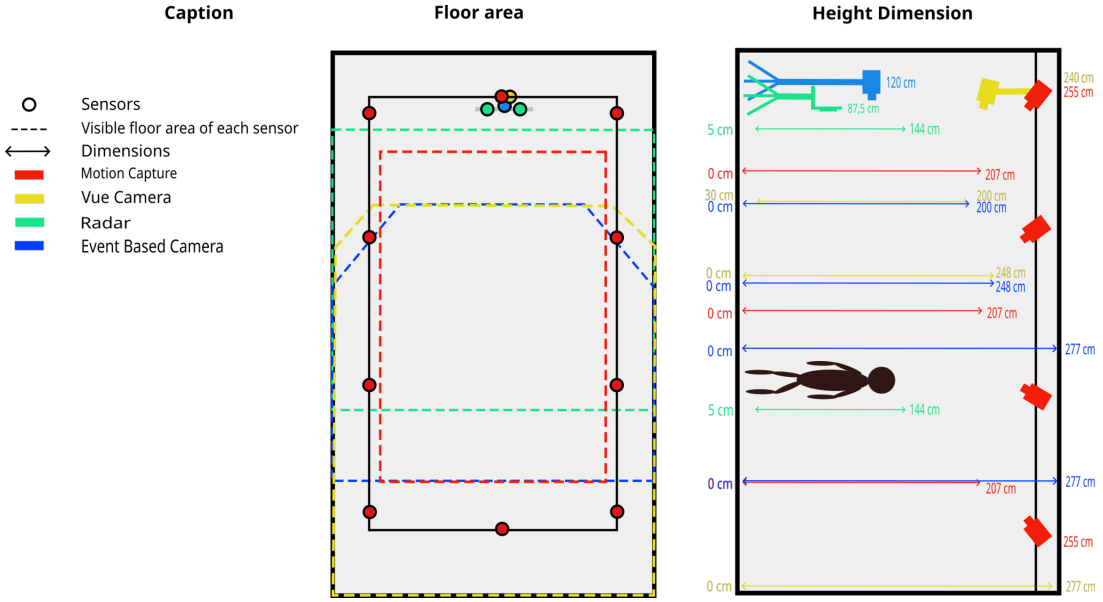

Figure 5.18 – *Movement Platform*'s Dimensions Measured in January 2024

To integrate the *Event-Based camera*, a fully wireless solution was implemented. An Arduino on the *Radar/Event-Based Camera* side receives a $433MHz$ trigger signal generated when the participant stands or via a remote control, as illustrated in Figure 5.19. A high-impedance switch enables safe voltage adaptation between the *Radar* ($3.3V$) and

the camera ($5V$). Upon receiving the trigger, both devices start recording simultaneously; the *Radar* controls the $10s$ recording duration and sends a stop signal to the camera.

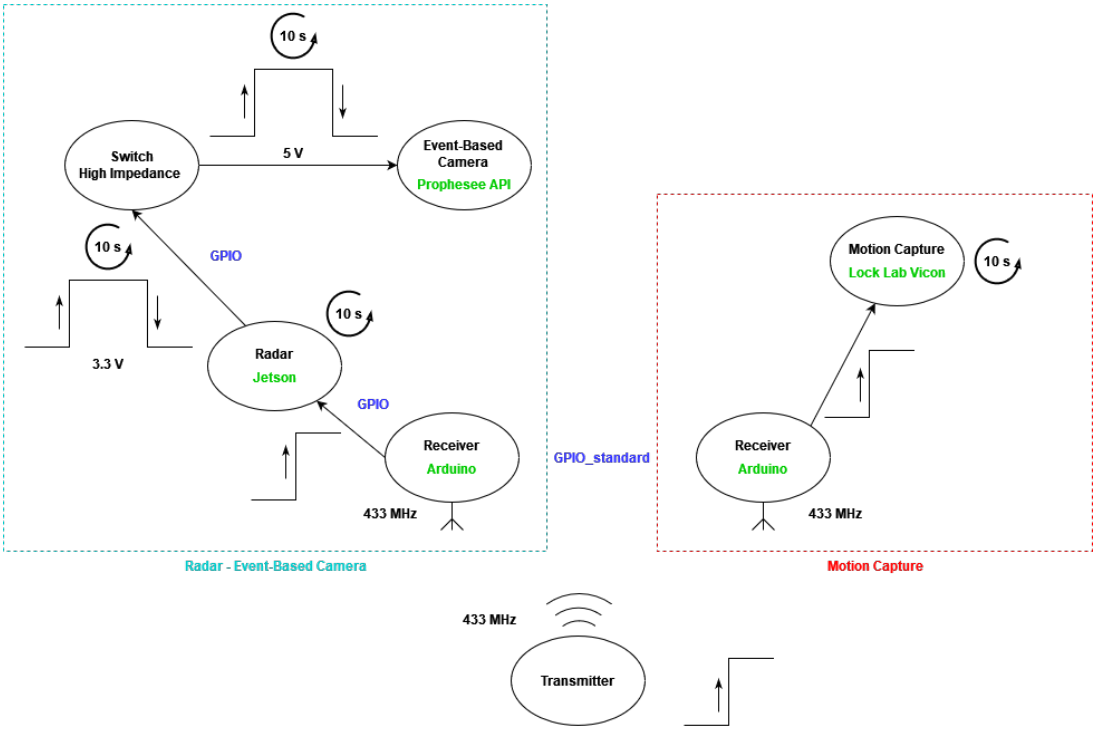

Figure 5.19 – Sensor's Configuration for Synchronisation

All preparatory measurements and initial configurations from 2023, minor calibration adjustments, and simulations conducted to refine sensor placement are described in Appendix B.1.

c) Experimental Accessories

Additional equipment handling, placement optimisation, and accessory simulations are detailed in Appendix B.2. A chair is used for "*Sit*" and "*Stand*", hurdles for "*Step Over*", steps for "*Up stairs*" and "*Down stairs*" activities, mats for "*Fall*", and a basketball for "*Pick*".

d) Session Schedule and Participant Protocol

Sessions were structured to balance physical load and ensure consistency. Participants began with "*Walk*" subcategories, followed by "*TUG* test", "*Up stairs*", "*Down stairs*",

"*Pick*", optional "*Arbitrary Walk*", and finally "*Fall*". Informed consent, equipment preparation, calibration, and end-of-session surveys were completed. The estimation duration is 2 hours.

Operational details for warm-up, hurdle adjustments per participant, trigger handling, and file verification are reported in Appendix B.3.

5.2.3 *MAD* Dataset: Management, Processing and Content

Throughout the participant recruitment process and the design of the datset, a data backup system was implemented. Its purpose is to centralise the data for the research team while ensuring its security within the laboratory.

a) Data Management and Storage

A centralised and secure data storage system was implemented to manage all session files and participant information. Full details of its operation are provided in Appendix B.4.

b) Data Estimation and Structuring

Each activity comprises three *subcategories* (*Toward*, *Backward*, and *Right–Left*), except for "*Walk*", which has 6. The minimum number of recorded subcategories per participant is therefore $8 \times 3 + 6 = 30$, resulting in $30 \times 3 = 90$ activity files per participant.

These files are organised by sensor as follows:

- **Radar:** produces three types of data (raw data, instantaneous distance images, and spectrograms), generating four files per *subcategory*, for a total of 450 files per participant.
- **Event-Based Camera:** records two files per *subcategory* (captured data and configuration), yielding 180 files per participant.
- **MoCap:** generates one *.csv* file per *subcategory*, totalling 90 files per participant.
- **Conventional RGB Camera:** records one video per *subcategory*, adding 90 files per participant.

The total number of files per participant is $450 + 180 + 90 + 90 = 810$, and for 100 participants, the expected volume is approximately 81,000 files.

Given this scale, a clear naming convention and structured storage system are essential. The hierarchy must enable quick identification by participant, activity, subcategory, and

sensor, ensure traceability across sensors, and support automated processing and future dataset expansion.

The nomenclature adopted is therefore: AY1Y2PXXXRZSUDV where:

- Y1 is the code for the activity [1–9], and Y2 is the code for the subcategory of the activity [1–5] in the context of this dataset.
- XXX is the participant identifier [1–100];
- Z is the repetition number [1–3];
- U is the code of the sensor used [1–5], where 1: eb_cam, 2: cam, 3: mat, 4: mocap, and 5: radar;
- V is the code of the antenna used [0–4]. This code is 0 if the sensor is not the radar or if raw radar data are being used; otherwise, the use of the receiving antenna Rx1 or Rx2 must be specified for the radar. For the radar, the mapping is as follows: radial distance image Rx1: 1, radial distance image Rx2: 2, spectrogram image Rx1: 3, and spectrogram image Rx2: 4.

For example, for participant 63, who performed activity 2 and sub-activity 1 for the second repetition, the data recorded by *MoCap* will have the following file name: A21P063R2S4D0.csv. The data tree is structured as illustrated in Figure 5.20, where user_id corresponds to the identifier of the participant.

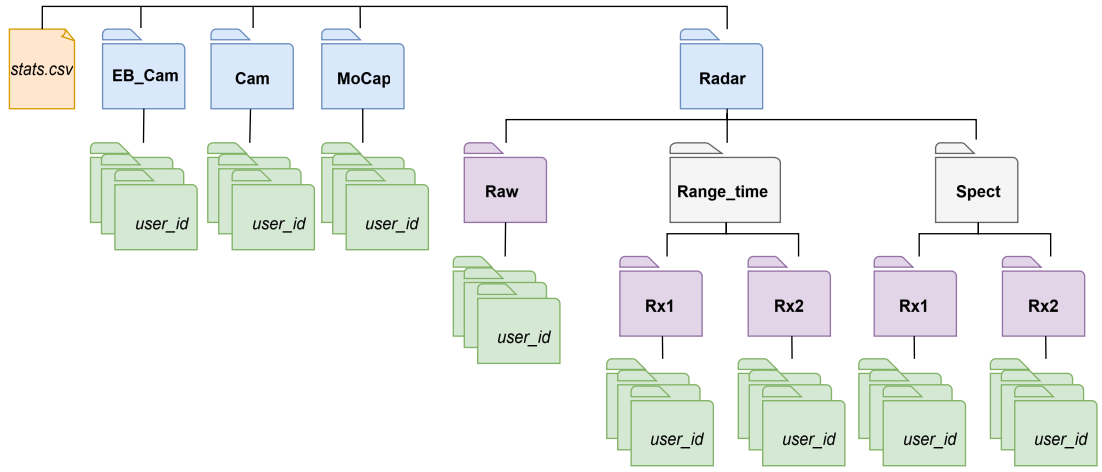

Figure 5.20 – *MAD* Data Tree Structure

Once the server and *API* were fully operational for data submission and participant form completion, comprehensive tests were conducted to validate the system. Following collaboration between engineers and doctoral students, the platform was stabilised,

and the *API* code was distributed to the team to allow modifications. This enabled the resolution of connection issues, *DNS* errors, format incompatibilities, and file transfer limitations.

c) Dataset Statistics

We successfully recruited 100 participants with a balanced distribution across *Age* and *Gender* categories, as summarised in Table 5.2. This unique cohort facilitates fair comparisons in classification analyses and provides a real-world basis for further studies. Although recorded in the laboratory, the data allow advanced testing of clinical measures and could be extended to broader applications.

Table 5.2 – Participant Distribution

Age	**18-40**	**41-60**	**Total**
Female	25%	25%	**50%**
Male	25%	25%	**50%**
Total	**50%**	**50%**	**100%**

Key participant characteristics, including *Body Mass Index* (BMI), physical activity levels, and sedentariness are provided in Appendix B.5.

d) Dataset Content

Our dataset *MAD* contains data from 100 participants evenly distributed across *Age* and *Gender* categories. Each participant performed a set of activities in various directions: B (*Backward*), T (*Toward*), RL (*Right–Left*), and D (*Diagonal*), as summarised in Figure 5.21.

Figure 5.22 presents a spectrogram mosaic illustrating the *Radar* data for the main activities: "*Walk*", "*Stand*", "*Sit*", "*Up and Down Stairs*", "*Pick*", "*Step Over*", "*Semi Turn Around*", and "*Fall*". Velocities from markers placed on the left (LANK) and right (RANK) ankles, as well as front-right (FEMUR) and rear-left belt (LPSI), shown in Figure 5.13b, were superimposed to compare *MoCap* ground truth with *Radar* data. Left-to-right movements are included in the mosaic but not overlaid with *MoCap* data due to their low radial velocity. Future work could transform *MoCap* Cartesian coordinates (x, y) into the *Radar* polar system (r, θ) for improved alignment.

The relative error for each peak in the "*Walk*" activity is defined as the difference between the maximum velocity peaks measured by the *Radar* and the reference *MoCap*

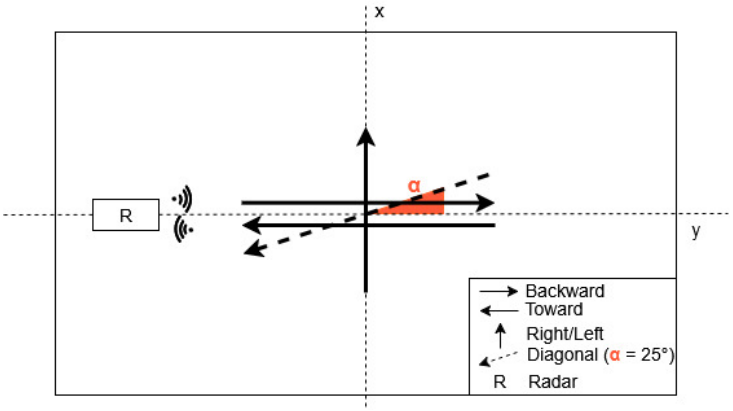

Figure 5.21 – Top-View Schematic of Activity Directions within the *Movement Platform*

system, normalised by the *MoCap* peak value:

$$e_i = \frac{|v^{\max}_{\text{radar},i} - v^{\max}_{\text{MoCap},i}|}{v^{\max}_{\text{MoCap},i}}, \tag{5.2}$$

where $v^{\max}_{\text{radar},i}$ and $v^{\max}_{\text{MoCap},i}$ denote the maximum velocities of the i-th peak detected by the *Radar* and *MoCap*, respectively. The mean relative error over all peaks is then computed as:

$$E = \frac{1}{N}\sum_{i=1}^{N} e_i, \tag{5.3}$$

with N being the total number of peaks. Using the "*Walk*" recordings, which provide the most extensive data relevant to fall risk assessment, an average relative error of 9.23% was obtained. In addition, *MoCap* provides a confidence measure to assess the reliability of *Radar* data and the presence of noise.

For quality assessment, we compared our dataset to *Radar Signatures of Human Activities* [57], in which 13% of the images are considered problematic, meaning that the activity is either not discernible due to excessive noise or the recording contains errors such as sensor glitches or incomplete captures. Restricting *MAD* to two subcategories, "*Toward*" and "*Backward*", we observed only 12 missing data points and 85 noisy images out of 5,400, leaving 98% of the data usable (5,303 images).

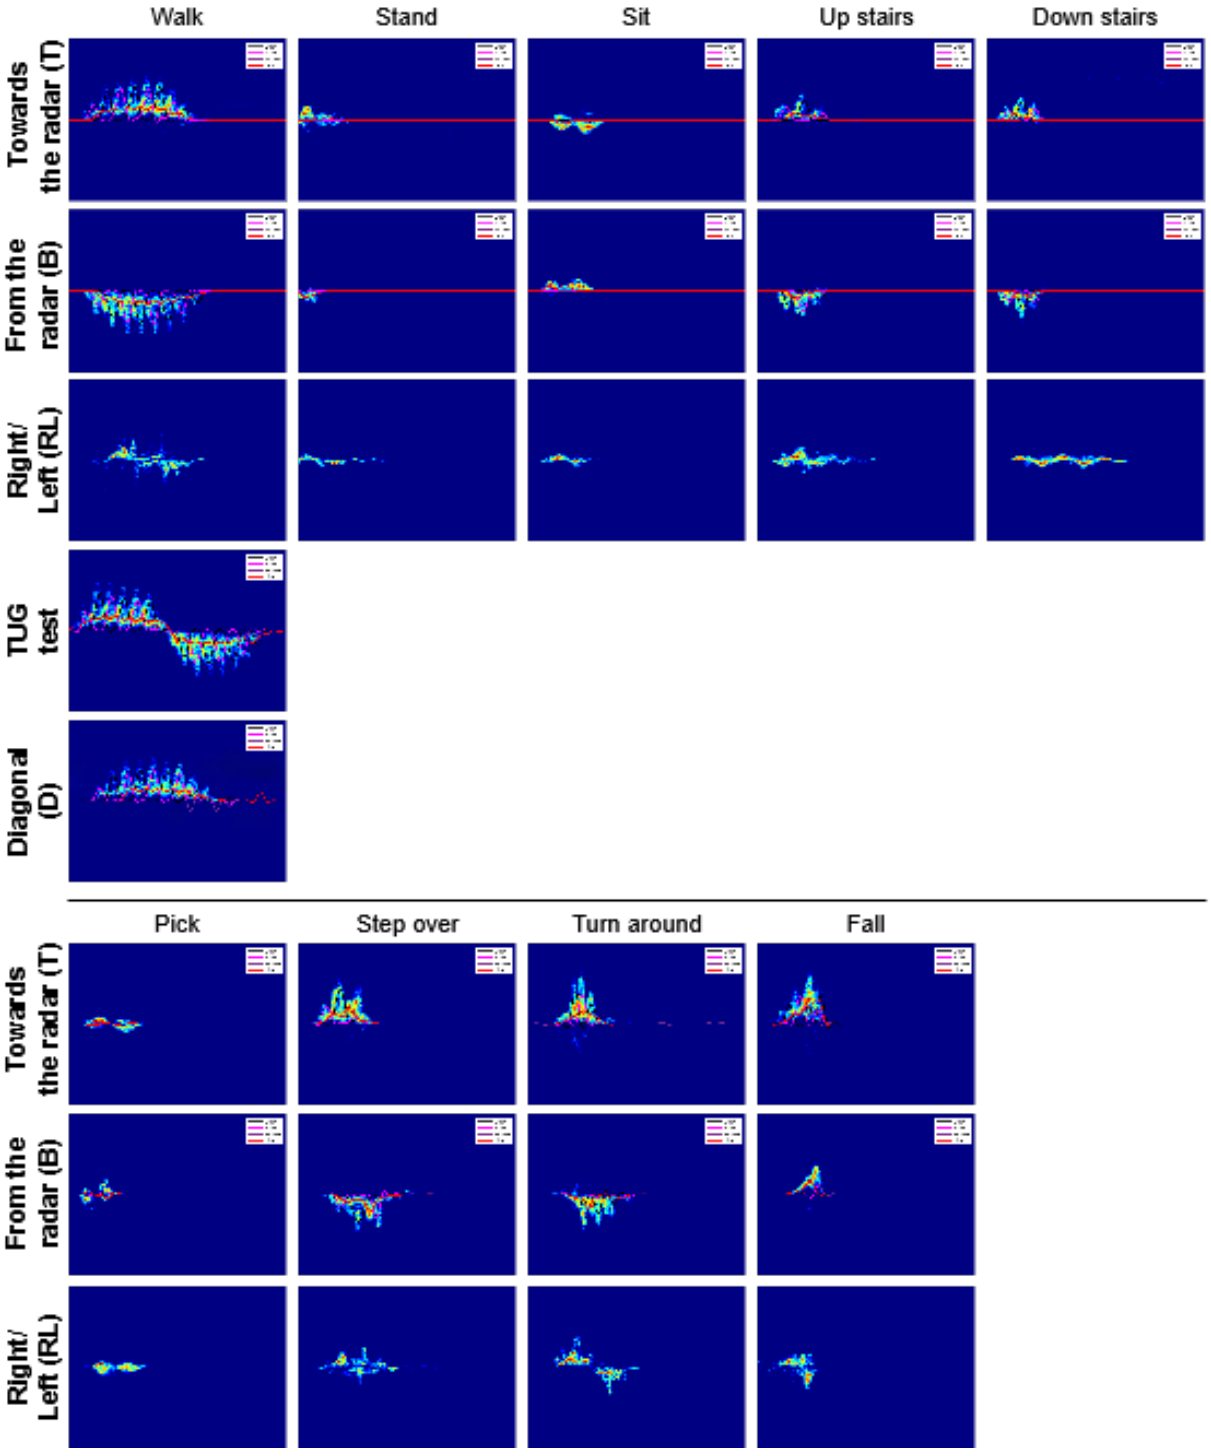

Figure 5.22 – Mosaic of our 9 Activities, Performed in All Directions (T, B, RL, D)

5.3 Evaluation of Models and Pre-Processing Strategies on the MAD Dataset

The newly created dataset enables evaluation of the performance and robustness of our approach under conditions close to real-world scenarios. Using tools similar to those and methods equivalent to those employed in the *Radar Signatures of Human Activities* dataset [57] (Chapter 4), we first analyse 5 comparable activities before extending to all 9 recorded activities. Performance is assessed with *ResNet-18*, *SVM*, and *KNN*, at two classification levels: initially with two subcategories (*Toward* and *Backward*) and then including three (*Toward*, *Backward*, and *Right–Left*) to evaluate the impact of lateral movements. Finally, majority-vote fusion is applied and its performance compared with results from Section 4.3.4.

5.3.1 Data Description

In total, *MAD* contains 8,874 *spectrograms* recorded from one of the two receiving antennas. Although data were collected from both antennas, only a single antenna is considered in this thesis, as *interferometric* processing is not applied. Further details on the setup can be found in Section 5.1.1. The spatial distribution of the data within the room, illustrated in Figure 5.14, is summarised in Table 5.3.

Table 5.3 – Dataset Overview

Activity Name	**Route**[1]	**Number of Samples**[2]
Walk	$A \rightarrow C$	T(298) B(300) RL(298) D(300)
Stand	B	T(300) B(300) RL(300)
Sit	B	T(300) B(300) RL(300)
Up Stairs	$B \rightarrow C$	T(300) B(299) RL(300)
Down Stairs	$B \rightarrow C$	T(299) B(300) RL(299)
Pick	B	T(298) B(300) RL(300)
Step Over	$B \rightarrow C$	T(299) B(300) RL(295)
Semi-Turn Around	Around B	T(299) B(299) RL(300)
Fall	B	T(298) B(299) RL(300)
TUG test	NA	300
Arbitrary Activity	NA	192

[1] NA: Not Applicable

[2] T: Towards, B: Backwards, RL: Right/Left, D: Diagonal

To enable comparison with the state of the art, we define two sub-datasets: *MAD5*, containing 5 activities comparable with the *Radar Signatures of Human Activities* dataset [57], and *MAD9*, containing all 9 recorded activities. The characteristics of these sub-datasets are summarised in Table 5.4.

Table 5.4 – MAD5 and MAD9 Specifications

SubSet	**Activities**	**Subcategories**
MAD5 2sscateg	*Walk*, *Stand*, *Sit*, *Pick*, *Fall*	Toward, Backward
MAD5 3sscateg		Toward, Backward, Right-Left
MAD9 2sscateg	*Walk*, *Stand*, *Sit*, *Up Stairs*, *Down Stairs*, *Pick*, *Step Over*, *Semi-Turn Around*, *Fall*	Toward, Backward
MAD9 3sscateg		Toward, Backward, Right-Left

For both sub-datasets, training is performed either using two subcategories, "*Toward*" and "*Backward*", or the three common subcategories, "*Toward*", "*Backward*", and "*Right-Left*". The "*Diagonal*", "*Arbitrary*", and "*TUG*" subcategories are excluded, as they overlap multiple activity categories.

Table 5.5 – Activity Distribution for 2 and 3 Subcategories Classification

(a) Towards and Backwards

Id	**Activities**	**Train**	**Test**
1	Walk	509	89
2	Stand	510	90
3	Sit	510	90
4	Up Stairs	510	89
5	Down Stairs	510	89
6	Pick	509	89
7	Step Over	510	89
8	Semi-Turn Around	509	89
9	Fall	508	89

(b) Towards and Backwards and Right-Left

Id	**Activities**	**Train**	**Test**
1	Walk	762	134
2	Stand	765	135
3	Sit	765	135
4	Up Stairs	765	134
5	Down Stairs	764	134
6	Pick	764	134
7	Step Over	760	134
8	Semi-Turn Around	764	134
9	Fall	763	134

Training set distributions for two subcategories are summarised in Table 5.5a, and for three subcategories in Table 5.5b. With two subcategories, *MAD* contains 5,400 spectrograms for a single antenna, approximately three times more data than [57]. Moreover,

the activity distribution is balanced, with around 600 images per activity, compared with 200–300 in the previous dataset, which was unbalanced.

5.3.2 Evaluation of Model Performance and Pre-Processing Strategies on the MAD Dataset

One of our main objectives is to assess whether our dataset can effectively support our learning and data processing methods. As part of this effort, Alexandre Bordat [23] has been evaluating several filtering techniques in his pre-processing chain, providing multiple versions of the generated spectrograms.

The noise level within the pre-processing chain can be tuned to control spectrogram quality. In the *Radar Signatures of Human Activities* dataset [57], it is fixed at $50dB$. In our case, the initial configuration used a lower threshold of $35dB$ combined with a *Finite Impulse Response* (FIR) filter.

a) Optimisation of Pre-Processing Chain Parameters

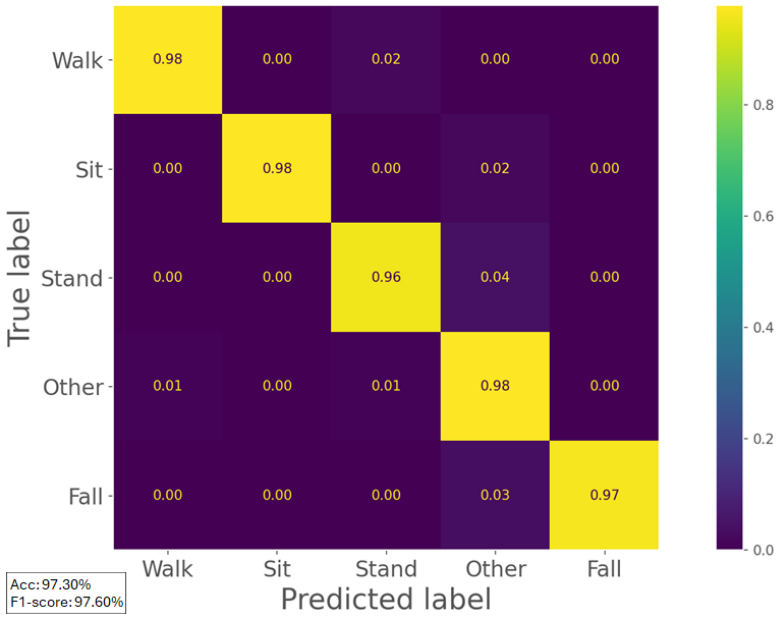

Figure 5.23 – Normalised Confusion Matrix of *Radar Signatures of Human Activities* Dataset [57] as Input to our Pre-Processing Chain

To validate the use of the pre-processing chain for generating spectrograms, we re-generated the data from the *Radar Signatures of Human Activities* dataset [57] using our

own chain. With a noise level of $35dB$, the performance obtained is higher than that of the state-of-the-art data, approximately 10% greater than that achieved with their pre-processing chain. The confusion matrix in Figure 5.23 confirms this improvement, showing a clearly defined recognition diagonal.

Subsequently, Alexandre Bordat [23] proposed using the *5-pulse canceller* (5-PC). With this new filtering approach, we conducted tests at noise levels of $30dB$, $35dB$, and $40dB$. Visual differences can be observed between the *spectrograms* shown in Figure 5.24. When using the *FIR* filter, a central line appears in the spectrograms, which is absent with the *5-PC* method. Conversely, the *5-PC* filter reduces background noise at lower noise levels. However, if the noise level is set too low, there is a risk of losing information about the movements.

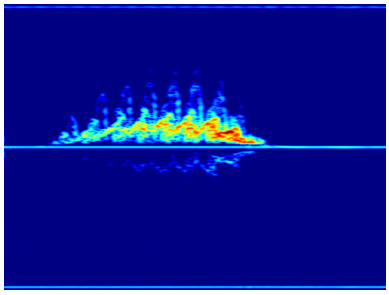

(a) Spectrogram with FIR, $35dB$

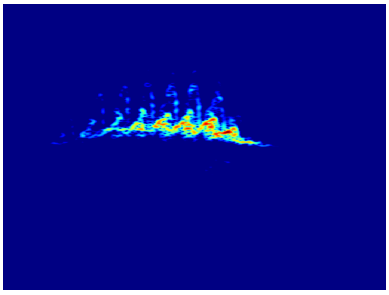

(b) Spectrogram with 5-PC, $30dB$

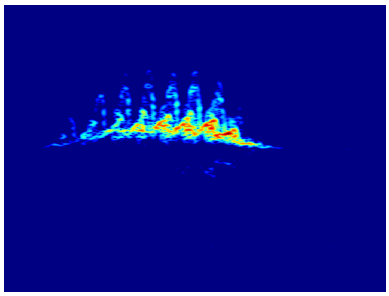

(c) Spectrogram with 5-PC, $35dB$

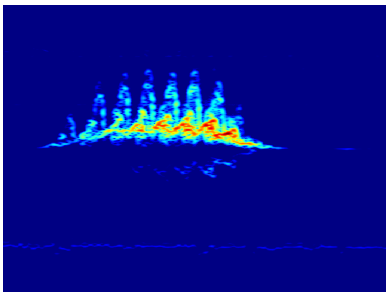

(d) Spectrogram with 5-PC, $40dB$

Figure 5.24 – Spectrograms Generated with Different Filters and Biases

Applying our *ResNet-18* model to the four datasets combining 5 activities and two subcategory configurations confirmed that the *5-PC* filter at $35dB$ provides the best average accuracy with 99.35%, as shown in Table 5.6. Consequently, this configuration was retained for all subsequent experiments.

Table 5.6 – Comparison of Accuracies for Different Biases for MAD5

(a) Comparison for 2 Subcategories

Filter	Bias-Level	Accuracy	F1-score
5-PC	35dB	**99.35%**	**99.20%**
5-PC	40dB	97.92%	**99.40%**
5-PC	30dB	**99.55%**	**99.40%**
FIR	35dB	98.66%	98.60%

(b) Comparison for 3 Subcategories

Filter	Bias-Level	Accuracy	F1-score
5-PC	35dB	**98.99%**	**99.00%**
5-PC	40dB	**99.55%**	98.00%
5-PC	30dB	98.51%	**98.60%**
FIR	35dB	97.32%	97.40%

b) Impact of the Number of Subcategories on Model Performance

We also want to assess whether or not the subcategories have a significant impact on learning accuracy. In Figure 5.25, we see that the confusion matrix obtained for two subcategories is very accurate. Indeed, with an accuracy of 99.35%, the performance becomes very interesting compared to the state of the art.

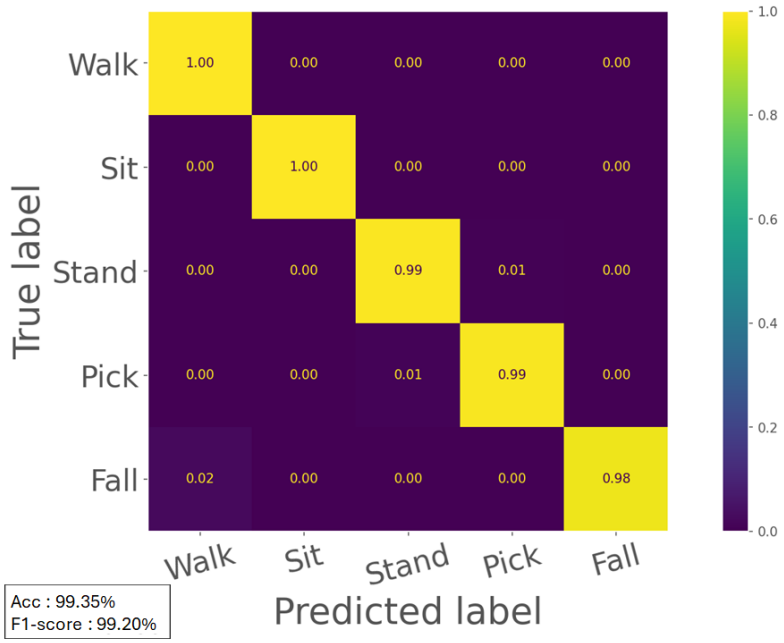

Figure 5.25 – Normalised Confusion Matrix of *ResNet-18* with *MAD5* 2sscateg

If we look at existing work in our field, the Table 3.3 allows us to compare ourselves. We can observe that the methods achieving accuracy rates above 90% rely on specific data processing techniques. With *MAD5* under two subcategories, we already surpass the state of the art, achieving an accuracy of 99.35%.

Similarly, with three subcategories, despite a slight decrease of around 1% in accuracy, the model appears to be only marginally affected by the addition of this third subcategory. The confusion matrix in Figure 5.26 confirms that class confusion remains very low, below 1% for all classes except for "*Pick*", which shows a 3% confusion rate.

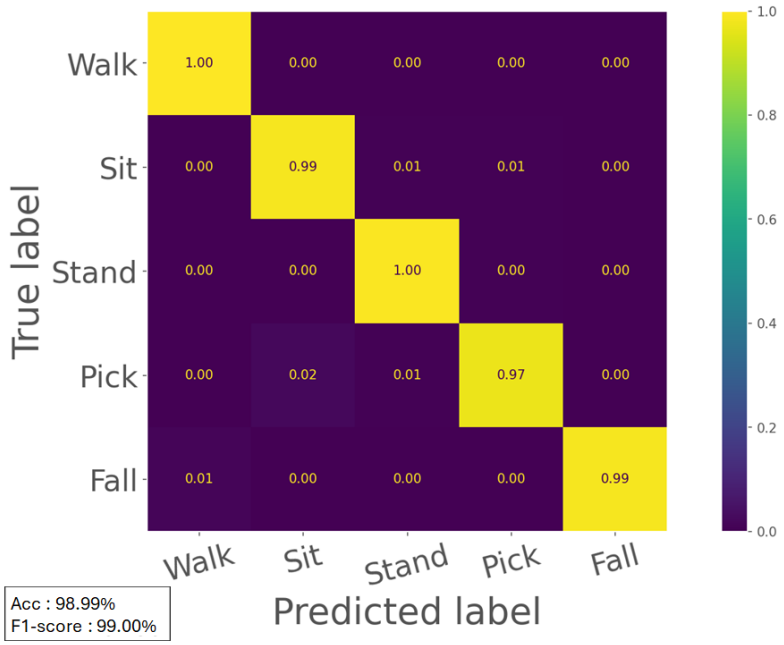

Figure 5.26 – Normalised Confusion Matrix of *ResNet-18* with *MAD5* 3sscateg

The *MAD* dataset comprises nine activities. Although it is not directly comparable with the *state of the art* in its current form, we conducted a preliminary test to assess whether the inclusion of additional activities could affect our activity recognition performance. For the two subcategory configuration, Figure 5.27 shows that an accuracy of 97.91% was achieved.

The confusion matrix in Figure 5.27 indicates that the "*Down Stairs*" activity is the most prone to misclassification, as it is closely related to the "*Up Stairs*" and "*Pick*" activities. This mirrors the similarities previously observed between "*Sit*", "*Stand*", and "*Pick*" in the *Radar Signatures of Human Activities* dataset [57]. This confusion arises from the fact that the motion signatures of these activities share similar movement patterns. Nevertheless, the level of confusion remains insufficient to significantly affect overall performance.

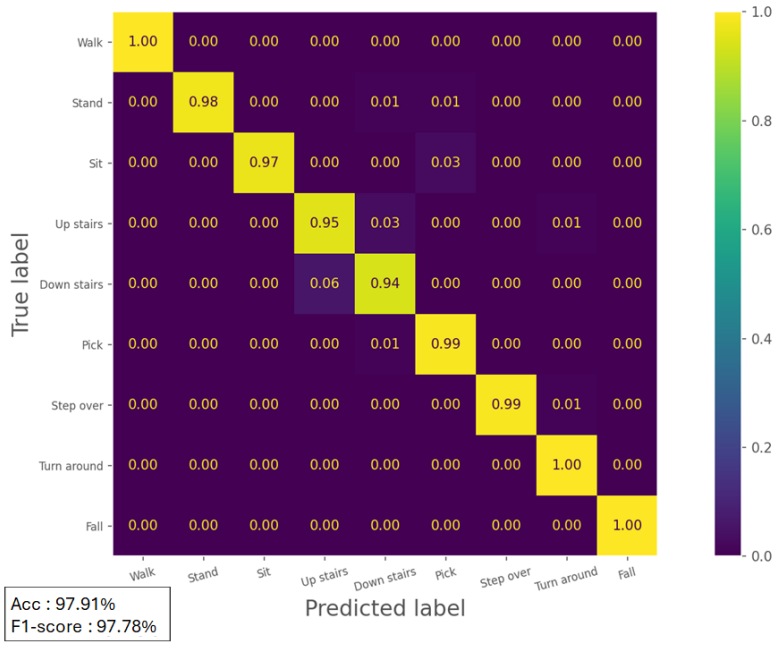

Figure 5.27 – Normalised Confusion Matrix of *ResNet-18* with *MAD9* 2sscateg

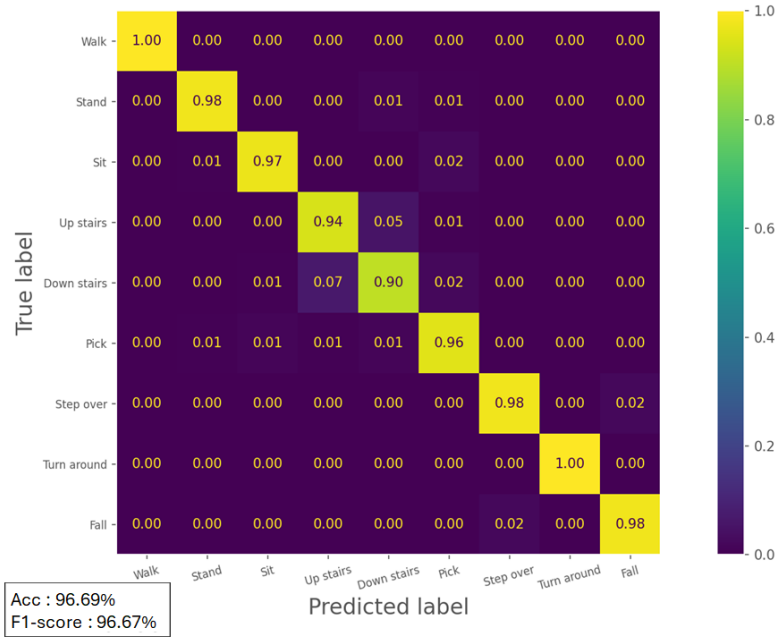

Figure 5.28 – Normalised Confusion Matrix of *ResNet-18* with *MAD9* 3sscateg

For the three subcategory configuration shown in Figure 5.28, the same phenomenon observed with the five activity case can be seen. The overall performance decreases by

less than 1%, reaching an accuracy of 96.69%. Confusion increases slightly for the "*Down Stairs*" activity. However, this remains acceptable, with an accuracy of 90% for this activity.

c) Evaluation of the Contribution of Image Processing Methods

As with the data from the *Radar Signatures of Human Activities* dataset [57], we conducted tests using our *denoising* and *cropping* methods. Table 5.7 shows that the results obtained are comparable to, or in some cases slightly lower than, those achieved with data that did not undergo additional processing. These findings support our hypothesis presented in Section 4.3.1, where we assumed that our methods are not systematically required to enhance the performance of *ResNet-18*.

Table 5.7 – Performance Comparison Between Image Processing Methods

(a) Comparison for 2 Subcategories

Method	**# of Activity**	**Accuracy**	**F1-score**
Denoise	5	**98.25%**	**98.20%**
Denoise	9	**95.83%**	**95.89%**
Crop	5	96.93%	96.80%
Crop	9	95.22%	95.22%

(b) Comparison for 3 Subcategories

Method	**# of Activity**	**Accuracy**	**F1-score**
Denoise	5	**97.02%**	**97.20%**
Denoise	9	**94.95%**	**95.11%**
Crop	5	95.68%	95.60%
Crop	9	91.47%	91.22%

The inclusion of additional activities will be essential in future work, particularly for real-life studies conducted in environments such as nursing homes. However, the primary objective at this stage is to validate our dataset and to understand its limitations in comparison with the state of the art. For all these reasons, we focus on the *MAD5* dataset, processed at $35dB$ with an *5-PC* filter.

5.3.3 Learning Chain: Model Parameter Optimisation and Performance Evaluation

In Section 4.3, we established a processing chain for learning. Specifically, we employed a *ResNet-18* model that takes spectrogram images as input, compressed into 32×32 thumbnails. In parallel, an *SVM* model was configured, for which the most relevant kernels were selected and the corresponding hyperparameters optimised. A *KNN* model was also tuned to determine the most appropriate hyperparameters. Finally, a majority-voting *fusion* method was implemented.

The results obtained with the *Radar Signatures of Human Activities* dataset [57] were highly promising. Consequently, this section applies the same procedure, enabling a direct comparison of performance with the results obtained using our own dataset.

a) ResNet-18 Parameters and Performance

For the remainder of the tests, we chose to retain the *ResNet-18* model trained on spectrograms of five activities, without any additional processing beyond the established pre-processing chain. With a noise level of $35dB$ and the *5-PC* filter, this model achieves an accuracy of 99.35%, which we consider highly satisfactory when compared with the state of the art. The corresponding confusion matrix is presented in Figure 5.25.

b) SVM Parameter Selection and Accuracy

For the *SVM*, we retain the same hyperparameters as those identified in Section 4.3.2, namely $coef0 = 1.0$, $degree = 4$ and $C = 100$. The objective is to ensure comparable configuration conditions with those used for the tests performed on the *Radar Signatures of Human Activities* dataset [57].

We implemented a method for automatically identifying the most optimal parameters through a combination search and *PCA*, as described in Section 4.3.2. For *MAD5*, the optimal parameters obtained are identical with and without the application of *PCA*. These parameters: *Area*, *Perimeter*, *Orientation*, *Major*, *Minor*, *Centroid X*, and *Centroid Y*, are summarised in Table 5.8.

Table 5.8 – *SVM* Performance Comparison Between Combination Search and *PCA* Methods

Method	**# of Activity**	**# of Subcategory**	**Kernel**	**Optimum Parameters**	**Accuracy**	**F1-score**
Combination	5	2	Polynomial	"Area", "Perimeter", "Orientation", "Major", "Minor", "Centroid X", "Centroid Y"	80.03%	79.20%
PCA	5	2	Polynomial	"Area", "Perimeter", "Orientation", "Major", "Minor", "Centroid X", "Centroid Y"'	80.03%	79.20%

Once these optimal parameters were selected, the learning process yielded the corresponding confusion matrix shown in Figure 5.29. Compared with the results obtained using the *Radar Signatures of Human Activities* dataset [57], overall confusion is approximately 8% higher. In particular, the "*Pick*", "*Sit*", and "*Stand*" activities exhibit greater

similarity in their signature patterns. This level of confusion is entirely consistent with the inherent definitions and decomposition of these activities.

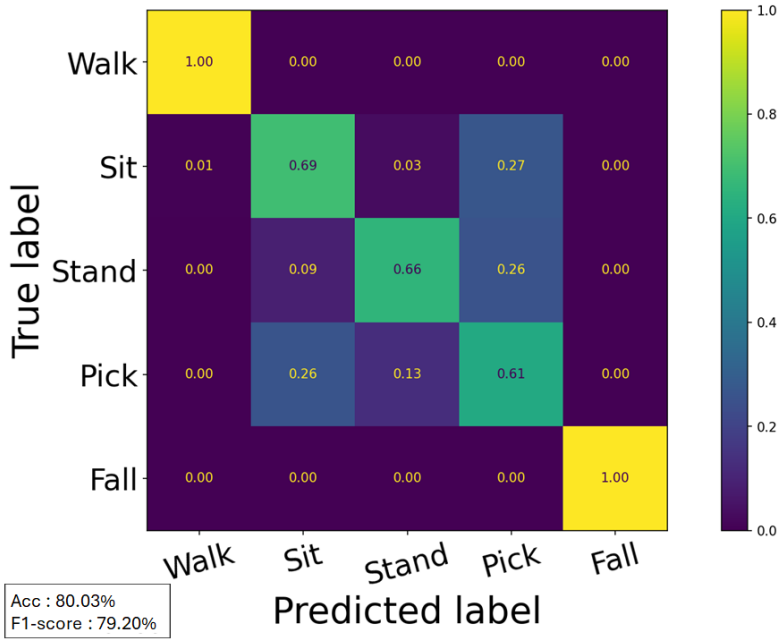

Figure 5.29 – Normalised Confusion Matrix of *SVM* with *MAD5* 2sscateg

Although this confusion is greater than that observed with the *Radar Signatures of Human Activities* dataset [57], the main diagonal of the confusion matrix remains clearly defined, with recognition rates exceeding 60% for the three most problematic categories. Overall, the model still achieves an accuracy of 80%.

c) KNN Parameter Selection and Accuracy

As with the *SVM*, we retain the hyperparameters identified in Section 4.3.3 to ensure methodological consistency. As summarised in Table 4.13, the *KNN* model employs $weights = distance$, $metric = manhattan$, and K-Neighbours $= 22$.

Similarly, we implemented a method for automatically identifying the most optimal parameters through a combination search and *PCA*, as described in Section 4.3.3. For *MAD5*, the optimal parameters obtained are identical across both methods. These parameters: *Dissimilarity*, *Correlation*, *Contrast*, and *ASM*, are summarised in Table 5.9.

We also applied an automated parameter optimisation combining grid search and *PCA*, as described in Section 4.3.3. For *MAD5*, both methods yielded identical optimal parameters: *Dissimilarity*, *Correlation*, *Contrast*, and *ASM*, summarised in Table 5.9.

Table 5.9 – *KNN* Performance Comparison Between Combination Search and *PCA* Methods

Method	# of Activity	# of Subcategory	Optimum Parameters	Accuracy	F1-score
Combination	5	2	"Dissimilarity", "Correlation", "Contrast", "ASM"	87.01%	82.60%
PCA	5	2	"Dissimilarity", "Correlation", "Contrast", "ASM"	87.01%	82.60%

Training with these parameters produced the confusion matrix shown in Figure 5.30, which reveals notable confusion between the "*Pick*", "*Sit*", and "*Stand*" classes. This behaviour is consistent with observations made for *SVM*, *ResNet-18*, and in the *Radar Signatures of Human Activities* dataset [57].

Despite this, the model trained on *MAD5* achieves a markedly higher overall accuracy of 87.01%, compared with 67.18% for the *Radar Signatures of Human Activities* dataset, representing an improvement of nearly 20%.

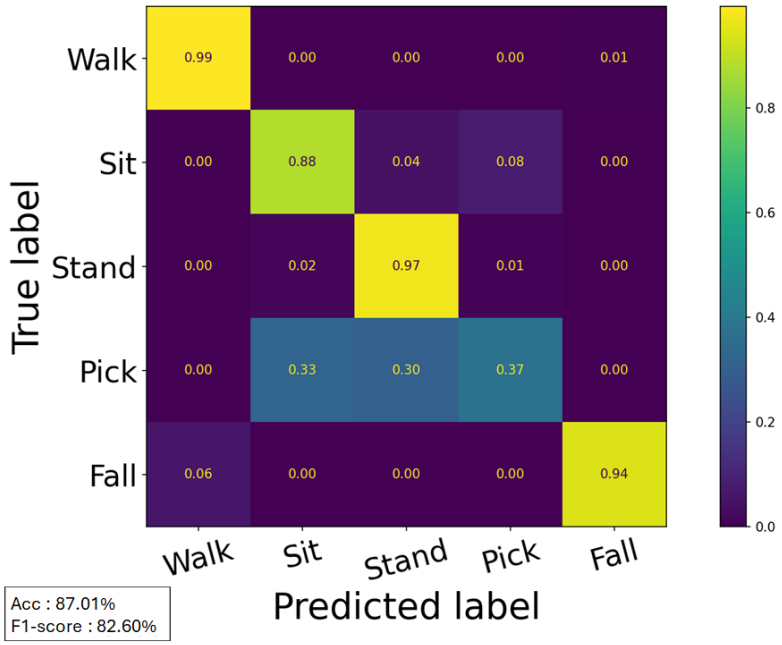

Figure 5.30 – Normalised Confusion Matrix of *KNN* with *MAD5* 2sscateg

Overall, all three classifiers exceed 80% accuracy on *MAD5*, whereas only two models reached similar performance in the reference dataset, confirming the robustness of our approach.

5.3.4 Fusion of Models by Majority Vote

Section 4.3.4 presents the complete processing chain designed to handle cases where the model may be either overconfident or uncertain. In particular, we implemented a *majority-voting* scheme combining the three *classification models*.

Table 5.10 – Summary of the Highest Accuracies Achieved on the *MAD5* 2sscateg Dataset

Model	Accuracy	F1-score
ResNet-18	99.35%	99.20%
SVM	80.03%	79.20%
KNN	87.01%	82.60%
Fusion	95.94%	95.00%

Using the *MAD5* sub-dataset at $35dB$ with two subcategories, we obtained highly promising results. As summarised in Table 5.10, *ResNet-18* achieved 99.35% accuracy, *SVM* 80.03%, and *KNN* 87.01%. These results contrast with those obtained on the *Radar Signatures of Human Activities* dataset [57], where model performance was more uneven.

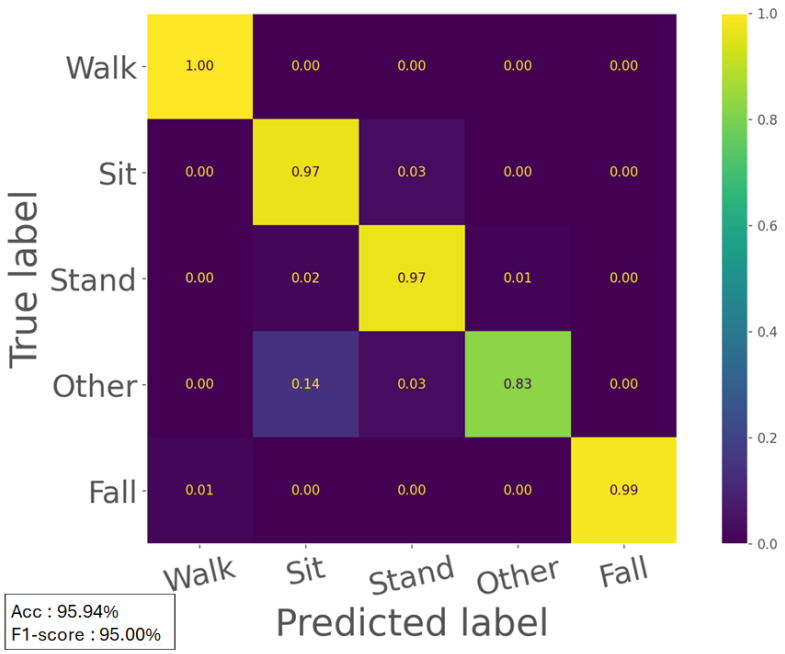

Figure 5.31 – Normalised Confusion Matrix of *Fusion* with *MAD5* 2sscateg

The fusion principle aims to combine models that are both accurate and complementary. When individual classifiers make distinct predictions, *majority voting* helps correct

isolated errors and improves robustness. Here, *ResNet-18* provides the strongest individual performance (99.35%), while *SVM* and *KNN* offer complementary decision patterns, as reflected in their respective F1-scores (79.20% and 82.60%).

Applying the *majority-voting* scheme yields a fused model with 95.94% overall accuracy and a 95.00% F1-score, as reported in Table 5.10. This confirms that the combination effectively mitigates individual misclassifications, as shown by the balanced confusion matrix in Figure 5.31.

Although the current fusion relies on Top-1 predictions, extending it to Top-N strategies could further improve performance. Overall, the high and consistent scores demonstrate that *majority voting* provides a reliable and robust method for combining these models on the *MAD5* dataset.

5.4 Conclusion

To address the *limitations* of existing *radar datasets*, particularly in the clinical context of *fall-risk detection*, we developed and experimentally validated a new dataset: the *Multimodal Activity Dataset* (MAD). Designed to ensure the *reliability* and *quality* required for clinical use, this dataset was collected using an acquisition platform integrating a *radar* and a *MoCap* system, enabling synchronised multimodal recordings. The study was approved by an *Ethics Committee*, ensuring full compliance with ethical standards.

The *MAD* dataset comprises 9 daily-life activities that can reveal *early signs of frailty*, and allows direct comparison with the *Radar Signatures of Human Activities* dataset [57], while improving data quality and consistency. The *learning and optimisation chain* presented in Chapter 4 was adapted to ensure a coherent evaluation across both datasets.

Experimental results confirmed the robustness of our approach. The *ResNet-18* model achieved over 99% accuracy with *resized* input images, while classical models, *SVM* (*geometrical forms* features) and *KNN* (*textural* features), exceeded 80% accuracy. *Majority-vote fusion* further demonstrated the complementarity of the models, reaching over 95% accuracy and balancing performance across methods.

These results validate the methodological choices and highlight the potential of *radar*-based approaches for *unobtrusive* and *reliable human activity recognition*. Nevertheless, challenges remain, including inter-subject variability, environmental effects, and the still limited volume of collected data. Addressing these issues will require further validation in realistic clinical settings, in close collaboration with healthcare professionals. A summary

of the *MAD* dataset and the experimental workflow is illustrated in Figure 5.32.

This chapter completes the experimental validation of the proposed system, from the design of the acquisition platform to the evaluation of the learning models. Chapter 6 will synthesise the contributions of this research, emphasise current limitations, and outline perspectives for clinical deployment.

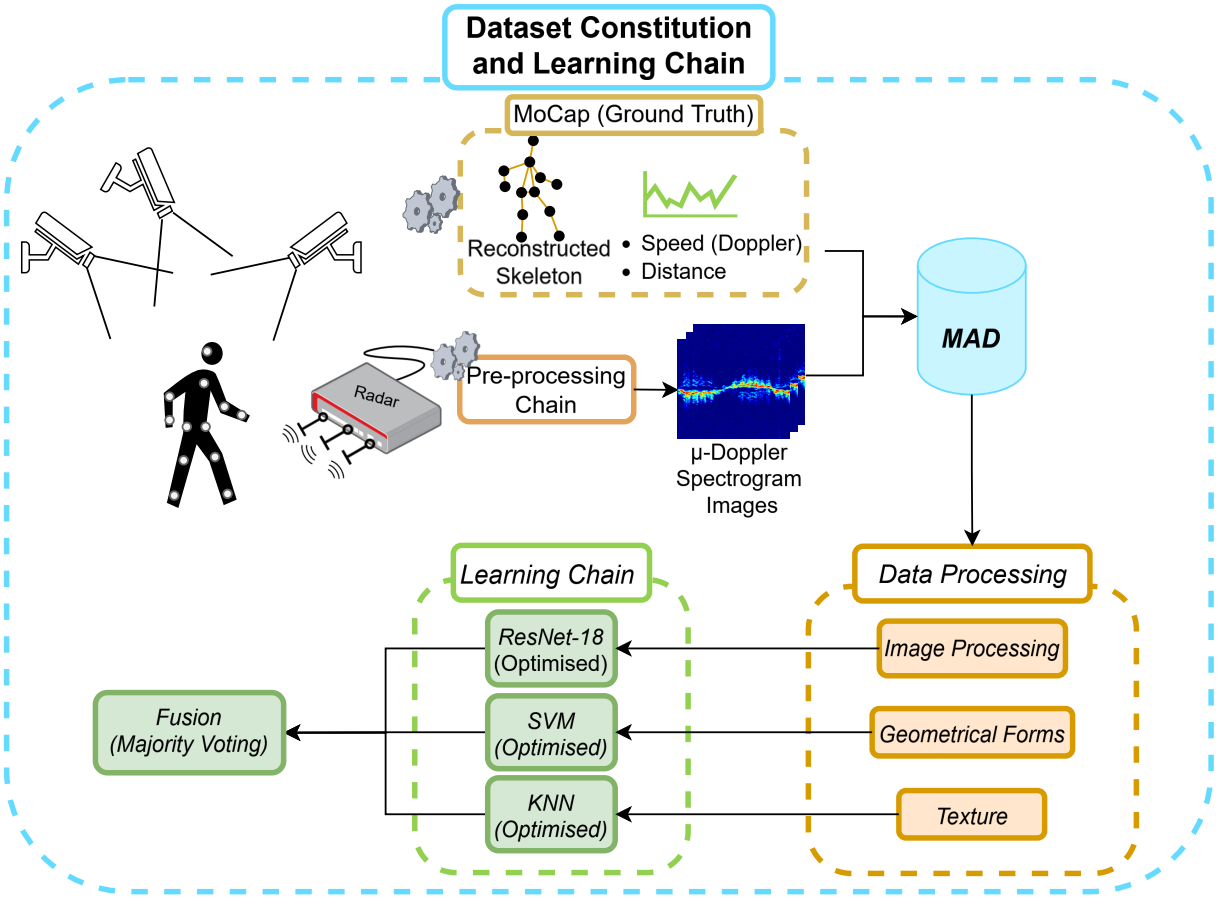

Figure 5.32 – Summary of MAD Constitution and Learning Chain, Chapter 5

Chapter 6

Conclusion and Perspectives

The rates of morbidity and hospitalisation associated with *falls* increase with age, underscoring the lack of appropriate measures to ensure effective monitoring of individuals at risk. Mobility difficulties linked to ageing require regular supervision by healthcare professionals in order to prevent and mitigate the risk of falls. In this context, the use of reliable and unobtrusive technologies to observe human movement and detect early signs of frailty has become an essential tool within clinical practice.

Among *clinical tests* used to assess mobility, gait analysis occupies a central role. Since the earliest studies on human biomechanics, it has been recognised as a valuable indicator of both health status and *fall* risk. Current approaches rely mainly on optical systems or wearable sensors, which have proven effective in laboratory settings and are occasionally employed in remote monitoring. However, these devices present several *limitations*, including technical constraints, privacy concerns, and reduced acceptability due to the need to wear sensors or the perception of being observed.

Given these *limitations*, *radar*-based detection technologies are attracting increasing interest due to their ability to capture movement information without physical contact while preserving data confidentiality. Their robustness under varying lighting conditions and their non-intrusive nature make them particularly promising solutions for use both in clinical environments and in domestic settings. It is within this societal and technological context that the present thesis is situated, with the primary objective of effectively recognising and detecting human activities through *classification* methods based on *radar* technology.

Main Contributions

The state of the art has underscored the inherent complexity involved in defining and interpreting human activity. *Clinical* observations reveal that certain analytical *criteria* overlap across the various tests used to assess patients' risk of falling. In this context, the

characterisation of movement constitutes an essential step towards providing quantitative and complementary data to the practitioners responsible for patient monitoring. By delivering *objective* parameters, technological solutions offer valuable support for the early *detection* of subtle signs that are often difficult for caregivers to perceive.

Radar technologies, and more specifically *FMCW radars*, offer considerable potential in this field. They enable accurate measurement of the *distance* and *velocity* of different body segments during an activity. By applying suitable signal transformations and filtering techniques, it is possible to extract *Doppler signatures* that are characteristic of each type of movement. These signatures reflect the temporal evolution of velocity and form the basis for the analysis and recognition of human activities.

This thesis addresses 4 interrelated scientific *challenges* aimed at developing *learning* methods tailored to *radar* data, while identifying features relevant to the analysis and *prevention* of *fall* risks. The overarching objective is to propose methodological approaches capable of effectively exploiting *radar* signals to extract clinically and preventively relevant information on human activity:

- **The first contribution** focuses on the analysis of existing *datasets* and *classification* models for *radar*-based human activity recognition. The lack of resources suited to the clinical context led us to employ the *Radar Signatures of Human Activities* dataset [57], whose experimental characteristics are comparable to those of our *platform*. We subsequently developed an initial end-to-end processing chain based on a *ResNet-18* model applied to *radar spectrograms* interpreted as *resized* images. This lightweight architecture, suitable for embedded implementation, achieved an accuracy of approximately 90%. Laboratory experiments confirmed the feasibility of this approach while also revealing certain *limitations* of the *dataset*, notably class imbalance and variability in the acquisition protocol. The results of this work were presented at the *DSD 2023 conference* [31].
- **The second contribution** concerns the characterisation of *radar* data prior to *training*. The *limitations* identified in the first study, particularly the presence of noise and the effects of *spectrogram resizing*, prompted the exploration of several preprocessing techniques (including *denoising* and *dynamic cropping*), none of which yielded a significant improvement in performance. To enhance data consistency, two activities from the *Radar Signatures of Human Activities* dataset [57] were combined. Subsequently, an approach based on the analysis of the *geometric form* of the extracted *signatures* was proposed, employing an *SVM* classifier that achieved

an accuracy of approximately 86%. An optimisation method adapted to this case study was also developed, resulting in two publications presented at the *Gretsi'2025 and DSD 2025 conference* [30, 29]. Finally, a complementary analysis of *spectrogram texture* using a *KNN* classifier achieved an accuracy of 67% on the same *dataset*, thereby confirming the impact of noise on the quality of the *textural* descriptors.

- **The third contribution** concerns the implementation of a complete *learning* chain combining the previously developed approaches. After reducing the *Radar Signatures of Human Activities* dataset [57] to five activity classes, the *ResNet-18* model achieved an accuracy ranging from 90% to 95%. A comparative analysis of the three methods revealed that *ResNet-18* exhibits a high level of confidence in its predictions, whereas the *SVM* (form-based) and *KNN* (texture-based) models are more sensitive to uncertainty. To exploit the complementarity of these approaches, a *majority-vote fusion* strategy was proposed, which stabilised the overall accuracy at 94%, with a slight decrease ($\sim$ 2%) for *ResNet-18* but a notable improvement (20–30%) for the other two models.
- **The fourth contribution** concerns the creation of a *new dataset*. One hundred participants were recruited, and nine distinct activities were identified. The *dataset* is modular and can be reduced to five activities, enabling direct comparison with the *Radar Signatures of Human Activities* dataset [57]. Applying our *classification* methods, the *ResNet-18* model achieved an average accuracy exceeding 98%, the *SVM* approximately 80%, and the *texture*-based *KNN* 87%, benefiting from the reduced noise in the data. However, the *majority-vote fusion* approach decreased the average accuracy to 60%, as this mechanism tends to reinforce dominant classes rather than compensate for uncertainties. These results demonstrate that while each model effectively captures its own distinctive features, combining their predictions does not necessarily mitigate data uncertainty.

The work presented in this thesis has resulted in the development of a complete processing chain for the *recognition* of human activities from *radar* signals. Effective *classification* methods, including *ResNet-18*, have been validated in an embedded context, while the characterisation of data has facilitated the proposal of practical *geometric* and *textural* representations. The *fusion* of approaches highlighted the complementarity of the models, and the creation of a *new dataset* enabled direct comparison with the state of the art. Collectively, these contributions demonstrate the potential of *radar* as a non-intrusive tool for activity analysis and *fall*-risk *prevention* in *clinical* settings.

Limitations and Research Prospects

Although the results obtained are promising, several *limitations* have been identified, highlighting potential avenues for future research:

- **Reliability of the *learning* chain**: Our experiments indicate that combining models can be highly effective on small *datasets* containing noisy data. However, when applied to larger *datasets*, this approach may encounter *limitations* or even exhibit reduced effectiveness. Further evaluation on additional *datasets* will therefore be necessary to fully assess the performance of the proposed methodology.
- **Variability of input data for *classification***: At present, *spectrograms* are used as the sole input for *classification*, whereas additional information, such as the *distance* travelled by the individual, could enhance the *learning* process. The exclusive reliance on *spectrograms* therefore represents a potential *limitation* on overall performance.
- **Improvement of texture parameter extraction:** Minor discrepancies were observed between the *MATLAB* and *Python* implementations of the *co-occurrence matrix*. Future work will involve a systematic analysis of these differences and the *optimisation* of the *Python* version to ensure full reproducibility and consistency of the extracted *texture* features.
- **Target population**: Although the *dataset* is highly valuable for our experiments, it does not include individuals over the age of 60 or those with specific pathologies. For *clinical* applications, the target population will need to be adapted to accurately assess *fall* risk. Nonetheless, this *dataset* provides a useful foundation for establishing a baseline of healthy participants.
- **Complex environment**: All data acquisitions were conducted in a *laboratory* setting, in a room with minimal clutter and controlled lighting. The presence of unobstructed windows or obstacles in *real-world environments* could affect participants' behaviour. Consequently, acquisitions in more representative contexts are required to fully validate the proposed approach.
- **Continuous *flow* and *multi-subject* analysis**: At present, each recording corresponds to a *single* activity performed by *one individual*. In *real-world* scenarios, a person may perform *multiple* activities in succession and may be *accompanied* by others, such as friends, family, or carers. It will therefore be necessary to develop

methods capable of handling continuous activity streams and the simultaneous presence of multiple subjects.

Based on the *limitations* identified, several avenues of research may be pursued to extend and build upon the work presented in this thesis:

- **Exploring multimodality**: At present, our work relies exclusively on *radar* data. In *clinical* or remote-monitoring contexts, patients may use additional measurement devices. Integrating such data would allow us to complement our observations and recalibrate our algorithms before full *autonomous* deployment.
- **Adaptable *classification***: At present, our models *recognise* activities such as *falling*. For effective risk *prevention*, it will be necessary to identify early signs of frailty, including changes in posture, slight tremors, or reliance on objects for support. With data collected in a *clinical* setting, our models could be adapted to detect both activities and the onset of risky behaviours.
- **Reliable data characterisation**: The calculation methods and resulting data quality can *vary* depending on the programming *language* employed. *Standardisation* and rigorous *validation* of these methods are therefore necessary to ensure reliable parameter extraction and to produce *comparable* results.
- ***Clinical validation* on the target population**: Future studies should focus on large-scale *validation*, particularly in *real-world* environments such as assisted-living facilities or patients' homes. This would allow for the assessment of effectiveness, robustness, sensitivity to changes in gait, and overall user acceptability.
- **Continuous activity *detection* and *segmentation***: Current processing methods focus on data in which each activity is isolated. It will be essential to develop techniques capable of handling continuous *radar* data streams, with the ability to automatically *detect* and *segment* individual activities.
- **Multi-person handling**: Our current systems are limited to a *single individual* at a time. To be applicable in *real-world* scenarios, it will be necessary to manage the simultaneous presence of *multiple individuals*, for example, a patient accompanied by a carer. This will require advanced signal-processing techniques to differentiate and track each person within the observation area.

Final Reflections and Perspectives

This thesis investigates the use of *radar*-based *human activity recognition* for analysing the movements of *elderly* individuals. In the context of *fall*-risk *prevention*, it proposes a tool that is non-restrictive and non-intrusive for the patient. By developing a *learning* chain and implementing various methods for processing input data, this work has demonstrated the feasibility of recognising complex activities with an accuracy of approximately 90%.

Beyond the results obtained, this research emphasises the importance of complementary methods, data quality, and the adaptation of tools to the target population. It also highlights concrete avenues for *improving* models, *managing* continuous activity *flows*, integrating *multimodal* signals, and addressing *complex* scenarios, such as the simultaneous presence of multiple individuals or more *realistic environmental conditions*.

It should be noted, however, that the proposed solution has not yet been evaluated in real-world conditions. Consequently, the *optimisation strategies* developed in this thesis, while effective in laboratory settings, remain to be tested in a practical environment. Investigating how the system behaves under real-life constraints, including furniture layouts, multiple occupants, and uncontrolled movement patterns, represents a valuable and necessary next step.

In the longer term, this work lays the foundation for *real clinical integration*, enabling the *monitoring* of patient mobility, the *detection* of early signs of frailty, and the support of *fall*-prevention strategies. It also provides a robust methodological framework for future research on *radar*-based activity recognition, while highlighting the *challenges* associated with the *generalisation*, *interpretability*, and *reliability* of the systems developed.

Finally, this thesis makes substantial *scientific* and *technical* contributions and opens promising avenues for *clinical application* and *future research*. It forms part of a collaborative and sustainable research initiative, strengthening the *SmartGait* project and exploring new areas of investigation. Two new doctoral projects and the recruitment of an engineer are now continuing these efforts, with the aim of developing a comprehensive, *embedded*, *low-carbon* solution for *fall detection and prevention*, thereby demonstrating both the *scientific vitality* of the project and the *relevance of the research directions* pursued. I was honoured with the "*Femmes et Sciences*" award (2025) by CY Alliance in recognition of my academic achievements.

PUBLICATIONS

Published Journal Paper

[24] A. Bordat, **C. Béranger**, M. Chapron, P. Dobiáš, J. Le Kernec, D. Guyard and O. Romain, "*Toward Automated Fall Risk Assessment: Validation of an FMCW Radar-Based Timed Up and Go Test*," IEEE Transactions on Radar Systems, 3 (2025), pp. 1086-1102. https://ieeexplore.ieee.org/document/11104103

Published International Conference Paper

[31] **C. Béranger**, A. Bordat, M. A. Khelif, P. Dobiáš, N.-S. Vu, J. Le Kernec, D. Guyard and O. Romain, "*Radar-based Human Activity Acquisition, Classification and Recognition towards Elderly Fall Prediction*," in 2023 26th Euromicro Conference on Digital System Design (DSD), Albania, Sept. 2023, IEEE, pp.95-102. 10.1109/DSD60849.2023.00023

[29] **C. Béranger**, A. Bordat, P. Dobiáš, N.-S. Vu, J. Le Kernec, D. Guyard and O. Romain, "*Human Activity Recognition Using SVM-based on micro-Doppler Radar Data Classification*," WiPiEC Journal - Works in Progress in Embedded Computing Journal, 11 (2025), p4. in 28th Euromicro Conference Series on Digital System Design (DSD), Salerno, Italy, Sept. 2025. https://doi.org/10.64552/wipiec.v11i1.92

Published National Conference Paper

[30] **C. Béranger**, A. Bordat, P. Dobiáš, N.-S. Vu, J. Le Kernec, D. Guyard and O. Romain, "*Classification SVM d'activités humaines acquises avec radar*," in 30è Colloque sur le traitement du signal et des images, no. 2025-1760 in Gretsi Colloque, Strasbourg, Août 25 - Août 29 2025, GRETSI - Groupe de Recherche en Traitement du Signal et des Images, p. p. 1185–1188. https://gretsi.fr/data/colloque/pdf/2025_dobias1760.pdf

Ongoing Papers

C. Béranger, A. Bordat, P. Dobiáš, N.-S. Vu, J. Le Kernec, D. Guyard and O. Romain, "*A New Dataset for Radar-Based Human Activity Recognition and Classification Towards Fall Risk Prevention*,". (Journal)

C. Béranger, A. Bordat, L. Acin, P. Dobiáš, C. Simon Chane, N.-S. Vu, J. Le Kernec, D. Guyard, K. Linardopoulou, A. J. Tsagmo Denkeng, L. M. Rampre, O. Ghoneim, A. Histace and O. Romain, "*MAD Scientific Data Descriptor*". (Data Descriptor, provisional title)

BIBLIOGRAPHY

[1] Académie Francophone de la Télémédecine & e-Santé, *Historique et situation actuelle de la téléassistance*, 2020. https://academie-francophone-telemedecine-et-e-sante.org/histoire-teleassistance/.

[2] C. Achard, *Cours de traitement d'images*, 2002–2003. http://cqcb.mm.free.fr/cours/M1S1/TSI/cours%20vision.pdf.

[3] L. Acin, *Machine Learning Algorithms for Event-Based Camera*, PhD thesis, CY Cergy Paris Université, 2025. https://theses.fr/s307416.

[4] Agence régionale de santé Auvergne-Rhône-Alpes, *Plan antichute des personnes âgées*, 2025. https://www.auvergne-rhone-alpes.ars.sante.fr/plan-antichute-des-personnes-agees-0.

[5] S. Agrawal, *Hyperparameter tuning of KNN classifier*, 2023. https://medium.com/@agrawalsam1997/hyperparameter-tuning-of-knn-classifier-a32f31af25c7.

[6] A. Ahmed and Y. D. Zhang, *Radar-based dataset development for human activity recognition*, in 2020 IEEE Signal Processing in Medicine and Biology Symposium (SPMB), 2020, pp. 1–4. https://ieeexplore.ieee.org/document/9353619.

[7] K. Ahuja, Y. Jiang, M. Goel, and C. Harrison, *Vid2Doppler: Synthesizing Doppler Radar Data from Videos for Training Privacy-Preserving Activity Recognition*, in Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems, Yokohama Japan, May 2021, ACM, pp. 1–10. https://dl.acm.org/doi/10.1145/3411764.3445138.

[8] Ameli, *L'exercice physique recommandé au quotidien et la lutte contre la sédentarité*, 2025. https://www.ameli.fr/assure/sante/themes/activite-physique-sante/exercice-physique-recommande-quotidien.

[9] M. Amin, B. Erol, M. Francesco, and A. Ravisankar, *Realization of radar-based fall detection using spectrograms*, in Compressive Sensing VII: From Diverse Modalities to Big Data Analytics, 05 2018,

p. 10. https://www.researchgate.net/publication/325132328_Realization_of_radar-based_fall_detection_using_spectrograms.

[10] M. H. Arshad, M. Bilal, and A. Gani, *Human activity recognition: Review, taxonomy and open challenges*, Sensors, 22 (2022), p. 6463. https://pmc.ncbi.nlm.nih.gov/articles/PMC9460866/.

[11] Artisight, Inc., *Patient room – remote monitoring & fall prevention | artisight*, 2025. https://artisight.com/solutions/patient-room/.

[12] N. Ashfaq, I. Sadiq, M. Junaid, M. H. Khan, and M. S. Farid, *Optimizing human activity recognition with ensemble deep learning on wearable sensor data*, International Journal of Innovations in Science & Technology, 6 (2024), p. 1751–1767. https://journal.50sea.com/index.php/IJIST/article/view/1069.

[13] F. Ayaz, B. Alhumaily, S. Hussain, M. A. Imran, K. Arshad, K. Assaleh, and A. Zoha, *Radar Signal Processing and Its Impact on Deep Learning-Driven Human Activity Recognition*, Sensors, 25 (2025), p. 724. https://www.mdpi.com/1424-8220/25/3/724.

[14] O. Banos, J. M. Galvez, M. Damas, H. Pomares, and I. Rojas, *Window size impact in human activity recognition*, Sensors, 14 (2014), pp. 6474–6499. https://www.researchgate.net/publication/300786728_Multiwindow_Fusion_for_Wearable_Activity_Recognition.

[15] M. Bantuelle, F. Baudier, C. Begin, V. Boudreault, H. Bourdessol, P. Dejardin, C. Ferron, F. Loew, M. Parisien, B. Petit, S. Pin, M. Saint-Laurent, A. Sizaret, and F. Trickey, *Prevention of falls in the elderly living at home*, Réseau francophone de prévention des traumatismes et de promotion de la sécurité, éditions INPES, 2008. https://www.educasante.org/wp-content/uploads/2018/03/RBP-English.pdf.

[16] M. Barnachon, *Reconnaissance d'actions en temps réel à partir d'exemples*, PhD thesis, Université Claude Bernard - Lyon I, 2013. https://theses.hal.science/tel-00820113v1.

[17] A. Benmahoum, *Analyse des paramètres de haralick pour la classification des anomalies mammaires*, mémoire de fin d'études (master), Université 8 Mai 1945 – Guelma, Département d'Informatique, 2020. https://dspace.univ-guelma.dz/jspui/bitstream/123456789/10083/1/BENHAMAHOUM_ABDELFETEH1602970775.pdf.

[18] K. Berg, *Measuring balance in the elderly: preliminary development of an instrument*, Physiotherapy Canada, 41 (1989), pp. 304–311. https://utpjournals.press/doi/10.3138/ptc.41.6.304.

[19] G. Bhavanasi, L. Werthen-Brabants, T. Dhaene, and I. Couckuyt, *Patient activity recognition using radar sensors and machine learning*, Neural Computing and Applications, (2022), pp. 1–16. http://dx.doi.org/10.1007/s00521-022-07229-x.

[20] S. Björklund, H. Petersson, and G. Hendeby, *Features for micro-doppler based activity classification*, IET Radar, Sonar & Navigation, 9 (2015), p. 1181 – 1187. https://doi.org/10.1049/iet-rsn.2015.0084.

[21] H. Blain, C. Annweiler, G. Berrut, P. L. Bernard, J. Bousquet, P. Dargent-Molina, P. Friocourt, F. Puisieux, J.-B. Robiaud, and Y. Rolland, *Anti-fall plan for the elderly in France 2022-2024: objectives and methodology*, Gériatrie et Psychologie Neuropsychiatrie du Vieillissement, 21 (2023), pp. 286–294. http://www.john-libbey-eurotext.fr/medline.md?doi=10.1684/pnv.2023.1122.

[22] J. M. Blodgett et al., *One-legged balance performance and fall risk in mid-life: Longitudinal evidence from a british birth cohort*, American Journal of Preventive Medicine, 63 (2022), pp. 946–954. https://pmc.ncbi.nlm.nih.gov/articles/PMC10499759/.

[23] A. Bordat, *Early Estimation of Signs of Biomechanical Failure by Non-Conventional Radar Imaging. Applications to Robust Fall Prediction in the Elderly*, PhD thesis, CY Cergy Paris Université, 2025. https://theses.fr/s296911.

[24] A. Bordat, C. Béranger, M. Chapron, P. Dobias, J. Le Kernec, D. Guyard, and O. Romain, *Toward automated fall risk assessment: Validation of an fmcw radar-based timed up and go test*, IEEE Transactions on Radar Systems, 3 (2025), pp. 1086–1102. https://ieeexplore.ieee.org/document/11104103.

[25] A. Bordat, P. Dobias, J. L. Kernec, D. Guyard, and O. Romain, *GPU Based Implementation for the Pre-Processing of Radar-Based Human Activity Recognition*, in 2022 25th Euromicro Conference on Digital System Design (DSD), Maspalomas, Spain, Aug. 2022, IEEE, pp. 593–598. https://ieeexplore.ieee.org/document/9996765/.

[26] A. Bordat, P. Dobias, J. Le Kernec, D. Guyard, and O. Romain, *Towards Real-Time Implementation for the Pre-Processing of Radar-Based Human Activity Recognition*, in 2022 IEEE 31st International Symposium on Industrial Electronics (ISIE), Anchorage, AK, USA, June 2022, IEEE, pp. 635–638. https://ieeexplore.ieee.org/document/9831677/.

[27] D. Bouchaboua, S. M. Nguyen, C. Lohr, B. LeDuc, and I. Kanellos, *Reconnaissance d'activités de la vie quotidienne au moyen de capteurs domotiques et d'apprentissage profond : lorsque syntaxe, sémantique et contexte se rencontrent*, 2025. https://inria.hal.science/hal-04401882v1/document.

[28] L. Brault Eloy, *Évaluation des risques de chute d'une personne âgée : 5 tests à réaliser*, 2024. https://www.filien.com/evaluation-risques-chute-personne-agee/.

[29] C. Béranger, A. Bordat, P. Dobias, N.-S. Vu, J. Le Kernec, D. Guyard, and O. Romain, *Human activity recognition using svm-based on micro-doppler radar data classification*, WiPiEC Journal - Works in Progress in Embedded Computing Journal, 11 (2025), p. 4. https://wipiec.digitalheritage.me/index.php/wipiecjournal/article/view/92.

[30] C. Béranger, A. Bordat, P. Dobias, S. VU, julien Le Kernec, D. Guyard, and O. Romain, *Classification svm d'activités humaines acquises avec radar*, in 30è Colloque sur le traitement du signal et des images, no. 2025-1760 in Gretsi Colloque, Strasbourg, Août 25 - Août 29 2025, GRETSI - Groupe de Recherche en Traitement du Signal et des Images, pp. p. 1185–1188. https://gretsi.fr/data/colloque/pdf/2025_dobias1760.pdf.

[31] C. Béranger, A. Bordat, M. A. Khelif, P. Dobias, N.-S. Vu, J. Le Kernec, D. Guyard, and O. Romain, *Radar-based human activity acquisition, classification and recognition towards elderly fall prediction*, in 2023 26th Euromicro Conference on Digital System Design (DSD), Golem, Albania, Sept. 2023, IEEE, pp. 95–102. https://ieeexplore.ieee.org/document/10456803/.

[32] J. D. Cardenas, C. A. Gutierrez, and R. Aguilar-Ponce, *Deep learning multi-class approach for human fall detection based on doppler signatures*, International Journal of Environmental Research and Public Health, 20 (2023). https://doi.org/10.3390/ijerph20021123.

[33] Centre hospitalier de l'Université de Montréal (CHUM), *Guide de pratique clinique pour la prévention des chutes*, 2015. https://24jours.ca/files/guide_de_pratique_clinique_chutes_2015_09_24.pdf.

[34] Centre intégré de santé et de services sociaux de Chaudière-Appalaches, *Programme de prévention des chutes*, 2017. https://www.cisss-ca.gouv.qc.ca/files/Programme_de_prevention_des_chutes-juillet_2018.pdf.

[35] Centre intégré de santé et de services sociaux de la Montérégie-Est (CISSSME), *Règle de soins infirmiers rsi-006 : Évaluation du risque de chute*, July 2022. https://letourdelest.ca/wp-content/uploads/rsi-006-cisssme-evaluation-du-risque-de-chute-2022-07-07-vf.pdf.

[36] M. Chakraborty, H. C. Kumawat, S. V. Dhavale, and A. A. B. Raj, *DIAT-μ RadHAR (Micro-Doppler Signature Dataset) & μ RadNet (A Lightweight DCNN)—For Human Suspicious Activity Recognition*, IEEE Sensors Journal, 22 (2022), pp. 6851–6858. https://ieeexplore.ieee.org/document/9715052/.

[37] V. Chen, F. Li, S.-S. Ho, and H. Wechsler, *Micro-doppler effect in radar: phenomenon, model, and simulation study*, IEEE Transactions on Aerospace and Electronic Systems, 42 (2006), pp. 2–21. https://ieeexplore.ieee.org/document/1603402.

[38] V. C. Chen and H. Ling, *Time-Frequency Transforms for Radar Imaging and Signal Analysis*, Artech House, Boston, MA, USA, 2002. https://us.artechhouse.com/Time-Frequency-Transforms-for-Radar-Imaging-and-Signal-Analysis-P1002.aspx.

[39] Z. Chen, G. Li, F. Fioranelli, and H. Griffiths, *Personnel Recognition and Gait Classification Based on Multistatic Micro-Doppler Signatures Using Deep Convolutional Neural Networks*, IEEE Geoscience and Remote Sensing Letters, 15 (2018), pp. 669–673. https://ieeexplore.ieee.org/document/8307105.

[40] H. Cho, S. Kang, Y. Sim, S. Lee, and Y. Jung, *Fall Detection Based on Continuous Wave Radar Sensor Using Binarized Neural Networks*, Applied Sciences, 15 (2025), p. 546. https://www.mdpi.com/2076-3417/15/2/546.

[41] C. Clemente, A. W. Miller, and J. J. Soraghan, *Robust principal component analysis for micro-doppler based automatic target recognition*, in 3rd IMA conference on Mathematics in Defence, GBR, October 2013. https://strathprints.strath.ac.uk/51885/.

[42] I. Cohen and N. Levanon, *Adjusting 3-pulse canceller to enhance slow radar targets*, in 2015 IEEE International Conference on Microwaves, Communications, Antennas and Electronic Systems (COMCAS), 2015, pp. 1–5. https://ieeexplore.ieee.org/document/7360360.

[43] Conseil Dépendance, *L'histoire des maisons de retraite : de la création à nos jours*, 2025. https://www.conseildependance.fr/lhistoire-des-maisons-de-retraite-de-la-creation-a-nos-jours/.

[44] M. Conte, *Téléalarme aujourd'hui télécontact demain [texte imprimé] / michèle conte*, 1988. https://kolibris.univ-antilles.fr/discovery/fulldisplay?docid=alma991000108779705746&context=L&vid=33UAG_INST:33UAG_INST&lang=fr&adaptor=Local%20Search%20Engine&tab=TOUT&query=sub,exact,Personnes%20%C3%A2g%C3%A9es%20%20--%20Protection,%20assistance,%20etc,AND&mode=advanced&offset=0.

[45] J. E. Cutting and L. T. Kozlowski, *Recognizing friends by their walk: Gait perception without familiarity cues*, Bulletin of the Psychonomic Society, 9 (1977), pp. 353–356. http://link.springer.com/10.3758/BF03337021.

[46] Daexal, *Équipement d'intervention secouriste / urgence*, 2025. https://www.daexal.fr/12-intervention-urgence?gad_source=1&gad_campaignid=22808936878&gclid=EAIaIQobChMIgbDUudGFjwMVg6n9BR1Efx_uEAMYASAAEgJNjvD_BwE.

[47] D. S. Dareus, *Reconnaissance d'activité en utilisant un smartphone*, 2024. https://www.uqar.ca/app/uploads/2024/11/Reconnaissance_d_activite.pdf.

[48] S. U. de Brest, *Reconnaissance d'activités quotidiennes à partir de données de capteurs*, 2025. https://siia.univ-brest.fr/w/images/5/57/DonneesCapteur.pdf.

[49] C. Debes, A. Merentitis, S. Sukhanov, M. Niessen, N. Frangiadakis, and A. Bauer, *Monitoring activities of daily living in smart homes: Understanding human behavior*, IEEE Signal Processing Magazine, 33 (2016), pp. 81–94. https://ieeexplore.ieee.org/document/7442720.

[50] A. Dey, S. Rajan, G. Xiao, and J. Lu, *Measurement methodology: Radar-based human activity recognition: Is it ready for aging in place?*, IEEE Instrumentation & Measurement Magazine, 26 (2023), pp. 12–19. https://ieeexplore.ieee.org/document/10238364.

[51] Direction interministérielle de la transformation publique (DITP), *La prévention des chutes des personnes âgées — volume 1 : Étude-action*, 2023. https://www.modernisation.gouv.fr/files/2023-12/Volume1_Plan-antichute-personnes-agees-ETUDEACTION-DITP.pdf.

[52] Ehpadia, *Histoire accélérée des ehpad*, 2022. https://www.ehpadia.fr/Histoire-acceleree-des-EHPAD_a888.html.

[53] R. Eisen and H.-C. Mager, eds., *Pflegesicherung in Frankreich*, VS Verlag für Sozialwissenschaften, Wiesbaden, 1999, pp. 263–278. https://doi.org/10.1007/978-3-663-11035-4_9https://doi.org/10.1007/978-3-663-11035-4_9.

[54] Essentiel Autonomie, *Aidants naturels ou aidants familiaux : quelle définition, quel statut juridique*, 2025. https://www.essentiel-autonomie.com/vos-questions-juridiques/aidants-naturels-aidants-familiaux-quelle-definition-quel-statut-juridique.

[55] K. N. Faisal, H. S. Mir, and R. R. Sharma, *Human activity recognition from fmcw radar signals utilizing cross-terms free wvd*, IEEE Internet of Things Journal, 11 (2024), pp. 14383–14394. https://ieeexplore.ieee.org/document/10365493.

[56] D. Fajardo-Delgado, A. Y. Rodríguez-González, S. Sandoval-Pérez, J. E. Molinar-Solís, and M. G. Sánchez-Cervantes, *Genetic programming to remove impulse noise in color images*, Applied Sciences, 14 (2024). https://www.mdpi.com/2076-3417/14/1/126.

[57] D. F. Fioranelli, D. S. A. Shah, H. Li1, A. Shrestha, D. S. Yang, and D. J. L. Kernec, *Radar sensing for healthcare: Associate Editor Francesco Fioranelli on the applications of radar in monitoring vital signs and recognising human activity patterns*, Electronics Letters, 55 (2019), pp. 1022–1024. https://researchdata.gla.ac.uk/848/.

[58] F. Foroughian, F. Quaiyum, P. Theilmann, B. Ghajari, J. E. Piou, O. Kilic, and A. E. Fathy, *Non-contact multi-subject human gait analysis using a state-space method with enhanced 1-d block representation*, IEEE Journal of Electromagnetics, RF and Microwaves in Medicine and Biology, 5 (2021), pp. 155–167. https://ieeexplore.ieee.org/document/9205223.

[59] N. Fusco, *Analyse, modélisation et simulation de la marche pathologique*, PhD thesis, Université Rennes 2, 2008. https://theses.hal.science/file/index/docid/293627/filename/theseFusco.pdf.

[60] P.-M. Gagey and I. Panaud, *Mort aux chutes !*, Éditions des Maisons-Alfort, Paris, France, 2016. Ouvrage sur la prévention des chutes et l'équilibre postural, https://prfssl.centredoc.fr/index.php?lvl=notice_display&id=43644.

[61] P. Gançarski, A. Cornuéjols, C. Wemmert, and Y. Bennani, *Clustering collaboratif : Principes et mise en oeuvre **, in BDA (Gestion de Données - Principes, Technologies et Applications), Nancy, France, Nov. 2017. https://pasteur.hal.science/ENGEES/hal-04338863v1.

[62] A. Ghaleb, *Analyse des micro-Doppler de cibles mobiles déformables en imagerie radar*, theses, Télécom ParisTech, Feb. 2009. https://pastel.hal.science/pastel-00634637v1/file/these_antoine_ghaleb.pdf.

[63] Gouvernement français, *Comment prévenir les risques de chutes chez les personnes âgées*, 2025. https://www.pour-les-personnes-agees.gouv.fr/preserver-son-autonomie/preserver-son-autonomie-et-sa-sante/comment-prevenir-les-risques-de-chutes-chez-les-personnes-agees.

[64] J. Guo, C. Shu, Y. Zhou, K. Wang, F. Fioranelli, O. Romain, and J. Le Kernec, *Complex field-based fusion network for human activities classification with radar*, in IET International Radar Conference (IET IRC 2020), vol. 2020, 2020, pp. 68–73. https://ieeexplore.ieee.org/document/9545567.

[65] S. Z. Gurbuz, M. M. Rahman, E. Kurtoglu, T. Macks, and F. Fioranelli, *Cross-frequency training with adversarial learning for radar micro-Doppler signature classification*, Radar Sensor Technology XXIV, (2020). https://doi.org/10.1117/12.2559155.

[66] S. général des ministères chargés des affaires sociales, *Plans regionaux antichute des personnes agees 2022-2024*, Ministère du travail, de la santé, des solidarités et des familles, (2024). https://solidarites.gouv.fr/sites/solidarite/files/2023-02/Plan%20antichute%20des%20personnes%20%C3%A2g%C3%A9es%20-%20Synth%C3%A8se%20des%20plans%20r%C3%A9gionaux%20-%20Janvier%202023.pdf.

[67] C. Hadjipanayi, M. Yin, A. Bannon, A. Rapeaux, M. Banger, S. Haar, T. S. Lande, A. H. McGregor, and T. G. Constandinou, *Remote gait analysis using ultra-wideband radar technology based on joint range-doppler-time representation*, IEEE Transactions on Biomedical Engineering, 71 (2024), pp. 2854–2865. https://pubmed.ncbi.nlm.nih.gov/38700960/.

[68] R. M. Haralick, K. Shanmugam, and I. Dinstein, *Textural features for image classification*, IEEE Transactions on Systems, Man, and Cybernetics, 3 (1973), pp. 610–621. https://ieeexplore.ieee.org/document/4309314.

[69] Haute Autorité de santé (HAS), *Prévention des chutes accidentelles chez la personne âgée*, 2005. https://www.has-sante.fr/upload/docs/application/pdf/Prevention_chutes_recos.pdf.

[70] ——, *Prescription d'activité physique chez les personnes âgées à risque de chute*, 2024. https://www.has-sante.fr/upload/docs/application/pdf/2024-04/synthese_aps_personnes_agees_a_risque_de_chute.pdf.

[71] K. He, X. Zhang, S. Ren, and J. Sun, *Deep Residual Learning for Image Recognition*, in 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Las Vegas, NV, USA, June 2016, IEEE, pp. 770–778. http://ieeexplore.ieee.org/document/7780459/.

[72] T. Hoshiga, K. Saho, K. Shioiri, M. Fujimoto, and Y. Kobayashi, *Simulation study of deep-learning-based gait classification of young/elderly adults using doppler radar*, Measurement: Sensors, 18 (2021), p. 100103. https://doi.org/10.1016/j.measen.2021.100103.

[73] C.-Y. Hsieh, H.-Y. Huang, K.-C. Liu, K.-H. Chen, S. J.-P. Hsu, and C.-T. Chan, *Subtask Segmentation of Timed Up and Go Test for Mobility Assessment of Perioperative Total Knee Arthroplasty*, Sensors, 20 (2020), p. 6302. https://www.mdpi.com/1424-8220/20/21/6302.

[74] S. Hu, S. Cao, N. Toosizadeh, J. Barton, M. G. Hector, and M. J. Fain, *Radar-Based Fall Detection: A Survey [Survey]*, IEEE Robotics & Automation Magazine, 31 (2024), pp. 170–185. https://ieeexplore.ieee.org/document/10420485/.

[75] IBM, *What is the k-nearest neighbor (knn) algorithm?*, 2025. https://www.ibm.com/think/topics/knn.

[76] S. Iloga, A. Bordat, J. Le Kernec, and O. Romain, *Human Activity Recognition Based on Acceleration Data From Smartphones Using HMMs*, IEEE Access, 9 (2021), pp. 139336–139351. https://ieeexplore.ieee.org/document/9557268/.

[77] V. Imaging, *Walabot - technical brief*, April 24, 2016. https://vayyar.com/technology/.

[78] INSEE, *Perte d'autonomie – france, portrait social*, 2024. https://www.insee.fr/fr/statistiques/8242365?sommaire=8242421.

[79] ——, *Pyramide des âges au 1er janvier 2025*, 2025. https://www.insee.fr/fr/outil-interactif/5014911/pyramide.htm.

[80] Inserm, *Activité physique et prévention des chutes chez les personnes âgées*, 2015. https://www.ipubli.inserm.fr/bitstream/handle/10608/6807/Synthese.html.

[81] M. E. Issa, A. M. Helmi, M. A. A. Al-Qaness, A. Dahou, M. A. Elaziz, and R. Damaševičius, *Human activity recognition based on embedded sensor data fusion for the internet of healthcare things*, Healthcare, 10 (2022). https://doi.org/10.3390/healthcare10061084.

[82] V. D. Jewell, K. Capistran, K. Flecky, Y. Qi, and S. Fellman, *Prediction of falls in acute care using the morse fall risk scale*, Occupational Therapy in Health Care, 34 (2020), pp. 307–319. https://pubmed.ncbi.nlm.nih.gov/32907452/.

[83] G. Johansson, *Visual perception of biological motion and a model for its analysis*, Perception & Psychophysics, 14 (1973), pp. 201–211. http://link.springer.com/10.3758/BF03212378.

[84] G. Johansson, *Visual motion perception*, Scientific American, 232 (1975), pp. 76–89. http://www.jstor.org/stable/24949822.

[85] H. J. Jung, Y. M. Lee, M. Kim, K. E. Uhm, and J. Lee, *Suggested Assessments for Sarcopenia in Patients With Stroke Who Can Walk Independently*, Annals of Rehabilitation Medicine, 44 (2020), pp. 20–37. http://e-arm.org/journal/view.php?doi=10.5535/arm.2020.44.1.20.

[86] C. Karabacak, S. Z. Gürbüz, and A. C. Gürbüz, *Radar simulation of human micro-doppler signature from video motion capture data*, in 2013 21st Signal Processing and Communications Applications Conference (SIU), 2013, pp. 1–4. https://ieeexplore.ieee.org/document/6531365.

[87] B. Khaleghi, A. Khamis, F. O. Karray, and S. N. Razavi, *Multisensor data fusion: A review of the state-of-the-art*, Information Fusion, 14 (2013), pp. 28–44. https://www.sciencedirect.com/science/article/pii/S1566253511000558.

[88] A. Khawani, *Human/Animal Activity Recognition Data Analysis - Classification and Identification*, 2024. https://essay.utwente.nl/100968/1/Khawani_BA_EEMCS.pdf.

[89] Y. Kim and T. Moon, *Human detection and activity classification based on micro-doppler signatures using deep convolutional neural networks*, IEEE Geoscience and Remote Sensing Letters, 13 (2016), pp. 8–12. https://ieeexplore.ieee.org/document/7314905.

[90] Kongakura, *Glcm python - matrice de co-occurrence | haralick | matrice de cooccurence*. https://kongakura.fr/article/Matrice-de-cooccurence-haralick, 2020.

[91] A. Krizhevsky, *Learning Multiple Layers of Features from Tiny Images*, Technical Report, University of Toronto, (2009). https://www.bibsonomy.org/bibtex/2fe5248afe57647d9c85c50a98a12145c/s364315.

[92] A. Krizhevsky, I. Sutskever, and G. E. Hinton, *ImageNet classification with deep convolutional neural networks*, Communications of the ACM, 60 (2017), pp. 84–90. https://dl.acm.org/doi/10.1145/3065386.

[93] J. R. Kwapisz, G. M. Weiss, and S. A. Moore, *Activity recognition using cell phone accelerometers*, ACM SIGKDD Explorations Newsletter, 12 (2011), pp. 74–82. https://www.researchgate.net/publication/220520200_Activity_Recognition_Using_Cell_Phone_Accelerometers.

[94] T. A. Lampert, *Spectrogram Track Detection: An Active Contour Algorithm*, PhD thesis, University of Southampton, Mar. 2010. https://bmva-archive.org.uk/theses/2010/2010-lampert.pdf.

[95] Y. Lang, Q. Wang, Y. Yang, J. Xu, and Z. Yang, *Person identification with limited training data using radar micro-doppler signatures*, Microwave and Optical Technology Letters, 62 (2020), pp. 2716–2722. https://onlinelibrary.wiley.com/doi/abs/10.1002/mop.32125.

[96] Y. Lang, C. Wu, W. Xia, Y. Fu, Y. Fang, and S. Li, *μd-diff: A universal blind denoising method for radar spectrograms using diffusion model*, IEEE Sensors Journal, 25 (2025), pp. 33212–33221. https://ieeexplore.ieee.org/document/11097056.

[97] H. T. Le, S. L. Phung, and A. Bouzerdoum, *Human gait recognition with micro-doppler radar and deep autoencoder*, in 2018 24th International Conference on Pattern Recognition (ICPR), 2018, pp. 3347–3352. https://ieeexplore.ieee.org/document/8546044.

[98] ——, *A Fast and Compact Deep Gabor Network for Micro-Doppler Signal Processing and Human Motion Classification*, IEEE Sensors Journal, 21 (2021), pp. 23085–23097. https://ieeexplore.ieee.org/document/9519700.

[99] H. T. Le, S. L. Phung, A. Bouzerdoum, and F. H. C. Tivive, *Human motion classification with micro-doppler radar and bayesian-optimized convolutional neural networks*, in 2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), 2018, pp. 2961–2965. https://ieeexplore.ieee.org/document/8461847.

[100] J. Le Kernec, F. Fioranelli, C. Ding, H. Zhao, L. Sun, H. Hong, J. Lorandel, and O. Romain, *Radar Signal Processing for Sensing in Assisted Living: The Challenges Associated With Real-Time Implementation of Emerging Algorithms*, IEEE Signal Processing Magazine, 36 (2019), pp. 29–41. https://ieeexplore.ieee.org/document/8746868/.

[101] J.-P. Lepoutre, *Modélisation biomécanique du mouvement : vers un outil d'évaluation pour l'instrumentation en orthopédie*, theses, Université du Sud Toulon Var, Dec. 2007. https://theses.hal.science/tel-00274709.

[102] B. Li, W. Ye, Y. Yang, T. Dong, X. Wang, and Y. Lang, *Universal blind-denoising method of radar spectrograms for unknown noise distribution*, IEEE Transactions on Aerospace and Electronic Systems, 60 (2024), pp. 2354–2366. https://ieeexplore.ieee.org/document/10552638.

[103] H. Li, A. Shrestha, F. Fioranelli, J. L. Kernec, and H. Heidari, *Hierarchical classification on multimodal sensing for human activity recogintion and fall detection*, in 2018 IEEE SENSORS, 2018, pp. 1–4. https://ieeexplore.ieee.org/document/8589797.

[104] Z. Li, *Radar sensing for Ambient Assisted Living application with Artificial Intelligence*, PhD thesis, James Watt School of Engineering, University of Glasgow, Glasgow, United Kingdom, 2024. https://theses.gla.ac.uk/84245/1/2024LiZhenghuiPhD.pdf.

[105] Z. Li, J. Du, B. Zhu, S. E. Greenwald, L. Xu, Y. Yao, and N. Bao, *Doppler radar sensor-based fall detection using a convolutional bidirectional long short-term memory model*, Sensors, 24 (2024). https://doi.org/10.3390/s24165365.

[106] Z. Li, J. Le Kernec, Q. Abbasi, F. Fioranelli, S. Yang, and O. Romain, *Radar-based human activity recognition with adaptive thresholding towards resource*

constrained platforms, Scientific Reports, 13 (2023), p. 3473. https://doi.org/10.1038/s41598-023-30631-x.

[107] LiNote, *Tout savoir sur la téléassistance et les chutes de nos aînés*. https://linote.fr/blog/teleassistance-et-chutes/.

[108] Z. Liu, M. Yang, Y. Yuan, and K. Y. Chan, *Fall detection and personnel tracking system using infrared array sensors*, IEEE Sensors Journal, 20 (2020), pp. 9558–9566. https://ieeexplore.ieee.org/document/9068256.

[109] F. Luo, *Human activity classification using micro-Doppler signatures and ranging techniques*, PhD thesis, Queen Mary University of London, 2020. https://qmro.qmul.ac.uk/xmlui/handle/123456789/68294.

[110] F. Luo, S. Poslad, and E. Bodanese, *Human Activity Detection and Coarse Localization Outdoors Using Micro-Doppler Signatures*, IEEE Sensors Journal, 19 (2019), pp. 8079–8094. https://ieeexplore.ieee.org/document/8716728/.

[111] ——, *Kitchen activity detection for healthcare using a low-power radar-enabled sensor network*, in ICC 2019 - 2019 IEEE International Conference on Communications (ICC), 2019, pp. 1–7. https://ieeexplore.ieee.org/document/8761484.

[112] I. E. López-Delgado, D. Wang, F. Fioranelli, and J. Grajal, *Gait symmetry analysis with fmcw mimo radar*, IEEE Transactions on Microwave Theory and Techniques, 73 (2025), pp. 1–0. https://ieeexplore.ieee.org/document/1234567.

[113] B. Ma and B. Chen, *A 1d cascaded denoising and classification framework for micro-doppler-based radar target recognition*, Remote Sensing, 17 (2025), p. 1515. https://www.mdpi.com/2072-4292/17/9/1515.

[114] X. Ma, R. Zhao, X. Liu, H. Kuang, and M. A. A. Al-qaness, *Classification of human motions using micro-doppler radar in the environments with micro-motion interference*, Sensors, 19 (2019). https://www.mdpi.com/1424-8220/19/11/2598.

[115] K. Maaref, *Biomécanique et analyse de la marche*, 2009. https://s502503c1c2c1db36.jimcontent.com/download/version/1348937570/module/6249217654/name/biomEcanique-et-analyse-de-la-marche%2520(1).pdf.

[116] Manuels MSD, *Présentation du vieillissement*, 2025. https://www.msdmanuals.com/fr/accueil/la-sant%C3%A9-des-personnes-%C3%A2g%C3%A9es/le-vieillissement/pr%C3%A9sentation-du-vieillissement.

[117] T. Marion, P. Gaëlle, L. Linda, C.-B. Laure, R. Annabel, and B. Nathalie, *Chutes des personnes âgées à domicile. Caractéristiques des chuteurs et des circonstances de la chute. Volet « Hospitalisation » de l'enquête ChuPADom, 2018*, Études et enquêtes, 2020. https://www.santepubliquefrance.fr/maladies-et-traumatismes/traumatismes/chute/documents/enquetes-etudes/chutes-des-personnes-agees-a-domicile.-caracteristiques-des-chuteurs-et-des-circonstances-de-la-chute.-volet-hospitalisation-de-l-enquete-chupadom.

[118] R. Mehta, S. Sharifzadeh, V. Palade, B. Tan, A. Daneshkhah, and Y. Karayaneva, *Deep learning techniques for radar-based continuous human activity recognition*, Machine Learning and Knowledge Extraction, 5 (2023), pp. 1493–1518. https://doi.org/10.3390/make5040075.

[119] D. Milhomme, *Processus de surveillance clinique par des infirmières expertes de la douleur*, Recherche en soins infirmiers, 135 (2018), pp. 48–54. https://stm.cairn.info/revue-recherche-en-soins-infirmiers-2018-4-page-48?lang=fr.

[120] d. l. e. d. P. h. Ministère des Solidarités, *Comment fonctionne la grille aggir ?*, 2022. https://www.pour-les-personnes-agees.gouv.fr/preserver-son-autonomie/perte-d-autonomie-evaluation-et-droits/comment-fonctionne-la-grille-aggir.

[121] ——, *Comment le gir est-il déterminé ?*, 2025. https://www.pour-les-personnes-agees.gouv.fr/preserver-son-autonomie/perte-d-autonomie-evaluation-et-droits/comment-le-gir-est-il-determine.

[122] Ministère des Solidarités et de la Santé, *Synthèse du plan antichute des personnes âgées*, 2022. https://sante.gouv.fr/IMG/pdf/plan_antichute_fiche_de_synthese_final_accessible.pdf.

[123] R. Moe-Nilssen and J. L. Helbostad, *Spatiotemporal gait parameters for older adults – an interactive model adjusting reference data for gender, age, and body height*, Gait & Posture, 82 (2020), pp. 220–226. https://doi.org/10.1016/j.gaitpost.2020.09.009.

[124] Monka Care, *Comment prévenir les chutes des personnes âgées ?*, 2025. https://www.monka.care/blog/chute-personne-agee.

[125] S. Monnier-Besnard, D. Soleymani, M. Torres, N. Beltzer, and L. Carcaillon-Bentata, *Identifying frailty to prevent the risk*

of falls in the elderly, La Santé en action, (2022), pp. 42–44. https://www.santepubliquefrance.fr/docs/identifying-frailty-to-prevent-the-risk-of-falls-in-the-elderly.

[126] National Council on Aging (NCOA), *The best medical alert systems with fall detection (2025)*, 2025. https://www.ncoa.org/product-resources/medical-alert-systems/best-medical-alert-systems-with-fall-detection/.

[127] National Library of Medicine, *After a fall in the hospital*, 2024. https://medlineplus.gov/ency/patientinstructions/000441.htm.

[128] A. Newell and H. Simon, *The logic theory machine–a complex information processing system*, IRE Transactions on Information Theory, 2 (1956), pp. 61–79. https://ieeexplore.ieee.org/document/1056797.

[129] N. Nguyen, M. Pham, V.-S. Doan, and V. Le, *Improving human activity classification based on micro-doppler signatures of FMCW radar with the effect of noise*, PLOS ONE, 19 (2024), p. e0308045. https://dx.plos.org/10.1371/journal.pone.0308045.

[130] N. B. Nguyen, A. N. Pham Minh, A. Phan Huy, and Pham Hoang Hung, *Enhance micro-Doppler signatures-based human activity classification accuracy of FMCW radar using the threshold method*, Journal of Military Science and Technology, 95 (2024), pp. 20–28. https://en.jmst.info/index.php/jmst/article/view/1142.

[131] T. Nguyen, *Human activity recognition: Everything you should know about*, 2025. https://www.neurond.com/blog/human-activity-recognition.

[132] H. F. Nweke, Y. W. Teh, M. A. Al-garadi, and U. R. Alo, *Deep learning algorithms for human activity recognition using mobile and wearable sensor networks: State of the art and research challenges*, Expert Systems with Applications, 105 (2018), pp. 233–261. https://www.sciencedirect.com/science/article/pii/S0957417418302136.

[133] F. J. Ordóñez and D. Roggen, *Deep convolutional and lstm recurrent neural networks for multimodal wearable activity recognition*, in Sensors, vol. 16, MDPI, 2016, p. 115. https://www.mdpi.com/1424-8220/16/1/115.

[134] P. W. Palumbo, P. Swaminathan, and S. N. Srihari, *Document Image Binarization: Evaluation Of Algorithms*, in Applications of Digital Image Processing

IX, A. G. Tescher, ed., vol. 0697, International Society for Optics and Photonics, SPIE, 1986, pp. 278 – 285. https://doi.org/10.1117/12.976229.

[135] D. Park, S. Lee, S. Park, and N. Kwak, *Radar-Spectrogram-Based UAV Classification Using Convolutional Neural Networks*, Sensors, 21 (2020), p. 210. https://www.mdpi.com/1424-8220/21/1/210.

[136] J. Park, R. Javier, T. Moon, and Y. Kim, *Micro-Doppler Based Classification of Human Aquatic Activities via Transfer Learning of Convolutional Neural Networks*, Sensors, 16 (2016), p. 1990. http://www.mdpi.com/1424-8220/16/12/1990.

[137] S.-H. Park, *Tools for assessing fall risk in the elderly: a systematic review and meta-analysis*, Aging Clinical and Experimental Research, 30 (2018), pp. 1–16. https://link.springer.com/article/10.1007/s40520-017-0749-0.

[138] J. S. Patel, *Innovative radar detection & AI classification techniques for small UAVs or drones*, PhD thesis, University of Glasgow, 2024. https://theses.gla.ac.uk/84685/4/2024PatelPhD.pdf.

[139] M. Patel, A. Pavic, and V. A. Goodwin, *Wearable inertial sensors to measure gait and posture characteristic differences in older adult fallers and non-fallers: A scoping review*, Gait & Posture, 76 (2020), pp. 110–121. https://doi.org/10.1016/j.gaitpost.2019.10.039.

[140] T. D. Piyadasa, *Image binarization in a nutshell*, 2023. https://medium.com/@tharindad7/image-binarization-in-a-nutshell-b40b63c0228e.

[141] D. Podsiadlo and S. Richardson, *The Timed "Up & Go": A Test of Basic Functional Mobility for Frail Elderly Persons*, Journal of the American Geriatrics Society, 39 (1991), pp. 142–148. https://agsjournals.onlinelibrary.wiley.com/doi/10.1111/j.1532-5415.1991.tb01616.x.

[142] V. B. S. Prasath, H. A. Abu Alfeilat, A. B. A. Hassanat, O. Lasassmeh, A. S. Tarawneh, M. B. Alhasanat, and H. S. E. Salman, *Effects of distance measure choice on k-nearest neighbor classifier performance: A review*, Big Data, 7 (2019), pp. 221–248. https://doi.org/10.1089/big.2018.0175.

[143] Prevenchute, *Matériel médical pour la prévention des chutes chez les personnes âgées*, 2025. https://www.prevenchute.com/.

[144] Programme ICOPE, *Présentation du programme icope*, 2025. https://www.icope.fr/pro/presentation.

[145] X. Qiao, Y. Zhang, L. Zhang, H. Zhang, and J. Li, *Human identification based on radar micro-doppler signal separation using deep convolutional neural networks*, Electronics Letters, 56 (2020), pp. 206–208. https://doi.org/10.1049/el.2019.3380.

[146] S. S. Ram, Y. Li, A. Lin, and H. Ling, *Doppler-based detection and tracking of humans in indoor environments*, Journal of the Franklin Institute, (2008), pp. 1054–1067. https://www.sciencedirect.com/science/article/abs/pii/S0016003208000446.

[147] C. Rasche, *Rapid contour detection for image classification*, IET Image Processing, 12 (2018), pp. 532–538. https://ietresearch.onlinelibrary.wiley.com/doi/full/10.1049/iet-ipr.2017.1066.

[148] S. R. Rath, *Training resnet18 from scratch using pytorch*, 2022. https://debuggercafe.com/training-resnet18-from-scratch-using-pytorch/.

[149] D. G. E. Robertson, G. E. Caldwell, J. Hamill, G. Kamen, and S. N. Whittlesey, *Research Methods in Biomechanics, Second Edition*, Human Kinetics, Champaign, IL, USA, 2 ed., 2013. http://www.gaitlab.ir/books/gaitlab_ref_26_Research_Methods_in_Biomechanics_2014.pdf.

[150] C. A. Ronao and S.-B. Cho, *Human activity recognition with smartphone sensors using deep learning neural networks*, in Expert Systems with Applications, vol. 59, Elsevier, 2016, pp. 235–244. https://dl.acm.org/doi/abs/10.1016/j.eswa.2016.04.032.

[151] F. Rosenblatt, *The perceptron: A probabilistic model for information storage and organization in the brain*, Psychological Review, 65 (1958), pp. 386–408. https://psycnet.apa.org/doiLanding?doi=10.1037%2Fh0042519.

[152] H. Sadreazami, M. Bolic, and S. Rajan, *Contactless Fall Detection Using Time-Frequency Analysis and Convolutional Neural Networks*, IEEE Transactions on Industrial Informatics, 17 (2021), pp. 6842–6851. https://ieeexplore.ieee.org/document/9314269/.

[153] K. Saho, M. Fujimoto, Y. Kobayashi, and M. Matsumoto, *Experimental verification of micro-doppler radar measurements of fall-risk-related gait differences for community-dwelling elderly adults*, Sensors, 22 (2022). https://doi.org/10.3390/s22030930.

[154] N. Salamat, M. M. S. Missen, N. Akhtar, M. Mustahsan, and V. B. S. Prasath, *Color image restoration by filtering methods: a review*, Soft Computing, 28 (2024), pp. 7755–7782. https://doi.org/10.1007/s00500-024-09673-4.

[155] A. Samudra, *k-nn with hyperparameter tuning*, 2019. https://www.kaggle.com/code/arunimsamudra/k-nn-with-hyperparameter-tuning.

[156] L. Santé, *Calcul du gir : Indicateur d'autonomie de la personne âgée*, 2025. https://www.lna-sante.com/conseils/calcul-du-gir/.

[157] M. Sarfraz, *Introductory Chapter: On Digital Image Processing*, IntechOpen, 05 2020, ch. . https://www.intechopen.com/chapters/71817.

[158] Securitas Healthcare, *Solutions de gestion des chutes pour les hôpitaux*, 2025. https://www.securitashealthcare.com/solutions/fall-management.

[159] N. Sedaghati, S. Ardebili, and A. Ghaffari, *Application of human activity/action recognition: a review*, Multimedia Tools and Applications, 84 (2025), pp. 33475–33504. https://link.springer.com/article/10.1007/s11042-024-20576-2.

[160] A.-K. Seifert, M. Grimmer, and A. M. Zoubir, *Doppler radar for the extraction of biomechanical parameters in gait analysis*, IEEE Journal of Biomedical and Health Informatics, 25 (2021), pp. 547–558. https://ieeexplore.ieee.org/document/9093181.

[161] A. Sen, A. Das, S. Pradhan, and S. Chakraborty, *A Dataset for Multi-intensity Continuous Human Activity Recognition through Passive Sensing*, July 2024. http://arxiv.org/abs/2407.21125.

[162] A. Shafi, *K-nearest neighbors (knn) classification with scikit-learn*, 2023. https://www.datacamp.com/tutorial/k-nearest-neighbor-classification-scikit-learn.

[163] D. Shah, *Human activity recognition (har): Fundamentals, models, datasets*, 2023. https://www.v7labs.com/blog/human-activity-recognition.

[164] S. A. Shah and F. Fioranelli, *Human Activity Recognition : Preliminary Results for Dataset Portability using FMCW Radar*, in 2019 International Radar Conference (RADAR), TOULON, France, Sept. 2019, IEEE, pp. 1–4. https://ieeexplore.ieee.org/document/9079098/.

[165] A. Shrestha, C. Murphy, I. Johnson, A. Anbulselvam, F. Fioranelli, J. Le Kernec, and S. Z. Gurbuz, *Cross-Frequency Classification of Indoor Activities with DNN Transfer Learning*, in 2019 IEEE Radar Conference (RadarConf), Boston, MA, USA, Apr. 2019, IEEE, pp. 1–6. https://ieeexplore.ieee.org/document/8835844/.

[166] G. Shufeng, S. Hanyin, and W. Zhefu, *Dataset for "human activity recognition with fmcw radar using few-shot learning"*, 2023. https://dx.doi.org/10.21227/pfzr-g090.

[167] SilverEco, *Chute de la personne âgée : causes, conséquences et prévention*, 2020. https://www.silvereco.fr/chute-de-la-personne-agee-causes-consequences-et-prevention/.

[168] R. Soubra, F. Mourad-Chehade, and A. Chkeir, *Automation of the Timed Up and Go Test Using a Doppler Radar System for Gait and Balance Analysis in Elderly People*, Journal of Healthcare Engineering, 2023 (2023), p. 2016262. https://onlinelibrary.wiley.com/doi/10.1155/2023/2016262.

[169] D. Spathis, I. Perez-Pozuelo, S. Brage, N. J. Wareham, and C. Mascolo, *Self-supervised transfer learning of physiological representations from free-living wearable data*, in Proceedings of the Conference on Health, Inference, and Learning, Virtual Event USA, Apr. 2021, ACM, pp. 69–78. https://dl.acm.org/doi/10.1145/3450439.3451863.

[170] S. Subramaniam, A. I. Faisal, and M. J. Deen, *Wearable Sensor Systems for Fall Risk Assessment: A Review*, Frontiers in Digital Health, 4 (2022), p. 921506. https://www.frontiersin.org/articles/10.3389/fdgth.2022.921506/full.

[171] C. Szegedy, Wei Liu, Yangqing Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich, *Going deeper with convolutions*, in 2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Boston, MA, USA, June 2015, IEEE, pp. 1–9. http://ieeexplore.ieee.org/document/7298594/.

[172] D. Tahmoush, *Review of micro-doppler signatures*, IET Radar, Sonar & Navigation, 9 (2015), pp. 1140–1146. https://doi.org/10.1049/iet-rsn.2015.0118.

[173] C. Tang, W. Li, S. Vishwakarma, K. Woodbridge, S. Julier, and K. Chetty, *Learning from natural noise to denoise micro-doppler spectrogram*, in

Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), 2021, pp. 1–5. https://arxiv.org/pdf/2102.06887.

[174] C. I. Tang, I. Perez-Pozuelo, D. Spathis, S. Brage, N. Wareham, and C. Mascolo, *Selfhar: Improving human activity recognition through self-training with unlabeled data*, Proc. ACM Interact. Mob. Wearable Ubiquitous Technol., 5 (2021). https://doi.org/10.1145/3448112.

[175] W. Taylor, K. Dashtipour, S. A. Shah, A. Hussain, Q. H. Abbasi, and M. A. Imran, *Radar Sensing for Activity Classification in Elderly People Exploiting Micro-Doppler Signatures Using Machine Learning*, Sensors, 21 (2021), p. 3881. https://www.mdpi.com/1424-8220/21/11/3881.

[176] T. Thayaparan, S. Abrol, E. Riseborough, L. Stankovic, D. Lamothe, and G. Duff, *Analysis of radar micro-doppler signatures from experimental helicopter and human data*, IET Radar, Sonar & Navigation, 1 (2007), pp. 289–299. https://doi.org/10.1049/iet-rsn:20060103.

[177] B. Thélot, L. Lasbeur, and G. Pédrono, *La surveillance épidémiologique des chutes chez les personnes âgées*, Bulletin Épidémiologique Hebdomadaire, (2017), pp. 328–335. https://beh.santepubliquefrance.fr/beh/2017/16-17/2017_16-17_5.html.

[178] M. E. Tinetti, *Performance-Oriented Assessment of Mobility Problems in Elderly Patients*, Journal of the American Geriatrics Society, 34 (1986), pp. 119–126. https://agsjournals.onlinelibrary.wiley.com/doi/10.1111/j.1532-5415.1986.tb05480.x.

[179] M. J. Torres, G. Pédrono, A. Rigou, L. Carcaillon-Bentata, and N. Beltzer, *Identifying profiles of people aged 65 and over who fall at home and associated falling-induced injuries: The French ChuPADom study*, Injury, 53 (2022), pp. 2511–2518. https://linkinghub.elsevier.com/retrieve/pii/S0020138322001267.

[180] Tous Ergo, *Équipements anti-chute pour personnes âgées*, 2025. https://www.tousergo.com/besoin/produits-anti-chute-age.

[181] F. Tupin, *"denoising" radar satellite images with artificial intelligence*, 2025. https://www.ip-paris.fr/en/news/denoising-radar-satellite-images-artificial-intelligence.

[182] VirtuSense Technologies, *Solutions pour hôpitaux*, 2025. https://www.virtusense.ai/solutions/hospitals.

[183] S. Vishwakarma, *Learning Algorithms For Micro-Doppler Radar Based Detection, Classification and Imaging of Humans in Indoor Environments*, PhD thesis, Indraprastha Institute of Information Technology, Delhi, 2020. https://repository.iiitd.edu.in/xmlui/handle/123456789/800.

[184] S. Vishwakarma, W. Li, C. Tang, K. Woodbridge, R. Adve, and K. Chetty, *SimHumalator: An Open-Source End-to-End Radar Simulator for Human Activity Recognition*, IEEE Aerospace and Electronic Systems Magazine, 37 (2022), pp. 6–22. https://ieeexplore.ieee.org/document/9664284/.

[185] B. Wang, Y. Liu, A. Lu, and C. Wang, *Application of wearable sensors in constructing a fall risk prediction model for community-dwelling older adults: A scoping review*, Archives of Gerontology and Geriatrics, 129 (2025), p. 105689. https://www.sciencedirect.com/science/article/pii/S0167494324003650.

[186] P. Wang, Q. Li, P. Yin, Z. Wang, Y. Ling, R. Gravina, and Y. Li, *A convolution neural network approach for fall detection based on adaptive channel selection of uwb radar signals*, Neural Computing and Applications, 35 (2022), pp. 15967–15980. https://doi.org/10.1007/s00521-021-06795-w.

[187] X. Wang, J. Cao, Q. Zhao, M. Chen, J. Luo, H. Wang, L. Yu, K.-L. Tsui, and Y. Zhao, *Identifying sensors-based parameters associated with fall risk in community-dwelling older adults: an investigation and interpretation of discriminatory parameters*, BMC Geriatrics, 24 (2024), p. 125. https://bmcgeriatr.biomedcentral.com/articles/10.1186/s12877-024-04723-w.

[188] X. Wang, S. Wang, P. Shao, L. Zhu, B. Jiang, and Y. Tian, *Event stream based human action recognition: A high-definition benchmark dataset and algorithms*, arXiv preprint arXiv:2408.09764, (2024). https://arxiv.org/abs/2408.09764.

[189] World Health Organization, *Falls*, 2021. https://www.who.int/fr/news-room/fact-sheets/detail/falls.

[190] ——, *Ageing and health*, 2024. https://www.who.int/fr/news-room/fact-sheets/detail/ageing-and-health.

[191] S. Yang, J. L. Kernec, O. Romain, F. Fioranelli, P. Cadart, J. Fix, C. Ren, G. Manfredi, T. Letertre, I. D. H. Sáenz, J. Zhang, H. Liang,

X. Wang, G. Li, Z. Chen, K. Liu, X. Chen, J. Li, X. Wu, Y. Chen, and T. Jin, *The Human Activity Radar Challenge: Benchmarking Based on the 'Radar Signatures of Human Activities' Dataset From Glasgow University*, IEEE Journal of Biomedical and Health Informatics, 27 (2023), pp. 1813–1824. https://ieeexplore.ieee.org/document/10032210/.

[192] X.-X. Yin and S. Hadjiloucas, *Digital filtering techniques using fuzzy-rules based logic control*, Journal of Imaging, 9 (2023), p. 208. https://doi.org/10.3390/jimaging9100208.

[193] Youngwook Kim and Hao Ling, *Human activity classification based on micro-Doppler signatures using an artificial neural network*, in 2008 IEEE Antennas and Propagation Society International Symposium, San Diego, CA, July 2008, IEEE, pp. 1–4. http://ieeexplore.ieee.org/document/4619933/.

[194] ——, *Human Activity Classification Based on Micro-Doppler Signatures Using a Support Vector Machine*, IEEE Transactions on Geoscience and Remote Sensing, 47 (2009), pp. 1328–1337. http://ieeexplore.ieee.org/document/4801689/.

[195] R. Yu, X. Zhang, Y. Wang, and Z. Zhang, *Radar-based human activity recognition using denoising techniques*, IET Remote Sensing, 13 (2024), pp. 123–134. https://ietresearch.onlinelibrary.wiley.com/doi/10.1049/rsn2.12501.

[196] J. Zabalza, C. Clemente, G. Di Caterina, Jinchang Ren, J. J. Soraghan, and S. Marshall, *Robust PCA micro-doppler classification using SVM on embedded systems*, IEEE Transactions on Aerospace and Electronic Systems, 50 (2014), pp. 2304–2310. http://ieeexplore.ieee.org/document/6965776/.

[197] H. E. Zein, F. Mourad-Chehade, and H. Amoud, *Reconnaissance intelligente de l'activité humaine en temps réel à l'aide de signaux wi-fi*, 2023. https://hal.science/hal-04220662/document.

[198] M. Zenaldin and R. Narayanan, *Radar micro-doppler based human activity classification for indoor and outdoor environments*, in Proceeedings of SPIE, 04 2016. https://doi.org/10.1117/12.2228397.

[199] B. Zhang, L. Zhang, M. Wu, and Y. Wang, *Dynamic gesture recognition based on rf sensor and ae-lstm neural network*, in 2021 IEEE International Symposium on Circuits and Systems (ISCAS), 2021, pp. 1–5. https://ieeexplore.ieee.org/document/9401065.

[200] J. Zhang, *Basic gait analysis based on continuous wave radar*, Gait & Posture, 36 (2012), pp. 667–671. https://doi.org/10.1016/j.gaitpost.2012.04.020.

[201] P. Zhang, C. Lan, J. Xing, W. Zeng, J. Xue, and N. Zheng, *View Adaptive Neural Networks for High Performance Skeleton-based Human Action Recognition*, arXiv:1804.07453 [cs], (2019). http://arxiv.org/abs/1804.07453.

[202] S. Zhang, Y. Li, S. Zhang, F. Shahabi, S. Xia, Y. Deng, and N. Alshurafa, *Deep learning in human activity recognition with wearable sensors: A review on advances*, arXiv preprint arXiv:2111.00418, (2021). https://arxiv.org/abs/2111.00418.

[203] Y. Zhao, H. Zhou, S. Lu, Y. Liu, X. An, and Q. Liu, *Human Activity Recognition Based on Non-Contact Radar Data and Improved PCA Method*, Applied Sciences, 12 (2022), p. 7124. https://www.mdpi.com/2076-3417/12/14/7124.

[204] Y. Zhou, X. Yu, M. Lopez-Benitez, L. Yu, and Y. Yue, *Radar-Based Swimming Activity Recognition with Temporal Dynamic Convolution and Spectral Data Augmentation*, Apr. 2024. https://www.techrxiv.org/users/298450/articles/691889-radar-based-swimming-activity-recognition-with-temporal-dynamic-convolution-and-spectral-data-augmentation?commit=507c076e3e1207111b926e6ee9aff1d3715c563a.

Appendix A

Ethical and Organisational Aspects of the Dataset Creation

A.1 Submission of the Dataset Creation Project to the Ethics Committee

CY Cergy Paris Université (CYU) has an *Ethics Committee* that ensures research involving human participants complies with ethical standards, particularly for non-interventional projects, i.e., studies that do not involve any direct intervention or manipulation of the participants themselves. Approval is required before collecting human data to protect participants' rights and well-being. Our project and protocol were approved on 11th October 2023, and the formal letter of approval is presented in Appendix A.6. The validated documents are as follows:

- **Application Form:** Document submitted to the *Ethics Committee*, summarising the project context, objectives, recording protocol, team members and supervisors, equipment, data storage conditions and duration, and an assessment of potential benefits and risks;
- **Information Letter:** Document provided to participants, including an invitation to participate, a detailed explanation of the project's objectives and expected outcomes, participants' right to withdraw, data usage and storage conditions, project supervisors, potential benefits, dissemination of results, and the right to ask questions. A dedicated website is available at: https://www.etis-lab.fr/platform/imaging-database-radar-event-based-camera-of-daily-movement-activities/
- **Informed Consent:** Document given to participants before the experiment, expressing gratitude, formalising their consent, and confirming their understanding and voluntary commitment to the protocol procedures;

- **Authorisation for Recording and Processing of Voice and Image:** Document given to participants before data collection, covering authorisation for the use of *Radar*, *Event-based camera*, *MoCap* systems, *GAITRite mat*, and *Survey*, and specifying how recorded voice, image, and movement data may be processed, stored, and used in scientific publications or communications;
- **GDPR File:** Document submitted to the *committee* outlining data protection and privacy measures in accordance with the *General Data Protection Regulation* (GDPR), including the types of data collected, responsible parties, retention period, storage conditions, and procedures for data access and consultation.
- **Survey:** Sample document completed by participants at the end of each session, collecting feedback on comfort, understanding of instructions, perceived effort, and overall experience, providing qualitative data to complement the sensor-based measurements.

With the dossier approved and all preparations in place, the study commenced in January 2024, ensuring that all ethical and administrative requirements were met.

A.2 Team and Session Management

This dataset project involves three theses and their supervisory teams, requiring careful organisation of recording sessions. While the goal is to accommodate 100 participants, managing both the *Movement Platform* and individual research work would be overly demanding for each PhD student. To address this, a dedicated team of nine people was established to oversee participants and ensure smooth operation of all sessions.

- Alexandre, Laure, and the author managed both the team and the recording sessions, overseeing operations and checking all equipment before each session to ensure readiness for data collection.
- Six additional team members volunteered during the acquisition period to assist with data collection. They operated the computers for the *MoCap* system and the *Radar* and *Event-based camera* systems, allowing session managers to supervise participants effectively. The contributions of Lili M. Rampre (PhD student), Konstantina Linardopoulou (Post-doc), Antoine Junior Tsagmo Denkeng (PhD student), Omar Ghoneim (PhD student), Petr Dobiáš (Dr.), and Olivier Romain (Prof.) were essential to the project's success.

During each recording session, at least three personnel were present: Alexandre, Laure, or the author managed the room, supervised the participant, and set up equipment, while two team members handled data acquisition, one for *Radar* and the other for *MoCap*.

A.3 Communication and Call for Participants

Participants were recruited to achieve a balanced cohort of 100 individuals across *Age* and *Gender* categories, representing diverse professional backgrounds and locations. Recruitment was conducted via institutional mailing lists (*ETIS Laboratory*, *ENSEA*, *CYU*, *ESIEE-IT*) and the doctoral school at *CYU*, complemented by targeted outreach to underrepresented groups.



Figure A.1 – Communication Poster for the MAD Creation Project

Communication was supported by a dedicated email address, a web page [1], and posters displayed on institutional screens and newsletters, as shown in Figure A.1. Participants registered through the *Evento* platform, providing their age group (*18–40* or *41–60*) and *Gender* (*Male* or *Female*), which enabled the selection of suitable candidates.

Recruitment began in November 2023, with the first participants welcomed in January 2024, and all data collection completed by July 2024. Some challenges were encountered in

1. `https://www.etis-lab.fr/platform/imaging-database-radar-event-based-camera-of-daily-movement-activities/`

recruiting women aged 18–40 and men aged 41–60 due to mailing list demographics. These were addressed through additional targeted outreach to ensure demographic balance.

A.4 Participant Selection Criteria

Participant recruitment followed predefined criteria to ensure balanced representation across *Age* and *Gender* categories within a cohort of 100 individuals:

- **Exclusion Criteria:** As the study involved physical activities such as "*Falls*", the *Ethics Committee* required the definition of conditions under which participation would not be permitted:
  * *Age:* Participants had to be between 18 and 60 years old; individuals outside this range were excluded to limit potential health risks.
  * *Health:* Individuals with mobility impairments or osteoporosis were not eligible.
- **Selection Criteria:** To ensure a balanced dataset, the following parameters were defined:
  * *Gender Parity:* Participants indicated their biological gender (*Female* or *Male*) to achieve equal representation.
  * *Age Parity:* Participants selected an age group (*18–40* or *41–60*) to balance both gender and age distributions.
  * *Professional Diversity:* Recruitment aimed to include individuals from a wide range of professional backgrounds through targeted communication efforts.

The final objective was to recruit 50 women and 50 men, equally distributed across the two age groups (*18–40* and *41–60*), ensuring a dataset fully balanced in terms of both *Age* and *Gender*.

A.5 Sessions Planning

The acquisition team operated on a rotation system, conducting an average of two sessions per day and around ten participants per week. Initial sessions were fully managed by the team to quickly address technical issues, including malfunctions of the *Radar*, *Event-based camera*, *MoCap* system, pressure sensors, or computers.

Subsequent sessions were organised to ensure at least one additional planning manager was present, allowing smooth operation even in case of last-minute cancellations. In total, 100 sessions were conducted, each lasting up to two hours for participants and 2.5 hours for session leaders. Alexandre, Laure, and the author oversaw session planning.

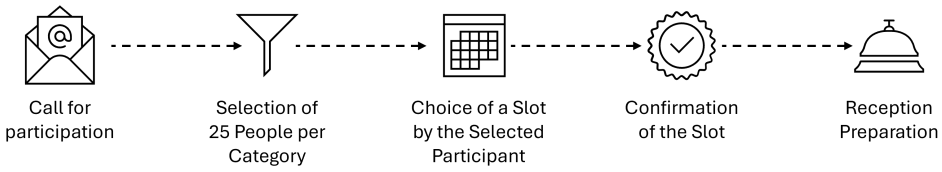

Figure A.2 – Registration Process for Participants

Team availability was assessed every two weeks via an *Evento* form, and slots with at least three available members were scheduled. Participants were informed of available times, sent reminders to maximise attendance, and received confirmations including session details and *Ethics Committee* consent documents. Alternative slots were offered for cancellations, and relevant *ENSEA* departments were notified in advance of external visitors. The main steps of this process are summarised in Figure A.2.

A.6 Research Ethics Committee Approval

Comité d'éthique de la recherche de CY
33 boulevard du Port
95011 Cergy-Pontoise cedex
cer-cy@ml.u-cergy.fr

M. Olivier ROMAIN
Laboratoire ETIS

À Cergy-Pontoise
Le 11 octobre 2023

Référence : protocole n° 202307 - 004

Cher Monsieur,

Nous vous informons que le Comité d'Ethique de la Recherche de CY Cergy Paris université donne un avis final **favorable**, au protocole de recherche soumis par vos soins.

Description du projet de recherche :

Titre : *Base de données d'imagerie (Radar, Caméra Événementielle) d'activités quotidiennes de déplacement*

Porteur et responsable : Olivier ROMAIN, Professeur des universités, laboratoire ETIS, CY Cergy Paris Université

Numéro de référence : 202307 - 004

Je vous prie de croire, Cher Monsieur, en l'expression de ma considération la meilleure.

Sabrina KELLOUCHE-GAILLARD

Présidente du CER-CY

Appendix B

Complementary Details on Experimental Protocol and MAD Dataset

B.1 Experimental Setup: Additional Logistical Details

We produced a diagram illustrating the room's dimensions and overlaid the detection ranges of each sensor, as summarised in Figure B.1.

This diagram revealed several issues, including an insufficient number of *MoCap* cameras and calibration gaps, particularly in the upper-left corner. To address this, we increased the number of *MoCap* cameras from 6 to 10, repositioned them to maximise coverage, re-centred the *Radar* and *Event-Based camera*, and mounted the conventional camera on a suspended rail, resulting in the configuration shown in Figure 5.2. Thermal noise from the *Radar* circuit board was mitigated by adding a fan, improving data consistency and quality.

Floor markings were placed to delineate the recording area, indicating maximum and minimum dimensions as well as the central reference point required for *MoCap* calibration and participant skeletal reconstruction. Additional markings later indicated start and end points of activities and equipment positions, establishing the room layout presented in Figure 5.18.

All sensors inherently collect anonymous data. For conventional RGB cameras, faces are blurred using an open-source script integrated into the processing chain [1]. To reduce noise in the *MoCap* data, masks were applied using the *Nexus* software to ignore unwanted bright spots that could be mistakenly detected as markers, such as reflections from walls,

1. `https://github.com/ORB-HD/deface`

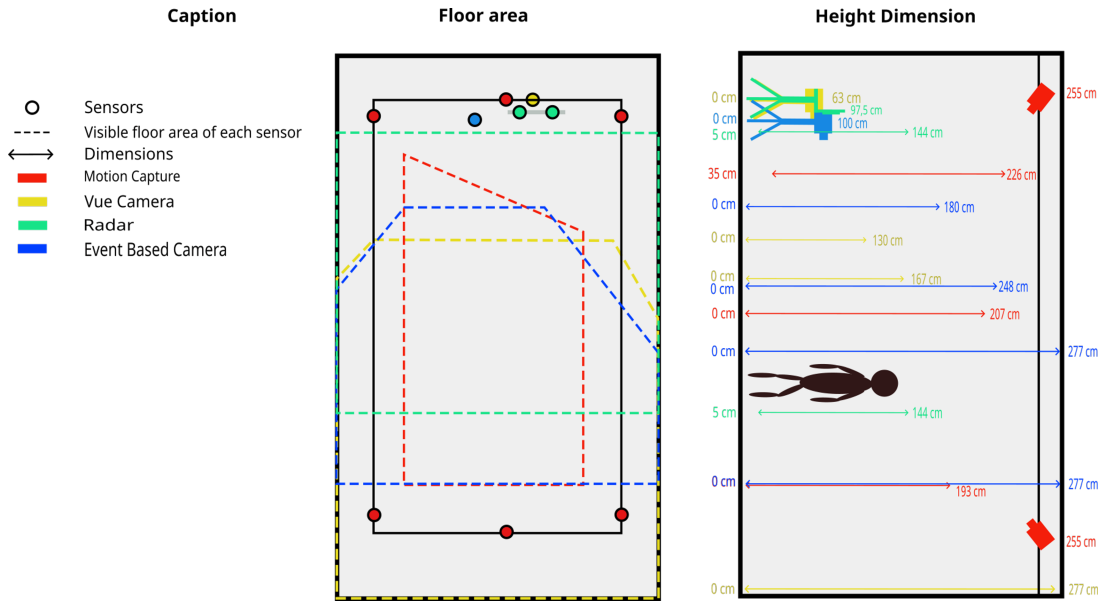

Figure B.1 – Dimensions of the *Movement Platform* Measured in March 2023

windows, or other objects in the room. In addition, a cork panel was placed in front of the *MoCap* screen to block residual reflections that could not be removed by software without affecting the capture of actual participant movements. Camera calibration, which had been affected by vibrations from the entrance door, was stabilised by keeping the door locked and using the adjacent room as an airlock for participants.

B.2 Experimental Accessories

A set of equipment was prepared to accommodate participants and conduct activities in the *Movement Platform* of the ETIS laboratory. Signs were installed to guide participants to the room. The *Radar* equipment, initially placed on a table, was enclosed in a ventilated case developed in collaboration with *BlueLinea* to secure the setup and make it transportable.

As the protocol was being developed, various pieces of equipment were procured, as shown in Figure B.2: a nursing-home chair for the "*Sit Down*" and "*Stand Up*" tasks, *hurdles* for "*Step Over*" an obstacle, a *mat* for "*Fall*", *steps* for "*Up Stairs*" and "*Down Stairs*", and a *Basketball* for "*Pick*" up an object. Straps and knee pads were also provided to secure the *MoCap* markers.

Once the equipment had been installed and *synchronisation* configured, simulations

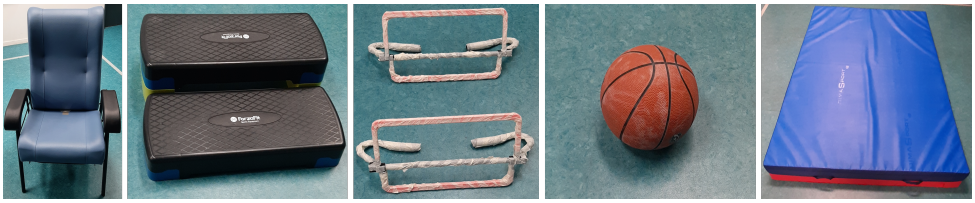

Figure B.2 – List of Equipment Used for Dataset Activities

were conducted to determine the optimal position of each object for every activity and to mark the start and finish areas on the floor. These simulations also allowed the synchronised recording system to be tested and any remaining issues to be resolved.

B.3 Session Schedule and Participant Protocol: Operational Details

Preliminary simulations were conducted as pilot tests to verify the feasibility and sequence of activities. Based on these, the session begins with the first 4 subcategories of the "*Walk*" activity, chosen for its simplicity, familiarity, and ability to familiarise participants with the recording process.

The "*TUG*" test follows, involving the use of a chair and a change in trigger mode, as the recording signal is generated by the chair's pressure sensor. For consistency, "*Stand Up*" uses the same triggering principle, followed by "*Sit Down*", "*Up Stairs*", and "*Down Stairs*" in logical sequence. More static activities, such as "*Pick*" up an object, are interspersed to balance physical load, while "*Step Over*" and "*Semi-Turn Around*", both involving hurdles, are grouped for efficiency. An optional "*Arbitrary*" walk was added later to allow free, relaxed movement, its identifier is 101. Sessions conclude with the "*Fall*" activity.

For "*Step Over*", hurdle distances are adjusted per participant according to stride length, and measured when changing direction to ensure data consistency. A warm-up is included before the "*Fall*" activity to minimise muscle soreness. Activities are assigned identifiers 01–10, with sub-activities numbered 1–5, allowing future expansion to 99 activities and 9 subcategories. The estimated session duration is 45–60 minutes, including preparation, equipment handling, and end-of-session surveys, as summarised in Table B.1.

In accordance with *Ethics Committee* requirements, informed consent is obtained at the start. Participants are introduced to the 3 experimenters, the sensors, and the types of data collected. The list of activities is explained before signing the *Informed Consent*,

Table B.1 – Protocol Stages with Approximate Durations

Protocol Stage	**Duration**
Introductions (experimenters, project, sensors, activities)	35 minutes
Signing of documents	5 minutes
Participant preparation and static calibration	5 minutes
Display of skeleton and explanation of recording procedure	2 minutes
First 4 subcategories of "Walk"	10 minutes
TUG setup and recording	3 minutes
"Stand Up" to "Semi-Turn Around"	42 minutes
"Arbitrary Walk"	2 minutes
"Fall" (setup and warm-up included)	9 minutes
End-of-session survey and equipment removal	5 minutes
Acknowledgements and compensation	2 minutes
Total	**2 hours**

Information Letter, and *Authorisation for Voice and Image Processing* documents.

Participants are then equipped and a static calibration is performed to record reference posture for automatic skeleton reconstruction in Vicon *Nexus*. The *MoCap* screen, normally hidden to avoid reflections, is briefly shown so participants can view their *skeleton*. The experimenter explains activity procedures, including floor markings, durations, and start/stop signals. After each activity, the number of files is checked to prevent data loss, and participants complete the end-of-session survey. Sessions conclude with equipment removal and presentation of a gratuity voucher.

B.4 Data Management Interface and API

Once approval from the *Ethics Committee* had been obtained and the participant recruitment campaign launched, attention turned to data storage. The scale of the project required a centralised, secure system with a monitoring interface capable of:

- **Efficient file transfer:** Session files are automatically transferred to the laboratory server.
- **Data integrity verification:** The system checks that all files have been correctly received.
- **Real-time monitoring:** Completion statistics are displayed for each session as data are added.

End-of-session forms are completed digitally via the interface shown in Figure B.3,

eliminating manual data entry and automating statistics on participants and session types.

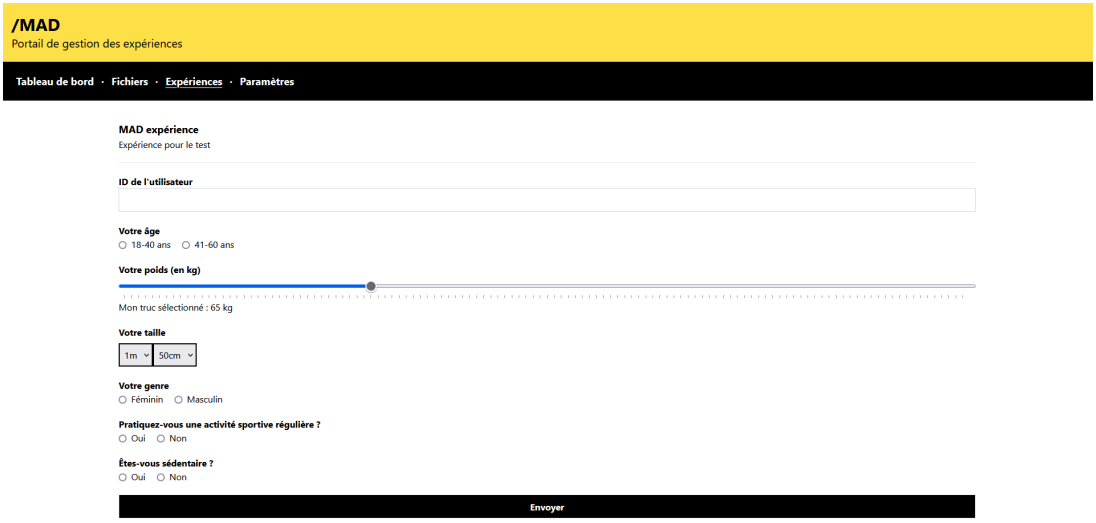

Figure B.3 – Interface Used for Completing the Form

A dedicated *API* was developed to meet these requirements, as illustrated in Figure B.4. All data are stored on a secure *ETIS laboratory* server, accessible only to authorised personnel. Two access modes are provided:

1. **SSH connection:** allows browsing folders, counting files, or manual deletion via the terminal;

2. **Web interface:** enables intuitive drag-and-drop uploads, automatic transfer to a *Virtual Machine* (VM) linked to the *API*, and real-time monitoring of folder creation and file population.

The *API* also integrates the data characterisation form completed at the end of each session. Once validated, the form is converted into a *.json* file containing all participant-entered parameters and stored automatically within the dataset structure.

Extensive testing validated system functionality, automated file creation, transfer reliability, and adherence to the defined dataset tree structure. An accurate estimation of the total expected file volume was also performed to optimise server architecture and storage management.

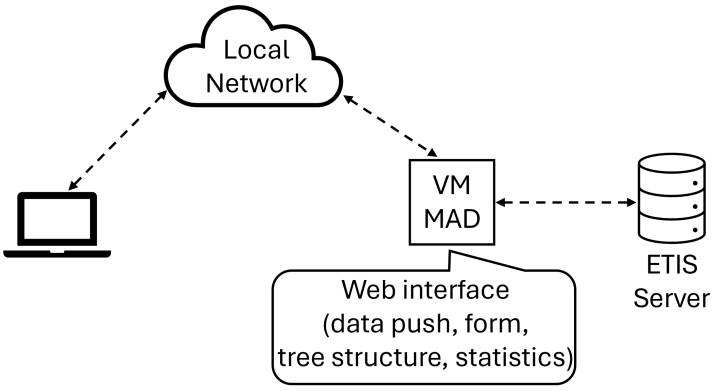

Figure B.4 – Functioning of the API Utilised for the Dataset

B.5 Detailed Participant Metadata and Survey Responses

Survey responses provided additional participant information, including engagement in sport and sedentary behaviour. Participants were classified as *physically active* if they regularly practiced club-based or other recurring physical activities, and as *sedentary* if their occupation or daily routine involved predominantly low physical effort.

Table B.2 – Percentage of Sedentary and Physically Active among the Population of MAD

Category	**S[1]& PA[2]**	**S & PNA[4]**	**NS[3]& PA**	**NS & PNA**	**Total**
Female	5%	15%	16%	14%	**50%**
Male	5%	6%	16%	23%	**50%**
Total	**10%**	**21%**	**32%**	**37%**	**100%**

[1] S: Sedentary
[2] PA: Physically Active
[3] NS: Non-Sedentary
[4] PNA: Physically Non-Active

Table B.2 shows that only 10% of sedentary participants engage in sport, evenly split between *Males* and *Females*, whereas 32% of non-sedentary participants are active, again balanced across genders. Although these variables are not currently used, they may inform future analyses.

Table B.3 summarises key characteristics, including *Body Mass Index* (BMI, calculated as $BMI = \frac{\text{Weight (kg)}}{\text{Height (m)}^2}$), physical activity levels, and sedentariness. This well-distributed

Table B.3 – Population Characteristics

Category	**BMI Range (kg/m²)**	**Value**
Underweight	< 18.5	4%
Normal Weight	18.5 – 24.9	56%
Overweight	25.0 – 29.9	33%
Obesity	≥ 30.0	7%

sample provides a robust foundation to study biomechanical variations and their influence on fall risk in realistic, controlled conditions.

Titre : Reconnaissance de la marche et prédiction de chute par apprentissage profond sur des signatures micro-Doppler

Mot clés : Intelligence Artificielle, Reconnaissance d'activité, Prévention des chutes, Radar

Résumé : Les chutes chez les personnes âgées constituent un enjeu majeur de santé publique, entraînant chaque année en France un nombre significatif de décès et d'hospitalisations. Bien que des systèmes de téléassistance aient été déployés pour réduire ces risques, leur efficacité reste limitée en raison d'un taux élevé de fausses alertes.

Cette thèse propose des méthodes de reconnaissance d'activités à partir de données radar non conventionnelles, permettant un suivi précis des mouvements des personnes âgées tout en préservant leur intimité. Les travaux portent sur l'extraction de caractéristiques pertinentes à partir de spectrogrammes radar et sur le développement de modèles de classification performants pour l'identification de l'activité quotidienne.

Un jeu de données multimodal expérimental a été constitué afin de soutenir l'analyse et l'évaluation des modèles. L'optimisation des approches et la fusion des résultats des modèles ont permis d'atteindre un taux de reconnaissance des activités proche de 95%, démontrant la robustesse et l'efficacité des méthodes développées. Ces résultats ouvrent des perspectives concrètes pour l'amélioration des systèmes de téléassistance et la prévention des chutes.

Au-delà de ces applications, ce travail illustre le potentiel des technologies radar pour l'étude des activités humaines et ouvre de nouvelles pistes pour la recherche en suivi comportemental et analyse biomécanique.

Title: Gait Recognition and Fall Prediction with Deep-Learning on Micro-Doppler Signatures

Keywords: Artificial Intelligence, Activity Recognition, Fall Prevention, Radar

Abstract: Falls among elderly people represent a major public health challenge, resulting each year in France in a substantial number of deaths and hospitalisations. Although teleassistance systems have been deployed to mitigate these risks, their effectiveness remains limited due to a high rate of false alarms.

This thesis presents methods for activity recognition using unconventional radar data, enabling precise monitoring of elderly people's movements while preserving their privacy. The work focuses on the extraction of relevant features from radar spectrograms and the development of robust classification models for identifying daily activities.

An experimental multimodal dataset was created to support the analysis and evaluation of the models. Optimisation of the approaches and fusion of model outputs achieved an activity recognition rate of approximately 95%, demonstrating the robustness and effectiveness of the proposed methods. These results provide concrete prospects for improving teleassistance systems and preventing falls.

Beyond these applications, this work highlights the potential of radar technologies for the study of human activities and opens new avenues for research in behavioural monitoring and biomechanical analysis.