Imperial College London
Department of Bioengineering

Learning Steps: Models of Intent-Driven Lower Limb Prosthesis Use

Bálint Károly Hodossy

Submitted in part fulfilment of the requirements
for the degree of Doctor of Philosophy in Bioengineering
at Imperial College London, November 2024

Abstract

Walking is often perceived as a simple and effortless process. However, it is generated by a complex and finely tuned neuromechanical system. Artificial devices that aim to seamlessly support or replace parts of it must adapt to the changing goals and environments of their users. Virtual environments provide an accessible and low-risk way to accelerate the early-stage development of these technologies. To capture the challenge inherent in this task, these simulations must aim to reproduce the complexities of our movement, our environment, and the dynamic control required assist it.

This thesis approaches this problem from two sides. First, we propose locomotion intent estimation models that regress biosignals to the desired future walking path. We investigate the sensitivity of our intent estimators to common error inducing factors such as electrode shifts, and propose ways to mitigate their effect. Second, we apply physics-based animation methods to create full-body motion corresponding to arbitrary walking path trajectories. We then use control policies optimised with reinforcement learning to trial prosthesis-assisted gait in non-steady-state locomotion settings. Conditioning the prosthesis' control on its user's walking intent is shown to mitigate the need for compensatory movements from the human agent.

The application of our gait synthesis agents to test prototype passive devices is also explored. We create a simulated model of a compliant prosthetic foot, and procedurally adapt our motion synthesis to tackle varying slopes and tripping hazards. We identify and quantify potential benefits provided by the compliant design and contrast it against data from an external experimental study.

In conclusion, we identify a series of challenges and solutions for dynamically simulating and evaluating intent-driven lower limb prosthetic interventions. Our methods contribute to a framework for refining assistive device design, facilitating the transfer of the next-generation assistive devices from the concept stage to real world systems.

Statement of Originality

This thesis presents the original work of the author. All other work and content included that was produced by others has been indicated and referenced as such.

Copyright declaration

The copyright of this thesis rests with the author. Unless otherwise indicated, its contents are licensed under a Creative Commons Attribution-Non Commercial 4.0 International Licence (CC BY-NC).

Under this licence, you may copy and redistribute the material in any medium or format. You may also create and distribute modified versions of the work. This is on the condition that: you credit the author and do not use it, or any derivative works, for a commercial purpose.

When reusing or sharing this work, ensure you make the licence terms clear to others by naming the licence and linking to the licence text. Where a work has been adapted, you should indicate that the work has been changed and describe those changes. Please seek permission from the copyright holder for uses of this work that are not included in this licence or permitted under UK Copyright Law.

Acknowledgements

It is appropriate that this thesis is written in first-person plural. Completing my PhD would not have been possible without my friends and family, so for the Statement of Originality to stand, I really ought to indicate this. Since essential contributions like recipe clarifying hotlines with my mother, or soul-healing boardgame nights with friends can't fit well into my reference manager, I must attempt to express my gratitude in this section.

Firstly, I wish to express my gratitude to my supervisor, Dario. Thank you for believing in me, and for showing how to rely on my curiosity to guide my work. Thank you for always making time to give advice, and for seeing everyone in your lab not only as researchers but as people.

Dario introduced me to many communities that welcomed me, without whom the PhD would have been dreadfully isolating. Thank you to all in our lab for treating each other with kindness and respect. I wish to extend special thanks to Deren, Hsien-Yung, the rest of the Eurobench team and Hristo for helping me find my footing in my first years. On the other end of my project, I wish to thank Agnese, Alex, Bruno, Irene, Pranav and Rejin for providing valuable feedback and comments for this thesis or otherwise. Many thanks to the student community from my CDT, specifically Michael for his advice on modelling socket properties. Thank you to all with whom I collaborated with and endured the process of publishing, including Annika and Yongkun.

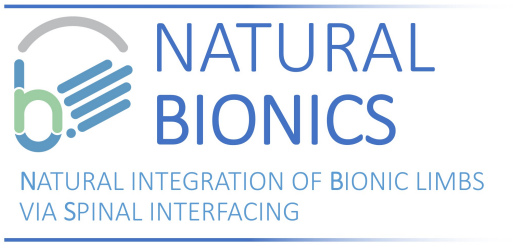
Meeting the Natural BionicS team during SSNR was a highlight every year. It took a while before I realised how my PhD could fit in well with the rest of NBS; once we figured it out, my PhD finally took proper shape. Special thanks to Matteo, Anna and Manuel from the IIT team, without whom the entire SoftFoot chapter would not have been possible.

Thank you Joan for sharing my passion for virtual worlds, and for commiserating over arcane bugs and merge conflicts. Being part of the Artanim team made for a memorable summer.

Thank you to my parents Ágnes and Károly, for their love and support. Together with my brother Ákos and sister Lilla, they showed the way leading to who I am today. I'm grateful to my partner Eszter, who is always in my corner. Thank you for brightening my days and for instilling in me the good sense to restrict myself to a single pun in this thesis. Szeretem mindannyiótokat, köszönöm hogy bíztatok bennem!

This work was supported by the UKRI CDT in Prosthetics and Orthotics (Grant No. EP/S02249X/1) and the Natural BionicS initative (Grant agreement ID: 810346).





Preface

Behind the community centre of my home town, was a small, overgrown sculpture park. Its centre piece was a mass of rusty metal geometric shapes, stacked seemingly haphazardly and covered in flaking red paint. I had been visiting it for years before I realised on a windy day that it was an articulated system of pendulums, and was moving. With a (cautious) push it came to life, the red pieces oscillating chaotically before joining together in sync. Eventually, they settled themselves back into an upright position, returning to a stability that appears inanimate on a cursory look.

The joy of applying motion to a static structure seemed like the perfect analogue to why I wanted to study the biomechanics of movement, and wanted to record the sculpture mid-swing for a figure in this thesis. By the time I could visit, the community centre with its garden had been converted into a four storey car park, and the sculptures were disassembled and put into a storage facility. However, I managed to find an old picture of it online. It didn't feel right to leave it at that, stuck in one pose. The picture served as reference. I modelled its structure and simulated it in motion. The red pendulum was swaying again. Only this time, I did not have to worry about breaking it when pushing.



"Red Pendulum" by István Harasztÿ, and its model in the MuJoCo physics engine. Photo reused with permission from user RTamas on www.kozterkep.hu.

Abbreviations

σ standard deviation.

AAN assist-as-needed.

ASIS anterior superior iliac spine.

BF Biceps Femoris.

CAD computer-aided design.

CNS central nervous system.

CoM centre of mass.

CoV coefficient of variation.

DoF degree of freedom.

DRL deep reinforcement learning.

EMG electromyography.

FSM finite-state machine.

GC gait cycle.

HD-EMG high-density electromyography.

HL high-level.

HMI human machine interface.

IED inter-electrode distance.

IIT Istituto Italiano di Tecnologia.

IMU inertial measurement unit.

IZ innervation zone.

KDE kernel density estimation.

LL low-level.

LSTM long-short term memory.

ML machine learning.

NN neural network.

P&O prosthetics and orthotics.

PCA principal component analysis.

PD proportional-derivative.

PPO proximal policy optimisation.

RF Rectus Femoris.

RL reinforcement learning.

RMS root mean square.

RMSE root mean squared error.

SCI spinal cord injury.

sEMG surface electromyography.

SO Soleus.

SPD stable proportional-derivative.

SVM support vector machine.

TA Tibialis Anterior.

TCN temporal convolutional network.

VR virtual reality.

ZMP zero moment point.

Contents

Abstract **3**

Acknowledgements **5**

Preface **7**

1 Introduction **35**

1.1 Motivation 35

1.2 What is intent? 39

1.3 Objectives and contributions 41

1.4 Outline of the thesis 42

1.5 Publications 44

2 Background **46**

2.1 Upper limb device simulation 49

2.2 Lower limb device simulation 54

2.3 Hierarchical control schemes 57

2.4 Neural-driven control 59

2.5 Simulation tools 60

2.5.1 The Modular Agents package 62

2.6 Virtual character animation 63

2.6.1 Motion matching 64

2.6.2 Closed-loop control 65

3 Myoelectric Knee Motion Estimation 67

3.1 Introduction 68

3.2 Materials and methods 71

3.2.1 Subjects 71

3.2.2 Experimental setup 71

3.2.3 Experimental procedure 73

3.2.4 Signal preprocessing 74

3.2.5 Bipolar sampling selection 75

3.2.6 Knee angle prediction model 78

3.3 Results 80

3.3.1 Connection between consistency metrics and model performance 80

3.3.2 Effect of the orthosis on the sEMG recorded signals 82

3.3.3 Spatially robust regression with location sampling 84

3.4 Discussion 87

3.5 Conclusion 91

4 High-level Locomotion Intent Estimation 93

4.1 Introduction 93

4.1.1 High-level locomotion intent representations 94

4.1.2 Input modalities . . . 97
4.1.3 Chosen representation . . . 98
4.2 Methods . . . 99
4.2.1 Walking path trajectory . . . 99
4.2.2 Temporal Convolutional Network . . . 103
4.2.3 Pilot study . . . 104
4.2.4 Follow-up study . . . 105
4.2.5 Conditions compared . . . 117
4.3 Results . . . 119
4.3.1 Eigenvalue extraction . . . 119
4.3.2 Walking path estimator . . . 120
4.4 Discussion . . . 122
4.4.1 Temporal Convolutional Network . . . 122
4.4.2 Limitations . . . 126
4.4.3 Assumption of nonholonomy . . . 128
4.4.4 Intent synthesis in simulation . . . 130
4.5 Conclusion . . . 131

5 Locomotion Synthesis for Simulated Prosthesis Use 133

5.1 Introduction . . . 134
5.2 Materials and Methods . . . 138
5.2.1 Locomotion intent synthesis . . . 138
5.2.2 Gait policy . . . 141

5.2.3 Virtual device . . . 147
5.2.4 Simulation conditions . . . 149
5.3 Results . . . 149
5.3.1 Gait policy . . . 149
5.4 Discussion . . . 156
5.5 Conclusion . . . 160

6 Modelling Locomotion with the IIT SoftFoot 161

6.1 Introduction . . . 161
6.2 Methods . . . 164
6.2.1 SoftFoot in MuJoCo . . . 164
6.2.2 Stance prediction and motion adaptation with footbase trajectories . . . 167
6.2.3 Locomotion task . . . 170
6.2.4 Modelling the mechanical interface . . . 171
6.2.5 Stability and evaluations metrics . . . 172
6.3 Results . . . 174
6.3.1 Gait policy with SoftFoot and a Rigid foot . . . 174
6.3.2 Comparison with experimental data . . . 178
6.4 Discussion . . . 179
6.4.1 Ablation studies . . . 184
6.4.2 Future work . . . 185
6.5 Conclusion . . . 187

7 Conclusion **188**

7.1 Summary of thesis achievements . . . 189

7.1.1 Understanding locomotion intent . . . 190

7.2 Future work and potential applications . . . 191

7.2.1 Robust bipolar sEMG interfaces . . . 191

7.2.2 High-level intent estimation . . . 192

7.2.3 Motion synthesis . . . 193

7.2.4 Humanoid locomotion model . . . 194

7.2.5 Architectural changes . . . 196

7.2.6 Other applications . . . 199

7.2.7 Learning environments . . . 201

7.3 Next steps: Intent-driven powered prosthesis control with empirical uncertainty 201

Bibliography **206**

Appendices **234**

A Motion Matching Implementation **235**

A.1 Feature extraction and labelling . . . 235

A.2 Root animation . . . 239

A.3 Nearest neighbour search . . . 239

A.4 Inertialisation blending . . . 242

B Pilot results **245**

C Motion Capture and sEMG Recording Protocol **250**

D Adversarial Imitation Learning for Policy Transfer **258**

E Motion Matching for Feedforward Control of Prosthetic Legs **262**

F Modelling Sensory Deficit in Prosthesis Use **265**

G RL configuration file for locomotion policy **269**

List of Tables

3.1 Median validation RMS error and σ for 100 models trained across subjects and muscles. Values in radian. 81

3.2 Median validation RMS radian error and σ for 100 models trained across subjects and muscles, when evaluated on all valid combinations. 84

4.1 The non-causal pilot model's euclidean distance error and its standard deviation, averaged across the five cross-validation folds. Rows represent the position of the sample in the trajectory horizon, in terms of time. 37188 frames were used in total to quantify performance. 105

4.2 Model-in-the-loop λ learned in 5-fold cross-validation (Method 2) and lowest cost λ found in bilevel optimisation (Method 3). 120

4.3 RMSE (m) of the 5-fold cross-validation results shown in Figure 4.13. 122

4.4 r^2 of the 5-fold cross-validation results shown in Figure 4.13. 122

4.5 Separating the r^2 along the anterior-posterior (AP) and mediolateral (ML) axes and calculating their ratio when using Method 2, for the EMG-only and the pose-only models. Mean values and one standard deviation reported. These results indicate contributions to inferring turning direction from the two input modalities. 126

5.1 Performance of non-prosthesis (condition 1), non-intent-driven prosthesis (condition 4) and intent driven prosthesis (condition 5) during straight walking. . . 155

6.1 Mechanism modelling complexity comparison. In digital editions of this thesis, most elements in this table are clickable, linking to relevant sections of the online MuJoCo documentation if available. Legend: *not yet implemented*, implemented, **selected**. 165

6.2 Pose error during straight walking and level ground walking with turns, calculated using Equation 5.7. Quantified for both SoftFoot and rigid foot useage, for unilateral (UL) and bilateral (BL) prosthesis use. 178

A.1 Chosen feature set for current pose description. Matching these features reduces discontinuity when switching between frames. 236

A.2 Chosen feature set for future state description. Matching these features allows adhering to directable goals (i.e. user input in desired walking path). 237

List of Figures

1.1 Failure of transferring motor skills from a full humanoid policy to a virtual prosthetic ankle controller trained offline. *a)* Very high validation performance can be seen making predictions with the prosthesis controller and compared against the output of the full humanoid policy offline. *b)* When applying the inferred joint torque the prosthesis, the user loses stability over time. The humanoid's suggestion for recovering balance shown in orange. The prosthesis controller model was not trained on recovery and keeps providing the stereotypical pattern it has seen offline. *c)* If we keep using the offline-trained prosthesis policy, we can update the humanoid policy to adapt to it. This converges to an improved, but still sub-optimal gait pattern. 38

1.2 Illustration of why relying solely on straightforward performance metrics such as root mean square angle deviation can lead to misleading evaluations. Above shown data are a generated example of joint angle prediction. Both prediction methods would score the same value on root mean square error, but the prediction shown in green is clearly more meaningful. 39

2.1 Flowchart of the sources queried to provide additional coverage during scoping review. 47

2.2 The role of visual and physics simulations in the P&O design process. Several development iterations can be performed with the virtual device while preparing for experimental validation. 50

2.3 Illustration of three different strategies for evaluating a myoelectric controller with a virtual environment for visual feedback. *a)* Cursor-based tracking of target time-series in cartesian space. *b)* Pose tracking in prosthesis' kinematic space. *c)* Functional tests in task space, position of participant's forearm synchronised to the virtual environment. MuJoCo model of the Johns Hopkins Modular Prosthetic Limb was used as the prosthesis model, imported to the Unity engine (John Hopkins Applied Physics Laboratory 2017; Kumar 2015). 51

2.4 Potential use cases for interactive, virtual and visual simulation of P&O systems. 53

2.5 The locomotion task separated into three levels of abstraction. The HL intent of planning to travel in a certain direction is translated to body segment trajectories, then muscle or motor activation (Flash et al. 1985). This structure is comparable to the three level model of voluntary control proposed by Hollerbach et al. (1982). 57

2.6 Relevant software tools for performing biomechanical simulations, and different roles they play. *Phyiscs engines* calculate and integrate multi-body dynamics over time, enforce constraints and contacts (Erez et al. 2015; Sherman et al. 2011; Geijtenbeek 2021). *Input/Output* protocols allow communication between external control systems and the simulated actuators. *Rendering engines* visualise and display the resulting movement. *Integrated Development Environments* provide tools to edit models and environments, script behaviours and controllers, and run simulations and learning (Geijtenbeek 2019; Unity Technologies 2021; Delp et al. 2007; Rohmer et al. 2013; Koenig et al. 2004). 61

3.2 The range of neighbouring electrodes considered for each electrode when determining the bipolar combinations. This schematic shows the combinations extracted for two example electrodes. The "origin" electrode (orange) was eligible for pairing with all the electrodes (blue) from its region of interest. This resulted in up to 16 possible bipolar combinations for each electrode. 74

3.3 Illustration of the three envelope measures within a gait cycle that were used to quantify signal consistency for the RF. Channels are considered more consistent if these measures don't change from cycle to cycle. Phasor representation of the continuous progress along gait cycle is shown above the curve. 76

3.4 Metric selection for further analysis based on model performance on reduced data set. Angle prediction error on validation data is compared for one model trained for each muscle of the first three subjects. The labels for each plot correspond to the metric used to determine the channel locations for that group's TCNs. The "-Agreement" label specifies the least consistent electrode location based on the "Agreement" method. 77

3.5 Schematic of the 1-dimensional TCN architecture, mapping the windowed sEMG signal $(1 \times W)$ to the user's knee angle (1×1). W stands for the width of the input window, either 1000 or 500 samples. 78

3.6 RMS error of angle predictions from TCNs trained on the combination with the best consistency score (blue) vs. the middle electrode combination (orange), for two different input window durations. The kernel density, the box plot and the raw data are shown for the different conditions. 81

3.7 Visualisation of the average cross-cycle values of the sEMG envelopes' (a) integrated area, (b) maximum peak amplitude, and (c) maximum peak location in terms of GC percentage (with 0% corresponding to the heel strike) for a subject's SO (1) without and (2) with the orthosis. Each of the graphs show the spatial distribution of the metrics across the HD-EMG grid, and the values shown correspond to the vertical differential between 2 monopolar electrodes (therefore showing a 12×5 grid as opposed to the original 13×5 monopolar grid. 83

3.9 Predicted knee angle RMSE of each muscle group across conditions when trained on the entire grid (Full) versus only the SENIAM bipolar combination in the centre of the grid (Single). The condition of wearing the knee orthosis is also shown. 85

3.10 Average model performance across all valid bipolar combinations of the grid. Models were trained either with electrode combinations sampled uniformly from the grid, or solely with the middle combination (electrodes 31 and 33). The hue determines whether the model was trained and evaluated with data recorded while wearing a knee orthosis. 86

3.11 Knee flexion predictions with data recorded with Delsys bipolar electrode, with models trained on signal from the HD-EMG grid. Two subjects shown, subject 3 with the RF (a) and TA (b), and subject 1 with the BF (c) and SO (d). "Full" refers sampling the input channel uniformly from the entire grid, "Middle" refers to only training with the middle location. The angle from motion capture is overlaid. 86

4.1 Concept for collecting the data set to construct a neural locomotion controller. Motion data can be paired with biosignals like sEMG and body posture during demonstrations of walking. Feature extraction can be performed on both motion modalities to focus the data to represent desired aspects of the behaviour. In the case of sEMG this is performed via TCNs, while for the motion capture the parameters of the critically damped trajectory model play this role ($\mathbf{v}_0$, $\mathbf{v}_d$ and $\mathbf{a}_0$). A mapping can be learnt from these paired embeddings. Once an intent estimator is trained, it may be applied on novel biosignal data coming from the same subject, to forecast their desired walking path. The predicted intent may then also be mapped back to full body kinematics via a method such as motion matching (covered in Appendix A). Anterior view of subject show, for full marker and electrode placements see Appendix C. 101

4.2 0.5 second window of 8 channels of bipolar sEMG, processed by the TCN and regressed to the walking path trajectory, represented in the local pelvis reference frame. Each trajectory has a corresponding 8 channel window of sEMG, in a sliding fashion. The right column shows the outcome variables: 1 second long horizon of the pelvis' position, relative to the pelvis' position and orientation at the trajectory's start. The event shown is a 90° turn to the right. 102

4.3 Experimental setup with one of the subjects at the Imperial College Biodynamics lab, showing motion capture system, the Conventional Gait Model markerset (Leboeuf et al. 2019), floor markers and bipolar sEMG electrodes. 107

4.4 Top down schematics of the surface markings applied on the motion capture volume, using which the motion trials could be guided. The A-D distance was 2.3 m, the A-B distance was 7 m, and A-AB was 4.7 m. To the right we show 2 out of the seven motion trial type descriptions. For the full list see Appendix C. 109

4.5 Bipolar sEMG signals recorded during one of the angular turns trials. The (L)eft and (R)ight side's Erector Spinae (ES) and oblique (OB) muscles' EMG is shown in the top four rows, the Soleus (SOL) and Peroneus (PER) muscles' signals shown in the bottom four. 110

4.6 The effect of the fixed eigenvalue λ on the velocity momentum/starting acceleration a_0, determined in the inner optimisation level. Values plotted against progress along one of the angular turn locomotion trials. Too high magnitude λ is mitigated by extreme values of a_0 to prevent too sharp turns, exceeding plausible values. 114

4.7 The effect of the fixed eigenvalue λ on the equilibrium/desired velocity v_d, determined in the inner optimisation level. With appropriate λ the fitted time series are continuous despite no constraint placed on this property. Too low magnitude λ is compensated by extreme values of v_d to reach the needed velocity by the end of the 1 second trajectory, exceeding plausible levels. 115

4.8 The effect of the fixed eigenvalue λ on the estimated starting velocity v_0, determined in the inner optimisation level. This parameter is less sensitive to λ than v_d and a_0. The model can take advantage of higher stiffness (and therefore faster convergence to v_d) by permitting greater variation in v_0. 116

4.9 Residual of the inner optimisation level across randomly sampled trajectories from one of subject 3's locomotion trials, conditioned on the eigenvalue λ. Finding the λ with the lowest expected residual is the goal of the outer optimisation level. A Matérn kernel was used, with a noise level estimated empirically by evaluating the loss function 10 times on the same value of λ (we selected -2.7). The length scale and output scale of the optimisation were adapted to maximise the marginal likelihood. 95% confidence interval shown 117

4.10 Three variant models to regress biosignals to walking path trajectories. All three may include either or both posture and EMG signals. We use an illustration of a virtual humanoid, but these models are applied on experimentally recorded data. 118

4.11 Training and validation learning curve examples for subjects 1 and 6. One standard deviation area from 5-fold cross-validation shown around curves. 120

4.12 Time series visualisation of the inference of the trajectory, plotted separately for the local forward and left/right directions at 1/3, 2/3 and 1s along the future trajectory horizon. A dashed line indicates on the time series results the time instance from which the trajectory on the left is shown. A virtual humanoid is shown as an illustration in the top left to aid interpretation, but these results are experimental with human subjects. This example is for subject 5 on multimodal validation data using the curve model, specialized for the angular-turns dataset. Errors in the prediction are correlated across trajectory points and depend on the type of motion performed (higher during turning). The inference also exhibits low-level constant biases, such as underestimating the magnitude of lateral motion. 121

4.13 Results comparing 5-fold cross-validation performance for the three Methods and three input conditions (organised per column), trained on all motion trial types and evaluated with reserved validation data from the "Angular turns" trials. We show all folds in each category, folds from a given subject labelled via colour, as well lines connecting the mean performance per subject to highlight cross-subject trends. *Top row:* Overall r^2 of the entire trajectory per subject. Estimating the mean trajectory from the dataset (which is not zero due to the non-holonomic nature of walking) would yield an r^2 of zero. *Middle row:* RMSE across the entire trajectory, with units of m. This measure is sensitive to the walking speed of the subject. *Bottom row:* Mean r^2 performance for the three Methods and three input conditions, quantified for the samples along the trajectory separately. Two σ from the 5-fold cross-validation of all subjects shown around each curve. It should be noted that most of the σ arises from inter-subject variability. 123

4.14 The activations of the second to last layer of the network trained in Method 2 can be interpreted as the parameters of the damped trajectory model. Here we compare on one of the "Angular Turns" validation trials of Subject 4 the parameters inferred by the network from EMG and posture, compared to directly fitting the parameters to the motion capture trajectory using least-squares optimisation. 126

4.15 Resulting path from synthesised motion from trajectory horizon from recorded motion capture, i.e. the ground truth of the intent estimator. The 180° turn is incorrectly generated as coming to a stop, then a 90° counterclockwise turn. . . 129

4.16 Target velocity as high-level control driven trajectory generation. Could be conceptualised as 2-D analogue joystick interface. 130

4.17 Target position as high-level control, as determined by the intersection of a ray cast from the camera based on the cursor's position and the ground plane. Target velocity can then be oriented towards target position, its magnitude is clamped to a configurable maximum, otherwise it equals the velocity that would reduce the distance to zero in one second. 131

5.1 Developing a virtual lower limb device testbed. A predictive walking policy was learnt, incorporating information from experimental data-driven references. A gait pathology was then simulated. Addressing the induced biomechanical deficit was the goal of the device's agent. Additional conditions and tasks can be introduced to assess the versatility of the controllers. Prosthesis mesh used from the Open Source Limb project (Azocar et al. 2020). 134

5.2 Illustration of the features used in the nearest neighbour search for the motion-matching system, and the local orientation frame they were quantified in. The orange trajectory consists of the position and normalised velocity of the pelvis center of mass, projected on the horizontal plane at 1/3, 2/3 and 1 1 s in the future. This curve represents the upcoming walking path in the motion capture data set. 138

5.3 The control problem solved using DRL. The target kinematic state from motion-matching is modulated by the gait policy, which receives observations from the simulation state. These targets are then converted into joint torques through stable PD. The simulation and reference states are then updated, and the process repeats. In prosthesis use conditions, below knee DoFs are controlled by a separate prosthesis policy, which receives its observations. Both gait and prosthesis policies are trained using the same reward signal. 140

5.4 The local reference frame and the bounding box surface points used in the reward calculation. Segments distal to the forearm were not included in the reward calculation. 144

5.5 Snapshot of the non-steady-state locomotion learning environment. The vector pointing from the reference to the target location is the high-level intent provided to the gait synthesis agent. Once a target location was reached, a new one was generated which may necessitate turning. 147

5.6 Path tracking ability of the gait policy. "Stroboscopic" visual trail of the agent's poses made by following the learned gait policy recorded at 1 Hz. The agent was capable of performing continuous circuits without falling or deviating from the path. 150

5.7 Ground reaction force during walking as measured by virtual sensors placed on both feet of the character. Recorded at 1.3 $\frac{\text{m}}{\text{s}}$ straight walking. 151

5.8 Plantar flexion torque in one of the ankles during simulated gait. Early stance dorsiflexive and late stance plantar-flexive torque can be observed (Winter 1984). 151

5.9 Phase profile of the rule-based classifier used in a traditional Finite State Machine controller reproduced from (Simon et al. 2014), applied to the virtual user using an active prosthesis. 152

5.10 Collaborative performance evolution of learning the (pretrained) gait policy and prosthesis policy simultaneously (conditions 2 and 3). Providing the intent as additional observation yields no improvement in these conditions. 153

5.11 Collaborative performance of user-prosthesis system, when the gait policy was pretrained without a prosthesis and could no longer learn. Only the prosthesis policy was adapted (Conditions 4 and 5). Note that the gait policy is deterministic here, as it no longer needs to explore by injecting noise in its actions. This yields the higher performance at the start compared to Figure 5.10. 154

5.12 Plantar flexion torque output of the ankle of the non-prosthesis-user (condition 1), the prosthesis with no intent input (condition 4), and intent-driven prosthesis (condition 5). 154

5.13 Trajectory of the equilibrium ankle angle in the PD control parameters, set by the intent-driven (condition 5) and non-intent-driven (condition 4) systems when there's a step change between the virtual user's intent (going from standing to walking). 156

6.1 CAD model of the IIT SoftFoot (Piazza et al. 2024), rendered by Matteo Crotti. 163

6.2 Tendon constrained model of the IIT SoftFoot in MuJoCo. The geometry mesh is simplified compared to Figure 6.1, leaving out extraneous elements such as screws and bolts. Elastic tendons shown in red, rigid tendon shown in yellow, and the sole-stabilizing spring shown in blue. The via-points of the piece-wise linear tendons shown with grey spheres between the segments. 166

6.3 Simplified model without tendons. The joint axes of the sole DoFs shown with blue vectors. 166

6.4 Implementation of the rolling motion of sole elements. *a)* Each sole element is modelled as having two parallel hinge degrees of freedom with a lateral axis. By constraining these two degrees of freedom to share an equal value (through MuJoCo's constraint enforcement tools (Google DeepMind 2024a)) no additional frictional contacts need to be simulated to model rolling without sliding between adjecent links. *b)* The same constraint system is used to enforce the connection between the metatarsal link and the foot arch in cartesian space, closing the kinematic loop. *c)* Collisions with the ground are simulated with mass-less colliders formed of the axis-aligned bounding box of each element. Inter-element collisions are disabled to speed up the collision detection. 167

6.5 *a)* MuJoCo recreations of the uneven terrain (bottom left) and obstacle geometries in (Pace et al. 2023), rendered in Unity. Additionally, procedural stair and ramp generation is also demonstrated. *b)* MuJoCo model of the SoftFoot stepping on an obstacle while used by a locomotion agent. 168

6.6 Height field stepping environment. By specifying the elevation in a grid, triangular prisms can be extruded to form a continuous smooth surface. As the locomotion task drives walking to arbitrary positions within the area of the valley, tackling slopes at various angles and directions is required from the agent. The height field data are shown on the left in a 65×65 grid, covering a $20m \times 20m$ area and a highest point corresponding to a 60 cm elevation. 168

6.7 Stepping environments. Similarly as with the sloped environment, due to the random traversal, this environment involves both ascension and descension at various angles. The elevation of the blue platform is 10 cm. 169

6.8 Height of the ankle when stepping on and off of obstacles in simulation. Shaded areas indicate period between the two swing phases preceding and following the obstacles. The agent with the rigid foot adopts greater foot clearance when unloading from an obstacle to reduce the risk of tripping. 175

6.9 Height of the pelvis when tackling obstacles with an prosthesis. *a)* Results from simulation, dashed lines indicate the mid-stance when stepping on the chord-type obstacle. *b)* Corresponding experimental results from IIT for descending from a single obstacle. Graph starts with the contact with obstacle, repetitions aligned for $t = 0$ at the mid-swing of the prosthesis side, determined by the minimum knee angle. Two σ interval shown. The difference between off-obstacle mean height can be attributed to inter-session motion capture marker placement changes. 176

6.10 Maximal Lyapunov exponent approximations by fitting least-squares linear regression to short-term (λ_S) and long-term (λ_L) sections of the divergence curves, calculated for the SoftFoot and rigid foot model. 177

6.11 Histograms of the percentage of strides spent in stance phase for the prosthesis and contralateral sides for the two types of simulated prostheses investigated. . . 178

6.12 Foot and knee kinematics and dynamics compared between the simulated SoftFoot, simulated rigid foot, experimental SoftFoot and experimental Triton foot data when unloading from the chord-type obstacle. Experimental data provided by IIT and Pace (2023), determined using inverse kinematics and dynamics from motion-capture and load-cell data. All plots are aligned for $t = 0$ at the mid-swing of the prosthesis side, determined by the minimum knee angle. Two standard deviation interval shown. For an illustration of the motion see Figure 6.13. 180

6.13 *a)* The simulated rigid foot and the SoftFoot stepping on the chord-type obstacle. *b)* Stepping on the same type of obstacle with the Triton foot and the SoftFoot during experiments. 181

6.14 Evaluating the locomotion agent on a dynamically shifting height field environment. 182

6.15 Comparing the effect the feedforward motion target and the ground reaction force (GRF) observations have on learning motion skills with RL. *Left:* Learning a continuous straight walking from a looping gait cycle animation. "Policy collapse" phenomenon can be observed later during the learning. 5 repetitions used for each condition, 1 standard deviation interval shown. *Right:* Training policies in a Level Ground Walking (LGW) setting involving turns with the same four conditions. 185

6.16 Including the inverse kinematics (IK) pipeline for adapting the level ground walking to the sloped environment shown in Figure 6.6 increases zero-shot and asymptotic performance. 186

7.1 *a:* Training and predicting with a steady state knee angle estimator TCN for 0.8 m/s walking speed (top row, validation data), and predicting on a dataset of 0.3 m/s (bottom row). *b:* Estimating the two dimensional gait phase (as defined in Equation 3.2) instead of the knee angle, trained for 0.8 m/s (top row, validation data), and transferring to 0.3 m/s (bottom row, only first component shown for clarity). Subject 5's Soleus middle electrode data was used. 192

7.2 Humanoid model before and after adjusted to match new anthropometrics using scaling tools released in the Modular Agents package (Hodossy et al. 2023b). . . 198

7.3 Locomotion inside a virtual home from an outside perspective, and from the perspective of a camera attached to a human approaching a stairwell. 199

7.4 Spatial awareness through height field vision. By intercepting rays with the surrounding geometry, with origins arranged in a grid above the the humanoid, we construct a height map based on the distance each ray travels before hitting the ground. The resulting grey-scale image can be processed by a convolutional neural network or other encoding technique to create rudimentary "vision" for the agent. Yellow markers around agent show the intercept points for the rays, and the resulting image is shown in the top left corner. 200

7.5 Proposal to evaluate the performance of an intent estimator trained from experimental data in simulation. The left side of the figure illustrates capturing biosignals and training the intent estimator to predict motion goals extracted offline (after the experiment). The intent estimator could then be used to provide input to the controller of an assistive device. The right side of the figure shows how a simulated locomotion is different; in virtual environments the ground truth motion goals are available even online (during the experiment), but there are no biosignals to use in the intent estimator. Instead, we can fit a model on the empirical performance of the intent estimator to predict the level of error given the true target. This residual model could then be used in the simulated environment to provide inputs to the virtual prosthesis at a realistic level of accuracy. 202

7.6 Subject 4's parameter space residual distributions for one of the EMG + Pose models trained with Method 3, approximated using KDE. In the left subfigure, we show the marginal distributions for each separate element. In the right subfigure we reduce the dimensionality to two via PCA and show their joint distributions. 203

7.7 Subject 4's conditional KDE relating the residual of the predicted lateral v_0 to the true lateral v_0 for a model trained with Method 3 and pose-only inputs. The starting lateral velocity is consistently underestimated. 204

7.8 Yule-Walker partial autocorrelation function (PACF) of the mediolateral (ML) component of the starting velocity parameter's (v_0) residual, calculated for one of the models of Subject 1 from Chapter 4 on a validation "Angular Turns" trial. Method 3 used with the multimodal datasets. Additionally we provided the Akaike information criterion (AIC) measure for AR models of different size, suggesting a window size of ~ 40 ms to be appropriate to model the autoregressive properties of the residual. Repeating this analysis for the v_d paramter's residual in the mediolateral direction suggests a ~ 90 ms window 204

7.9 Time varying Gaussian process with length-scales and covariance scales fitted on example residuals through optimising the likelihood. A Matérn kernel with $\nu = 1$ was used. We sampled a random trajectory from the process, which should be different than the experimental signals, but resemble them in signal characteristics.206

A.1 Flowchart of the data processing steps, the file formats and the software used. The recorded motion capture data is exported as a BioVision Hierarchical file format (BVH). Using Blender, we convert it to the binary FBX format for efficient loading and asset management in Unity. Parallel to this, we process the BVH files with our custom scripts to label every tenth frame with the motion matching features (Hodossy et al. 2024), and export these data in a JSON or HDF5 format. Inside Unity, each FBX animation file is associated with a corresponding metadata file, meaning that the motion matching features do not need to be computed during runtime for the animation database. 238

A.2 Top-down view of the resulting walking path (white) when directing the motion matching process by a user attempting to guide the walk along a trapezoid track (black). Despite using displacement from animation, user based inputs with joystick allow accurate following of a pattern on the ground. 240

A.3 Summary diagram of subsystem roles and interactions performed every frame of the simulation. The Controller waits 10 frames between constructing a query vectors, formed from the pose features of the currently active animation and the trajectory features from the user input. This is sent to the Motion Matcher object, which has processed the input dataset and paired it with the metadata features. The motion matcher optionally defines ranges of frames valid for matching, and normalises and rescales the features (based on the parameters supplied by the controller). Once the Motion Matcher receives a query vector from the Controller, it calculates a cost for each valid frame in its dataset, and returns the time and clip index of the found nearest neighbour. If the animation frame is sufficiently far away in time (determined another configurable parameter) than the currently playing one, the Controller notifies the Animator of the change, and the Inertialiser to perform the blend. 241

A.4 Three points of view shown for each condition (rows). *Top row:* Feature space trajectory when walking only straight. *Middle row:* Trajectories when randomly generating turns and pivots (see Section 5.2.2 for details). *Bottom row:* Same except when each feature dimension is normalised separately (instead of normalising vectors based on their average norm). 243

B.1 Ground truth prediction frame to be approximated for intent estimation. The left columns shows the predictors: 0.5 second sized window of 8 bipolar channels recording bilateral trunk muscle groups. ES stands for Erector Spinae. The right column shows the outcome variables: 1 second long horizon of the pelvis' position, relative to the pelvis' position and orientation at 0 seconds. The event shown is an obtuse right turn. 246

B.2 Model trained and tested on angular turns. Notice at the start the user readjusted their position, walking backwards. This behaviour was not included in the training set, hence the model is unable to predict that very different behaviour. 247

B.3 Model trained and tested on angular turns. The model can not only predict the direction of turn, but also distinguish between the different angles of turning. . . 248

B.4 Model trained and tested on circular walking. All or nothing response is shown, where the model is effectively classifying which direction the turn is going to be and predicting a stereotypical curve, with very little in between. When a misclassification occurs a short spike in error can be observed along the left/right axis. 249

D.1 Differences in the structure of motion tracking vs Adversarial Style Transfer methods. Traditionally motion capture data is used for reference in adversarial learning, here we instead use the movement generated by a motion tracking agent. 259

D.2 Two reference phase space distributions of the left knee. Motion from another agent trained with motion tracking is shown on the left, and one from motion capture shown on the right. Small differences arising from details such as the joint centre location can bias the shape of the phase diagram distribution making it more challenging to replicate using Adversarial Motion Priors. We also illustrate the direction of the state evolution with arrows overlayed, representing the state transitions included for the discriminator. The goal of the discriminator is to guess that a given state and transition belongs to the reference distribution or the one of the trained agent. 260

D.3 When the target horizontal walking velocity changes, the agent quickly reorients the walking direction, despite no such behaviour being available in example data. 261

E.1 Walking with high DoF prosthesis, with the missing kinematic targets recovered from a second subject's dataset using motion matching. The yellow character is the motion from Dataset A (missing the right leg), the red character is Dataset B (only right leg used) and the blue character is the physically simulated RL locomotion agent. 264

F.1 Added obstacles that require learning stepping skills. Agent is performing knee hyperflexion even when no stepping up is imminent, due to the uncertainty in the environment and its pose sensing. 266

F.2 Foot height in a straight walking scenario with the sensory deficit, with three non-consecutive steps. The agent selects a step height on the left side to be able to comfortably clear a potential step. 267

Chapter 1

Introduction

1.1 Motivation

Commercially available prosthetics and orthotics (P&O) have yet to realise the full benefits they could potentially provide. Only about 15% of individuals who require assistive technology like P&O have access to it, primarily due to cost and availability (World Health Organization 2017). Even among these patients, as many as one fifth eventually abandon using prescribed P&O devices (Phillips et al. 1993). A key contributing factor to the high abandonment ratio is unmet user expectations of reliable performance and functionality (Shah et al. 2009; Hill et al. 2017; Biddiss et al. 2007; Kilgore et al. 2001).

State-of-the-art human machine interface (HMI) based prosthetic limbs and exoskeletons aim to address this need. These devices have the potential to restore ability to intuitively participate in activities that require positive power output (e.g. grabbing heavy objects, walking up stairs) via precise and powerful actuators (Tucker et al. 2015). However, in both the 2016 and the 2020 instances of the international P&O technical competition, Cybathlon, the pilots winning the Powered Arm Prosthesis Race were not using these prostheses, but rather body-powered devices with purely mechanical interfaces (Riener 2016; Cybathlon ETH 2020).

Decoding the user's intent and translating it into commands for a digital device with multiple

degrees of freedom (DoFs) is often either slow, inaccurate or requires extensive training and continuous concentration from the user (Oskoei et al. 2007a). Due to limitations posed by the signal quality and the implemented control strategy, passive and body-powered upper limb prosthetics still outperform their more expensive, electronic counterparts in various contexts (Schweitzer et al. 2018). Similarly, actuated exoskeletons have yet to establish themselves as a reliable alternative to wheelchairs, despite their many advantages (e.g. better circulation, less reliance on accessible environments)(Arazpour et al. 2013). Therefore, the issue of implementing both robust signal processing and efficient control strategies needs to be addressed during the development of practical and safe P&O (Clement et al. 2011; Tucker et al. 2015). This could broaden the range of patients for whom actuated devices are a worthwhile intervention. Consequently, a larger user-base would promote further progress in making these devices cheaper, lighter and more accessible.

Developing and tuning control systems is hindered by limited access to hardware, and the risk of harming users or equipment during test deployment of incomplete systems (Colombo et al. 2013; John Hopkins Applied Physics Laboratory 2017). A virtual P&O device situated in a visually simulated environment allows for a low-cost, safe and iterative design process of both devices and control algorithms (Perry et al. 2018). Using this approach, intent-driven HMIs can be evaluated qualitatively during initial stages of development, mitigating many of the risks from hardware based testing.

Simply simulating P&O hardware on its own is only part of the picture. There is much to be learned by involving a user to interact with the simulation. Not only can they provide valuable feedback through their experience, evaluating the success of a P&O device hinges on restoring functionality to the person wearing it. Furthermore, the underlying control problem significantly changes with a second autonomous agent influencing it. As the human attempts to achieve their goal through using the device, they react to the assistance received from it and adapt their movement. It is therefore an incomplete model to only take offline recordings of the biosignals P&O users or subjects without limb difference and use that to evaluate control systems.

In upper limb prosthetic applications this can be resolved by recording the signals of a subject in real-time and tasking them with controlling the virtual bionic limb, the simulation of which is synchronised with their signals. By being able to observe the state of the device (e.g., through visual feedback) the control can be evaluated in closed loop setting.

There is no straightforward way to reproduce this for orthotic or lower limb devices. In this case, the only way to reliably simulate the interaction between a fully virtual device and its user is to model the human as well. We illustrate how such an approach can uncover the weakness of an offline-trained machine learning (ML) controller on Figure 1.1. In this example, we modelled the human as a dynamically simulated locomotion agent, trained in a physics simulation using reinforcement learning (RL) (details of this method will be presented in Chapter 5). This agent observed the current state of its body, and output torques for each of its joints. By recording signals of simulated embedded sensors worn on the right leg, and pairing them with the torque output of the humanoid agent for the right ankle, we can create a dataset suitable for training a joint torque estimator model. We fit a simple feed-forward neural network on this dataset, essentially distilling the control policy from the full humanoid to one that only controls and uses signals of the affected leg. In this offline model we can achieve very high validation performance. However, when inside the online simulation the ankle torque was determined using this offline model, the locomotion became stumbling and asymmetric. Small errors add up over time and drive the motion to situations not present during the offline training. When we permitted the locomotion agent to adapt to the virtual prosthetic ankle with RL, ambulation was restored. However, the resulting gait pattern was still asymmetric, and less stable than the original one. This example shows that while often only offline evaluation of controllers is feasible, we should strive to validate our systems online. Simulating virtual subjects is one of the ways to attempt this.

There is growing precedent for using physics simulations and virtual environments for designing and evaluating P&O devices and the HMIs that control them (Betthauser et al. 2019; Nakamura et al. 2017; Perry et al. 2018; Tkach et al. 2013). These environments are ideal for initially training and testing control methods in an online manner before they graduate to real hardware. Due to the accuracy of the physics simulations, transfer of control policies to real life in a non-

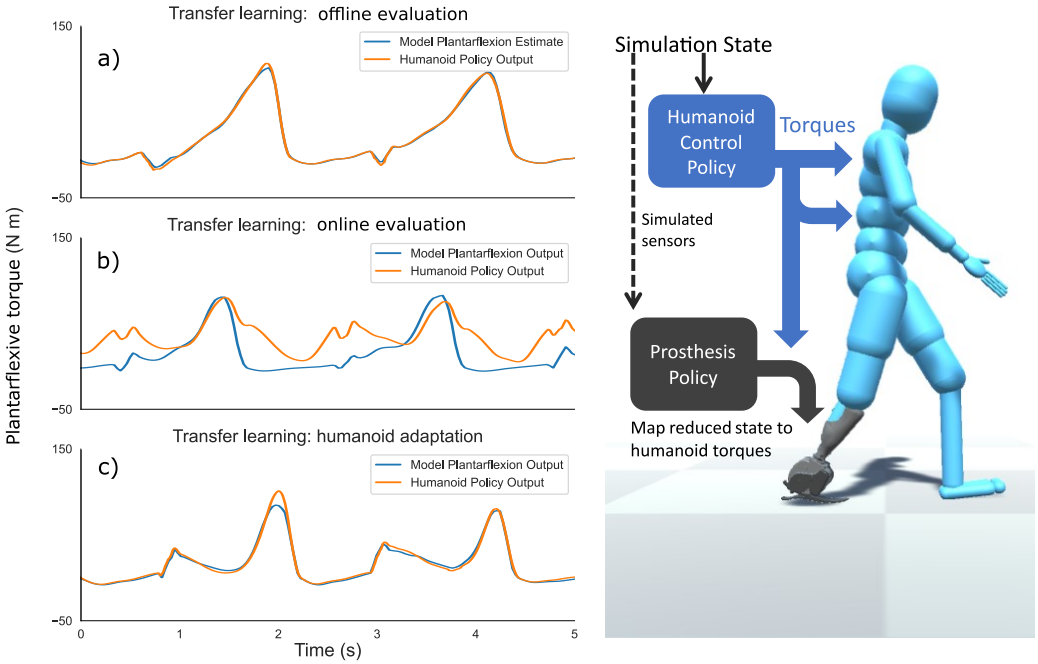


Figure 1.1: Failure of transferring motor skills from a full humanoid policy to a virtual prosthetic ankle controller trained offline. *a)* Very high validation performance can be seen making predictions with the prosthesis controller and compared against the output of the full humanoid policy offline. *b)* When applying the inferred joint torque the prosthesis, the user loses stability over time. The humanoid's suggestion for recovering balance shown in orange. The prosthesis controller model was not trained on recovery and keeps providing the stereotypical pattern it has seen offline. *c)* If we keep using the offline-trained prosthesis policy, we can update the humanoid policy to adapt to it. This converges to an improved, but still sub-optimal gait pattern.

wearable robotics context has already seen success (Barrett et al. 2010; Rusu et al. 2016; Peng et al. 2020), and recent developments show it is also possible for powered P&O devices (J. I. Han et al. 2022; Luo et al. 2024).

Developers of rehabilitative robotics gain important temporal and positional context of the synthesised motion. In simulation, functional tests can be performed as well. This type of evaluation can augment quantitative measures, which can be misleading if used in isolation (Fougner et al. 2012) (illustrated on Figure 1.2).

The limited access to patient participants for testing, and the costs of creating working prototypes imposes heavy constraints on the process of translating novel ideas for lower limb P&O (Windrich et al. 2016); more so than for upper limb settings. Computational models could be key to quickly iterate on experimental concepts, such as intent-driven powered P&O for locomotion activities that require a net-positive energy output (e.g., stair climbing). Restoring the

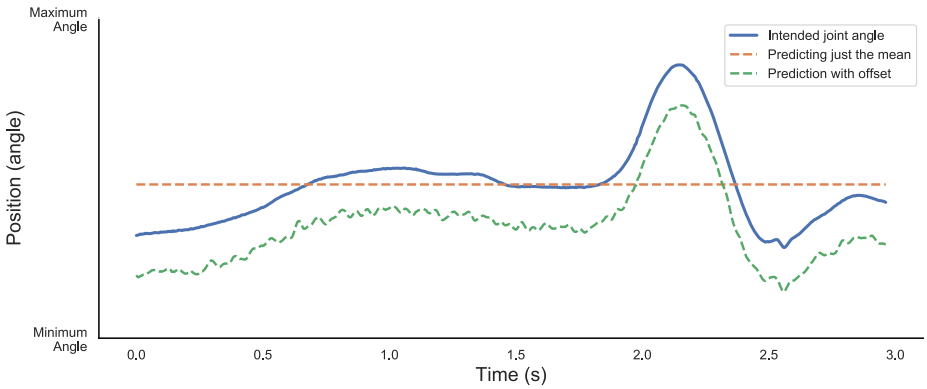


Figure 1.2: Illustration of why relying solely on straightforward performance metrics such as root mean square angle deviation can lead to misleading evaluations. Above shown data are a generated example of joint angle prediction. Both prediction methods would score the same value on root mean square error, but the prediction shown in green is clearly more meaningful.

ability to reliably ambulate is a top long term user priority for spinal cord injury (SCI) patients and neuroprosthesis users (Hill et al. 2017; Simpson et al. 2012; Kilgore et al. 2001; Ditunno et al. 2008). Meeting user expectations of functionality, speed and reliability would mean that P&O devices can be a viable path to improvement in their top life domain priorities, such as personal relationships, health, employability and access to services (Simpson et al. 2012; Hill et al. 2017).

1.2 What is intent?

Throughout this thesis we will frequently discuss motion intent. The concept of intent, and the process of estimating it form a key cornerstone of human-machine interfacing. However, its exact definition varies to suit different applications and author perspectives. Common examples include specifying desired motion directly through joint kinematics, through parametric description (e.g., walking speed) or selecting between discrete behaviours (e.g., different grasps or locomotion modes). Further examples are provided with more detail in Chapter 4.

Due to the flexibility with which this term is used in practice, we do not attempt to provide a rigorous unifying theory. However, to clarify its discussion, we describe our interpretation of it:

- We interpret intent as a signal (i.e., a codified message that contains useful information) in

systems that involve acting agents. This signal is permitted to (but does not necessarily) vary over time, and it may be conditioned on the state of the system. However, the intent signal should still retain a level of uncertainty even if the complete biomechanical state of the system is available (i.e., the intent signal is not determined fully as a function of the state).

- Intent signals should be closely related to how a movement behaviour is evaluated from the perspective of the agent performing the movement. In simpler terms, the signal should inform us regarding how someone wants to move, not just how they are moving right now.

- A good intent signal should be close to a free variable with low amounts of redundancy. This way, within some constraints, arbitrary samples of this signal should be able to be interpreted as valid inputs to control behaviours.

To summarise, an intent signal captures useful information regarding an acting agent's "free will" that is relevant to the studied behaviour and its evaluation. The exact encoding, dimensionality, distribution and constraints of this signal is up to us to define and measure, but should be chosen to suit the systems they influence.

As it is processed by the different stages of a control system, the information content of the intent signal may be split, compressed or decompressed, augmented by sensory information or re-encoded to different formats. In the case of hierarchical control strategies (discussed in greater detail in Section 2.3) it is useful to distinguish between high-level (HL), mid-level and low-level (LL) intent, containing the information relevant to the corresponding decision making stage.

In Chapter 3 we discuss a common LL representation of motion intent, described simply as the desired joint kinematics. A straightforward (but flawed) way to apply this signal for control would be to directly interpret it as target states for a wearable assistive robot, achieved via a method like proportional-derivative (PD) control. In contrast, Chapter 4 discusses the concept of estimating HL locomotion intent in more abstract formats, which would not directly influence actuation. Specifically, the current and desired horizontal walking velocity is used to describe

how walking behaviour varies over time based on what locomotion goals the subject wants to achieve. Chapter 5 shows how this more abstract representation can usefully condition the control of a simulated bionic ankle, while delegating the specific actuation necessary to drive the device to a separate control policy. In either case, an interface with the user is necessary to enable estimating their current state and future behaviour.

1.3 Objectives and contributions

The development of the next generation of P&O requires numerous design iterations and extensive user involvement. This is true for both low-cost passive devices, and active wearable robotics. Simulated models of P&O can catalyse this process. In particular for prototyping prosthetic devices, virtual environments can reduce the participation burden on the patient population. While our methods have broad relevance to modelling wearable assistive devices, we focus on lower limb prosthetic applications.

However, in an effort to model HMI-controlled lower limb prosthetics, we encounter an apparent contradiction. A real participant can not "wear" a virtual bionic leg and use it the way it is possible for virtual hands. Recreating a virtual model of the participant in the simulation restores the mechanical interactions necessary, but we lose access to biosignals used to drive the HMI controller.

In this thesis, we propose to resolve this contradiction through a hierarchical control approach:

- We identified a continuous HL representation of walking intent, as described by the current and desired horizontal walking velocity. This signal was able to be generated arbitrarily in simulation, and could direct diverse locomotion patterns in a virtual humanoid agent trained with RL.

- This HL intent signal served as a surrogate for the missing biosignals of the virtual agent. We applied it to control an intent-driven multi-DoF bionic ankle during non-steady-state

gait synthesised by its virtual user, and have shown it reduces the need for adapatation from the virtual user.

- We investigated the feasibility of estimating this HL intent representation outside of simulations using surface electromyography (sEMG) and body-posture measurements, through a convolutional network approach, and extracted interpretable gait parameters from it using physics-informed ML techniques.

- While versatile and powerful, ML models are prone to be sensitive to subtle changes between the conditions of training and deployment. We quantify the sensitivity of our ML models' susceptibility to common distributional shifts and suggest data collection methods to train robust estimators.

- The applications of P&O simulation go beyond modelling powered bionic devices. We demonstrated trialling a passive compliant prosthetic foot design in a virtual gait lab, and identified benefits that are paralleled from results of its real counterpart.

In summary, the aim of this PhD was to contribute dynamic simulations of the lower-limb prosthesis and its user in settings of changing locomotion goals, providing a virtual platform to test concepts for the control and structure of the assistive device. To model intent-driven powered devices, we also propose a low-dimensional quantitative representation of the changing locomotion intent, how to synthesize it in simulation and estimate it from sEMG and body posture in real life.

1.4 Outline of the thesis

Our objectives arose and evolved gradually, starting from our investigation of ML models of sEMG-based locomotion intent estimators. We detail the progression of our research questions in four technical chapters (3-6) that build upon each other's work.

Chapter 2 provides a broad overview of existing applications of simulated P&O, contrasting our objectives with more widespread upper limb implementations. Here we provide context

for the methods applied in this thesis from the state of the art, but the subsequent chapters contain their own literature reviews, providing motivation and background for their respective topics.

In **Chapter 3**, we introduce bipolar sEMG-based knee angle regressors using temporal convolutional network (TCN) architectures during steady-state level-ground gait. While the focus was on an exoskeleton control use case, the methods and findings are broadly applicable to powered lower limb P&O systems. This chapter also discusses the notion and feasibility of optimal electrode placements and how to leverage high-density electromyography (HD-EMG) data acquisition to maintain high performance across an area over the skin surface. Upon reviewing the outcomes, we questioned whether the model's success was primarily due to the consistent, repetitive nature of continuous straight walking. Could TCN-based motion intent estimators generalise effectively for time-varying behaviours?

Chapter 4 applies TCNs to a non-steady state locomotion dataset, where both walking speed and direction vary over time. We abstracted away the specific joint kinematics from the captured motion and instead predicted a low-dimensional representation of the forecasted walking path. Our investigation included both sEMG and body pose input modalities, discussing their strengths and limitations. Following promising results from offline analysis, we addressed subsequent research questions: How can the low-dimensional motion intent signal be related back to full-body kinematics? Is our low-dimensional representation of locomotion intent useful for controlling prosthetic legs? How can we trial our system in an online, closed-loop setting before involving patient participants?

Chapter 5 explores potential solutions to these issues, taking advantage of recent techniques in the related but often not examined field of virtual character animation. Simulated humanoid agents with responsive control schemes served as virtual participants to trial control policies for powered prosthetic ankles, optimised via RL. These agents paired full-body kinematics and kinetics to a parametrised representation of locomotion goals, which could serve as a surrogate signal for the virtual HMI of the prosthesis controller. We employed a very simple model of the prosthesis, and only evaluated it on level surfaces, therefore we adapted our research questions.

Can our approach predict the real-world impact of prosthetic design choices? How can we adapt our motion synthesis techniques to address everyday environmental challenges, such as slopes and trip hazards?

Finally, **Chapter 6** demonstrates a case study of testing a prototype prosthetic foot design, the SoftFoot from Istituto Italiano di Tecnologia (IIT) in our virtual environments. We extend our motion synthesiser to generate appropriate motion patterns for slopes, elevated surfaces and obstacles. We compare our results with experimental data recorded by IIT, which include trials of an amputee using a physical SoftFoot. This chapter demonstrates the application of our methods to passive prosthetics, which have a more direct path to impact than robotic ones. Limitations and suggested future additions to our methods are also discussed.

Chapter 7 provides a summary of the achievements and contributions of this PhD, and projects forward the ideas presented in the thesis to address limitations and translate results to the end-users' benefit. This chapter also introduces a series of appendices that summarise proof-of-concept experiments that highlight potential future applications of our methods.

1.5 Publications

First author publications reproduced in this thesis:

Chapter 3:

Hodossy, B.K., Guez, A.S., Jing, S., Huo, W., Vaidyanathan, R. and Farina, D. (2024). Leveraging high-density EMG to investigate bipolar electrode placement for gait prediction models. In: IEEE Transactions on Human-Machine Systems, 54, pp. 192-201.

Chapter 5:

Hodossy, B.K. and Farina, D., 2023. Shared Autonomy Locomotion Synthesis With a Virtual Powered Prosthetic Ankle. In: IEEE Transactions on Neural Systems and Rehabilitation

Engineering, 31, pp.4738-4748.

Relevant second author publications contributed to during the undertaking of the degree:

Guez, A., Hodossy, B., Farina, D. and Vaidyanathan, R., 2023. Transferring Gait Predictors Across EMG Acquisition Systems with Domain Adaptation. In: 2023 International Conference on Rehabilitation Robotics (ICORR) (pp. 1-6). IEEE.

Pereira, J., Halatsis, D., Hodossy, B. and Farina, D., 2024. Tackling Electrode Shift in Gesture Recognition with HD-EMG Electrode Subsets. In: ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP) (pp. 1786-1790). IEEE.

Zhao, Y., Hodossy, B.K., Jing, S., Todoh, M. and Farina, D., 2024. Delayed reinforcement learning converges to intermittent control for human quiet stance. In: Medical Engineering & Physics, 130, p.104197.

Caggiano, V., Durandau, G., Song, S., Tan, C.K., Wang, H., Hodossy, B., Schumacher, P., Gionfrida, L., Sartori, M. and Kumar, V., MyoChallenge 2024: Physiological Dexterity and Agility in Bionic Humans. In: NeurIPS 2024 Competition Track (under review).

Zhao, Y., Zhang, M., Hodossy, B., Jiang, J., Todoh, M., Farina, D., 2024. Neuromechanical Simulation of Human Postural Sway in the Sagittal Plane Based on a Hybrid Triple Inverted Pendulum Model and State-Dependent Intermittent Neural Control. In: IEEE Transactions on Biomedical Engineering

Chapter 2

Background

P&O devices have been applied as solutions to a wide range of injuries and health problems throughout our history; securing fractured bones, correcting posture, restoring missing limb function and aesthetics (Seymour 2002). Prior to the 20th century, most devices were unique, bespoke solutions crafted through artisanal skill. Standardisation and parametrisation of artificial limbs and their use was accelerated due to the increased patient population from the World Wars (Fliegel et al. 1966). Today, approximately 0.5% of the global population requires assistance from P&O, and the associated rehabilitation therapy (McDonald et al. 2021). This percentage is expected to double by 2050, in large part due to the demographic transition caused by increased longevity (World Health Organization 2017; Cohen 2003). This consequently correlates with an increased prevalence of age-related conditions that can benefit from P&O treatment, such as SCI, stroke, Parkinson's disease, and amputation due to trauma, dysvascularity and diabetes (Tucker et al. 2015). Global conflicts contribute to traumatic injuries and amputations in younger demographics (Goniewicz et al. 2023). Wearable robotic interventions offer greater restored functionality to a portion of these populations. Outside patient groups, they have the potential to extend the biomechanical capabilities or preventing injuries in their users (Bogue 2018). However, the development, clinical impact and adoption of these technologies have been slow-moving. This is partly due to the requirement for robotic solutions to be robust to the large degree of uncertainty they encounter during daily use.

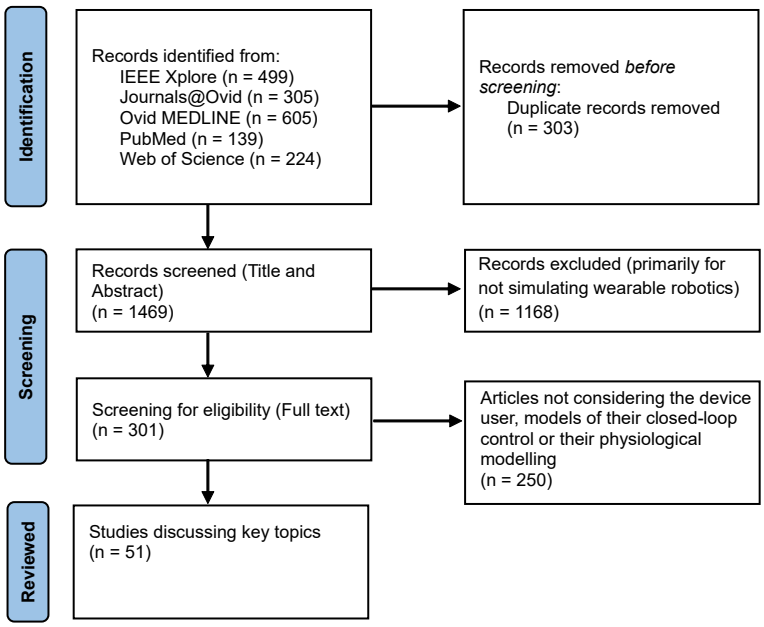


Figure 2.1: Flowchart of the sources queried to provide additional coverage during scoping review.

The standard engineering approach to tackle uncertain situations is through modelling (C. K. Liu et al. 2021). The design of the hardware and control of P&O is plagued by many such uncertainties, from both ends of the device; both our environments, and the humans using these devices are diverse and constantly changing. Attempting to bridge the gap between the two likewise offers an exponentially growing tree of choices. Modelling the behaviour of P&O systems, or creating holistic virtual representations of them through computational methods is a way to make sense of this problem and inform our decision-making.

We must clarify, that physical prosthetic simulators, which are worn parallel to the arm or leg of subjects without limb difference during experimental investigations are not the primary concern of this thesis (Sinke et al. 2022; S.-K. Wu et al. 2011; Shultz et al. 2013). Those devices are valuable tools to lighten the burden on the patient population required to test hardware and control strategies, but do not resolve many other hurdles involved with P&O development. Our focus is for solutions that partially or wholly model the device, its structure, control and usage computationally.

To assess the existing use cases and methods of computational P&O models, and to identify potential gaps we may address, we surveyed the literature. In addition to keyword-based

search and backward-forward citation search, we employed scoping literature review techniques to broaden our coverage of the prior work on P&O simulation. The following search terms were used for the title and abstract fields for journal articles, applied to the sources listed in Figure 2.1:

- (Prosthe* OR Orthosis OR Orthotic OR Exoskelet* OR Rehabilitation Robot* OR Wearable robot*)
- AND (Simulation OR Virtual)
- AND (Control OR Intent)

Simulation of upper limb exoskeletons and prostheses has greater precedence than their lower limb counterparts. In our literature search, we saw concepts relating to virtual bionic arms proposed as early as the 1980's (Abul-Haj et al. 1987), with similar work appearing in the 2010's for the lower limb (Hargrove et al. 2011). Within the filtered search results, 149 articles employed keywords relating to upper limb, in contrast to 75 for the lower limb. The following terms were used for this search in titles, keywords and abstracts of articles, combined with a boolean OR operator:

- *(upper limb, arm, shoulder, elbow, wrist, hand, forearm, transhumeral, transradial)*
- *(lower limb, leg, hip, knee, ankle, foot, thigh, transfemoral, transtibial)*

Most existing simulation work of P&O only involve modelling the device and its low-level actuation, either ignoring the user or modelling them through predetermined, open-loop control trajectories (determined e.g., via inverse kinematics and dynamics). This was the primary reason for excluding articles during the secondary screening of our search results. We also noted a particular challenge of involving users in a closed-loop online evaluation of a virtual lower limb P&O. Due to the greater potential for novelty and impact, as well as the identification of opportunities to apply recent techniques from the related field of character animation, this project pursued to contribute models of lower limb P&O use. A common pattern between

upper and lower limb settings is the additional design considerations needed to model orthotic systems. They involve overactuated parallel kinematic loops and may require modelling a large number of contact points for the device-user mechanical interface (Al-Shuka et al. 2019). Models of prosthetic systems instead can rely on series kinematic chains, reducing their complexity.

In addition to background on P&O simulation, we include below a brief introduction to relevant techniques from the field of humanoid animations that will be explored further in Chapter 5, alongside additional discussion of RL as a method to obtain motor control policies. We also identify several ML HMI techniques for processing myoelectric signals from the arm that merit adaptation and exploration for lower limbs, which became the focus of Chapter 3.

2.1 Upper limb device simulation

Virtual environments have been used by P&O developers and clinicians to model devices, with several decades of history (Hauschild et al. 2007; Soares et al. 2003; Abul-Haj et al. 1987), in particular for the simulation of upper limb devices. To validate a new concept for P&O hardware and its control, the patient, device, clinician and laboratory equipment all need to be available at the same place and same time. For this reason experiments are challenging to organise and results are hard to get and may be difficult to interpret. Evaluating and validating a novel design, from concept to results via a clinical study may take multiple years (Mitchell 2018). This leads to slow iterations on the design, which then need to go through the same process once more. To ease this issue, the problem of P&O engineering is often broken down into elements with more manageable sizes. Each of these subsystems can then be modelled and developed individually, in a modular fashion.

Through modular systems smaller, internal iterative development cycles can be formed (Figure 2.2). If issues are identified and resolved within these cycles, evidence based choices can be made earlier without the need of a full experimental study. By no means does this eliminate the need for validation and quantification with full scale user studies. It does, however, allow for more efficient use of the limited resources available to conduct them.

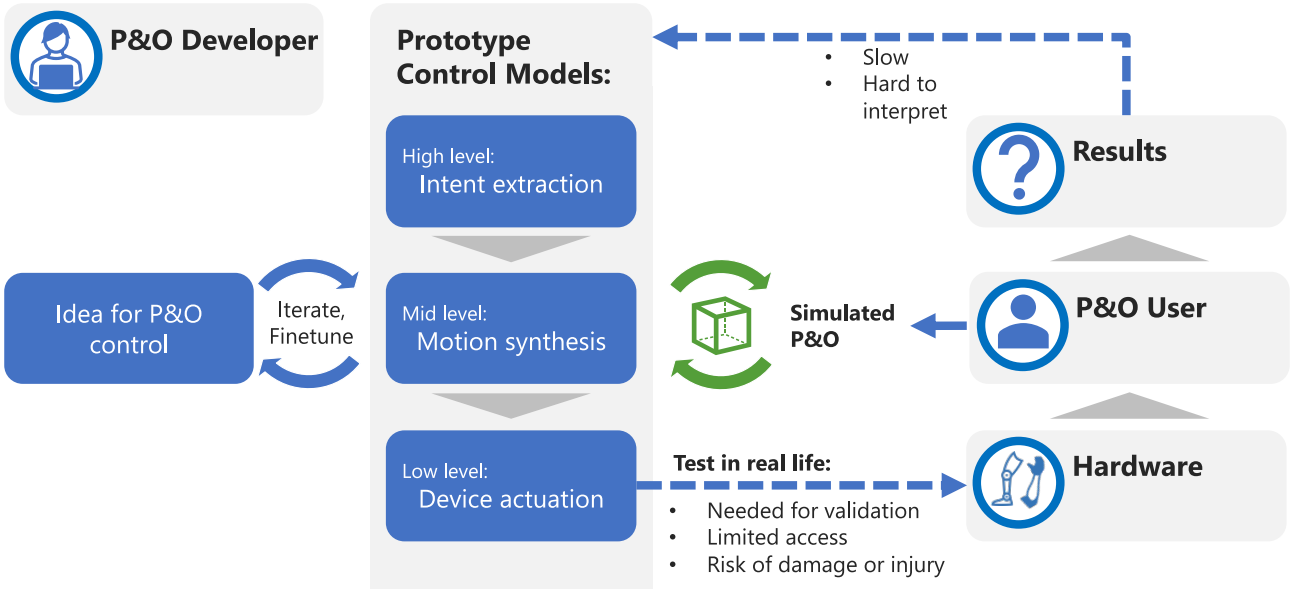


Figure 2.2: The role of visual and physics simulations in the P&O design process. Several development iterations can be performed with the virtual device while preparing for experimental validation.

The creation of modular control schemes requires the decomposition of the process into a semi-independent steps, forming a hierarchical structure. This concept will be further discussed in section 2.3. Aspects relating to the mechanical structure of the device too may be designed separately, determining key parameters independently of patient users (Della Santina et al. 2018; Fey et al. 2011).

HMIs that drive wearable robotics based on biosignals are commonly evaluated initially without consideration for the details of the hardware they'll control, and treated as an orthogonal problem. Virtual control spaces with two to three DoFs are commonly presented to users, e.g., as a task of controlling a cursor in a Cartesian coordinate system, either tracking a reference trajectories or following a "Fitt's law"-style target reaching scheme (Scheme et al. 2013). The virtual environment provides visual feedback, allowing for closed-loop control from the subject. Tracking error or target reaching success metrics are then calculated to quantify performance. The cursor-based approach liberates the researcher from the need of considering the end device, but it is not without disadvantage. The closed-loop nature of the system influences the accuracy, and therefore the type of visual feedback the subject receives will influence the outcomes. There's a notable drop in performance when the precise control space visualisation is swapped for a virtual prosthesis (Scheme et al. 2013). However, with evaluations through performance

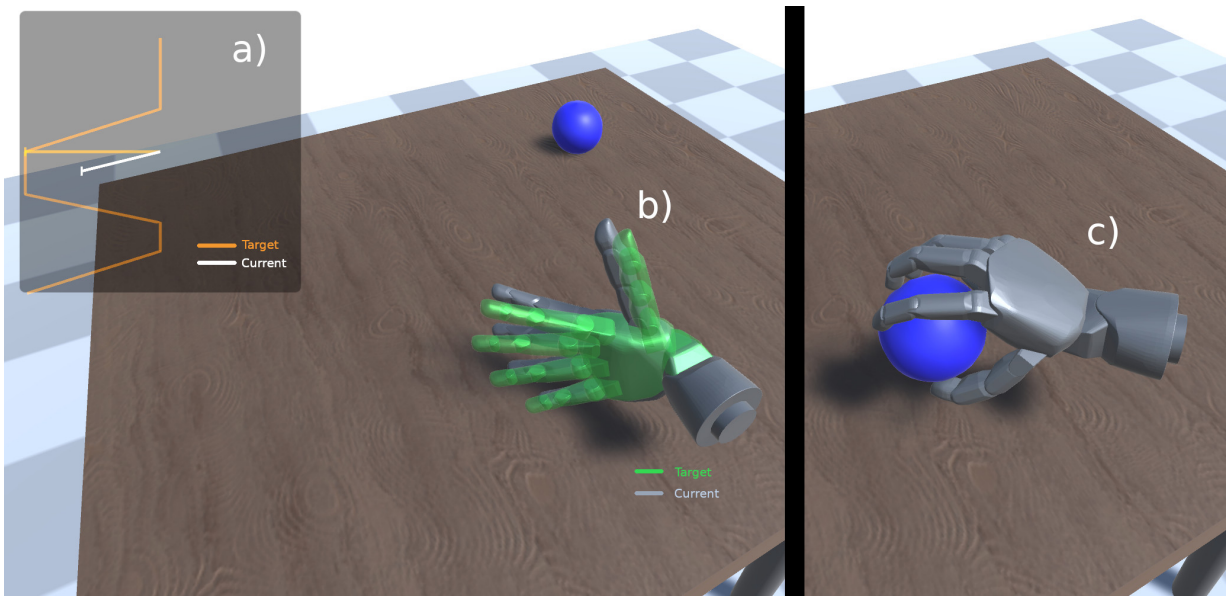


Figure 2.3: Illustration of three different strategies for evaluating a myoelectric controller with a virtual environment for visual feedback. *a)* Cursor-based tracking of target time-series in cartesian space. *b)* Pose tracking in prosthesis' kinematic space. *c)* Functional tests in task space, position of participant's forearm synchronised to the virtual environment. MuJoCo model of the Johns Hopkins Modular Prosthetic Limb was used as the prosthesis model, imported to the Unity engine (John Hopkins Applied Physics Laboratory 2017; Kumar 2015).

metrics we wish to indicate expected outcomes during prosthesis use. Therfore, it may in fact be appropriate to present the more challenging task of virtual limb control, instead of potentially overestimating accuracy with cursor environments (Gusman et al. 2017). Competency with virtual devices in target reaching tasks is shown to be correlated with capability with a real prosthesis (Simon et al. 2011; Hargrove et al. 2018). Even relatively high scores with cursor based methods are sometimes shown to pair with low performance when it comes to activities of daily living (Lock et al. 2005). In cases where the HMI is independent from the biosignals of the user and only rely on the mechanical interaction and configuration of the body, the user themselves is sometimes simulated prior to participant involvement. For upper limb applications, classic optimal control algorithms like Linear Quadratic Regulators may be used to create a closed-loop responsive user model (Ghannadi et al. 2018; Yang et al. 2022).

On-board sensing systems too can be simulated on virtual devices, modelling signals from sources such as inertial measurement units (IMUs) (Sharifi Renani et al. 2021; Vienne et al. 2017), distance and depth sensors (Zhang et al. 2019; Krausz et al. 2019), computer vision (Zhong et al. 2022), or even myoelectric sensing (Ma et al. 2024b). If a dynamic simulation is

used where the system evolves through forces and physical constraints, contact forces can be estimated for tactile sensing (Kluger et al. 2019). These sensory systems collect information both regarding the state of the user (proximate sensing) and that of the environment surrounding them (teleceptive sensing) (Krausz et al. 2019). The acquisition of this information is essential for any autonomous or semi-autonomous control of wearable robotics. Similarly, it is also of interest to the user, who, in cases of neurodegenerative disease or amputation, may have a loss off natural sensing capability, resulting in impaired motor function (Galan et al. 2015). Virtual environments also offer an excellent opportunity to test sensory feedback systems of bidirectional HMIs (such as vibrotactile, pressure or audio interfaces) with a user in the loop (Kluger et al. 2019; Hara et al. 2010; Chen et al. 2016). The properties of virtual sensor data will not match their real counterparts without error (Xiang et al. 2022), therefore any data synthesised this way will likely only augment datasets determined through experimental measurement. Despite this, there have been related cases in the field of robotics where even zero-shot (i.e., without re-tuning the controller) performance was successful with real users and devices, despite using synthetic data to learn the controller with (Akkaya et al. 2019; Radosavovic et al. 2024; Luo et al. 2024).

Virtual P&O have uses beyond determining control policies to be transferred to real life devices. Simulated prototypes can help uncover general insights regarding the class of controllers studied, for example hyperparameters such as the processing window size (Smith et al. 2011), or to test new ideas for the input signal processing and architecture of a control system (Karrenbach et al. 2022; Chu et al. 2006; Hauschild et al. 2007). Beyond developing, testing and evaluating controllers, virtual devices can also guide acquisition trials to collect examples for other data driven methods. If the conditions between data capture and deployment change, the controllers may have to extrapolate beyond observed examples and the performance may suffer due to this distributional shift (Subbaswamy et al. 2020). Some virtual environments can replicate functional tasks (such as pick up and place tasks), which may allow the collection of more suitable datasets (Kluger et al. 2019; Chappell et al. 2022; Bunderson 2014). The problem of distributional shifts is present in many aspects of HMIs. For example, the myoelectric skin-electrode interface may change its properties on the short- and long-term (discussed in greater

detail in Section 2.4), affecting the resulting signals. Some commercially available myoelectric prostheses offer virtual environment facilitated calibration trials, which can be initiated by the users of these devices to update the control parameters of their devices (Coapt 2024). During these trials not only the device adapts to its user, but also vice versa. As such, self-guided or teletherapy training sessions can take place using these methods. (Hargrove et al. 2018). The suitability of different devices and control strategies could also be assessed for a particular patient when they are referred to a prosthetist (Hauschild et al. 2007). These potential use cases are summarised in Figure 2.4.

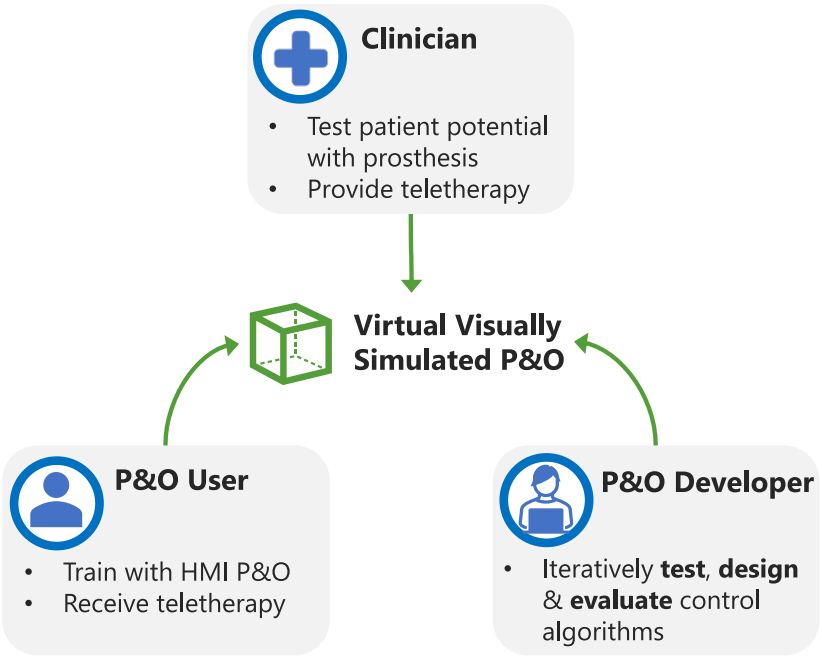


Figure 2.4: Potential use cases for interactive, virtual and visual simulation of P&O systems.

The level of fidelity of how the virtual environment is presented to their user is also worth considering. Relaying the virtual prosthesis' state through a screen necessarily leads to some loss of information. Depth information is lost (relevant for pick and place tasks), the subject cannot reposition the point of view to have a different perspective, and the embodiment illusion may be weaker as well (Nawfel et al. 2022; Kluger et al. 2019). Virtual reality (VR) headsets may instead be used to address these issues, which render immersive 3D scenes with head tracking for their user (Chappell et al. 2022). It is also notable that usually there is only kinematic coupling between the real residual limb and the virtual bionic limb via some form of motion capture. The extra mass and inertia is not experienced at the socket site. The lack of reaction

forces alters the user's perception of the activity and may influence behaviour (Simon et al. 2011). This issue can partially be remedied by the donning of a weighted cosmetic prosthesis, however, external forces arising from environmental interaction requires additional actuation (e.g. via an external robot arm) (Chappell et al. 2022). In lower limb device modelling this issue will be more relevant, due to the larger influence external forces have over the motion.

2.2 Lower limb device simulation

The development of lower limb P&O is both simpler than its upper limb counterpart in some aspects, while simultaneously being more challenging in others. On one hand, the task of locomotion readily lends itself for broad simplifications and assumptions. The diversity of tasks we perform with our hands is considered more varied than that of activities of daily living involving our legs (Jang et al. 2011), where the analysis of a steady state, level ground straight walking dominates the discussion (Schaik et al. 2019; Trent et al. 2020). This makes modelling locomotion tasks more straightforward. On the other hand, prototyping devices with novel actuation or control mechanisms faces additional challenges with prosthetic legs and feet. Not only are the power requirements greater for locomotion than for manipulation, so are the involved risks. Prototype devices will be sub-optimal and incomplete when they are first tested; a malfunctioning robotic hand may drop an object but a mistake in the control scheme of a leg could lead to a fall and injury. This necessitates additional risk mitigation strategies, and requires a higher level of technology readiness before lower limb devices are trialled. Paradoxically, the controllers need to be improved through the results from experiments, the lack of which prevents their acquisition in the first place. This is partially the reason why novel intent-driven control ideas are less frequently investigated for lower limb devices (Windrich et al. 2016; Cimolato et al. 2022), despite the greater prevalence of lower limb amputations by almost an order of magnitude (Farina et al. 2013a). For this reason there is a huge potential benefit to creating models of lower limb P&O use, lowering the barrier to entry of ideas to practice. A lot of the approaches and concepts from upper limb device simulation apply here as well, therefore the discussion below will focus on the differences and challenges.

In upper limb use-cases, one of the main benefits of virtual environments was the ability to provide visual feedback for user-in-the-loop testing. For activities involving the lower limb (i.e. standing and locomotion) visual feedback through VR can similarly guide the users in a controlled manner through virtual tasks, which has been applied for post-stroke rehabilitation (Deutsch et al. 2007b). However, while some sensory content can be replicated through these methods, kinaesthetics and the biomechanical state of the body cannot be reliably faked with a human interacting with the simulation. In the case of VR simulations of upper limb activities there is margin to deviate between the simulated state and the measured kinematics of the arm to reconcile the virtual environment and the user (e.g. to prevent penetration of the virtual arm with virtual objects). In contrast, even subtle attempts to show modified full body kinematics from the user's point of view during locomotion may result in motion sickness (Chang et al. 2020). Even more so, in many cases of locomotion assisting wearable robotics, the biomechanical deficit present in patient subjects simply prevents performing the movement in the first place without a real device. Outside non-weight-bearing settings, attempts for users to interact with virtual lower limb P&O are rendered impossible without ways to apply forces on their body.

Notable ways to escape this issue include device emulation (Caputo et al. 2014; Diaz et al. 2022; Bryan et al. 2021), where a more capable hardware system is being used to provide the missing external dynamics that should have arrive from the virtual device to the user. This still requires extensive instrumentation and hardware development, albeit one that is reusable to reproduce many different device behaviours. It also introduces additional constraints on the movement, e.g., allowing only straight walking on a treadmill to be investigated. A second approach is that of non-weight-bearing volitional lower limb control, in particular virtual ankle control. This use case is analogous to virtual hand environments previously mentioned, and is appropriate for human-in-the-loop development of devices (Hargrove et al. 2013; Chen et al. 2016) or for rehabilitation (Deutsch et al. 2007a; Deutsch et al. 2007b). Some positive effects of this type of training have been observed even on weight-bearing behaviours afterwards. However, this does not address the fact that lower limb devices need to be evaluated and tuned for use during locomotion.

The state of a wearable device is heavily determined by its user (either voluntarily or involun-

tarily). For this reason their models too must be evaluated in plausible kinematic and kinetic contexts. Speculative modifications of the human-device system's dynamics can be investigated by simulating the human body, the state of which can be constrained to match prescribed specifications, determined experimentally or analytically. This way the virtual device's operation can be kept in states relevant for its operation. Inverse dynamics methods can be used to recover forces and torques needed to explain the specified constraints, to which virtual devices can contribute to. However, this approach cannot allow deviation from the original motions, which could result from the assistance, nor can they model adaptation from the human (Aftabi et al. 2021; Shushtari et al. 2022).

This leads to the primary challenge of holistic virtual models of lower limb prosthesis use: the user's motion must be simulated as well, and their control policy approximated, removing the need to have a experimental subject in the loop. The level of detail of the humanoid model can range heavily based on the needs of the research performed, from simple double inverted pendulum "compass gait" models in the sagittal plane (Gregg et al. 2014; Gregg et al. 2012), to full body 3-D musculoskeletal models (J. Park et al. 2022). Virtual humanoid agents can then provide the necessary dynamics to the simulated P&O devices, for which novel control policies can be identified through the simulated coadaptive experience. This approach is finally seeing validation when transferring policies to real world environments, users and devices; primarily demonstrated for exoskeleton applications (J. I. Han et al. 2022; Luo et al. 2024). The most common approach for determining the dynamics of the human body is through cost-function minimising techniques such as trajectory optimisation or collocation (De Groote et al. 2021). Using them the generated gait can satisfy requirements of reducing metabolic cost and muscle effort, or to track experimental recordings of motion. This is often performed jointly with the optimisation of the simulated robotic device's control, if present. A limitation of performing optimisation directly on the dynamics is that it obfuscates the control policy of the device and its user. Due to their reliance on open-loop control, no additional observations can be made through them regarding the sensing and feedback mechanisms necessary to make control possible in stochastic environments (De Groote et al. 2021).

2.3 Hierarchical control schemes

The terminology of hierarchical lower-limb device control, as described for general motor control by Merel (2019) and for lower limb P&O by Tucker (2015) is used in this thesis (also see illustration on Figure 2.5):

- *HL control* refers to the control component responsible for discerning overall user intent, such as activity mode (e.g. walking, standing, sitting down), and parametrised description of the desired locomotion or movement (e.g. its direction and speed).
- *Mid-level control*, directed by the HL controller, defines the desired mechanical state of the device described in terms of angles, velocities, torques.
- *LL control* manages how the device reaches the selected mechanical state via actuation of the device's DoFs.

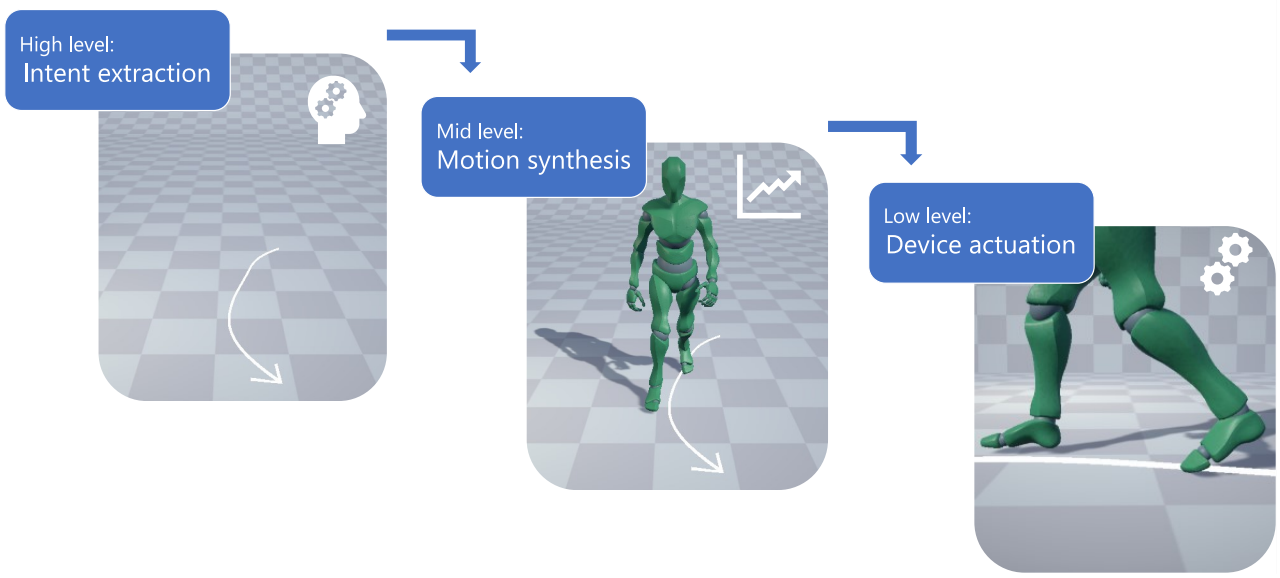


Figure 2.5: The locomotion task separated into three levels of abstraction. The HL intent of planning to travel in a certain direction is translated to body segment trajectories, then muscle or motor activation (Flash et al. 1985). This structure is comparable to the three level model of voluntary control proposed by Hollerbach et al. (1982).

By separating the motor control task into these levels, we can investigate them one by one. Certain assumptions must then be made regarding the levels not currently considered. Surrogate signal sources may, for example, play the role of HL controllers when testing lower level

processing (Bustamante et al. 2021; Gloumakov et al. 2022). However, hierarchical control is more than a convenient design strategy. There is strong evidence that an analogous division of responsibilities is present in biological motor control systems. Merel et al. (2019) provide an in-depth discussion on the emergence of hierarchy in control (anatomical and evolutionary), and the rationale behind the advantages they offer to drive the movement of living, robotic and computational systems. Summarising from Merel (2019), the benefits of hierarchical control strategies (in contrast to monolithic/flat approaches such as trajectory optimisation) include:

- *Efficient exploration.* The different layers of a control scheme act as signal processing stages, mapping sensory inputs and connections from other stages to an output. Not all information is relevant for achieving every element of a motion task, and what is a critical detail at one point may be considered as noise at another. In hierarchical settings, this mapping can be performed on an input and output space of reduced scale and dimensionality, making learning from trial and error more tractable.

- *Reduced effort.* Lower-level systems can simplify the task-segments delegated to them through assumptions using less costly operations, and act with partial autonomy, avoiding costly invocations of more complex control strategies. In the meanwhile, methods capable of planning and handling uncertainty, which require greater effort may be applied for higher-level tasks at a coarser resolution in time. They are utilised less frequently and with insignificant details abstracted away by lower-levels. The outcome of the higher-level control can then modulate the semi-autonomous operation downstream.

- *Skill reuse/generalisation.* Through the multiple layers of abstractions between the layers of hierarchy, a level of generalisation is inherently necessitated at each level. For example, HL locomotion goals do not necessarily have to be altered if the human picked up a backpack. The same course of action to plan a walking path can be used, relegating the responsibility to the lower levels to compensate. Furthermore, as the output of HL controllers may be interpreted by multiple subsystems, it must provide a representation that is useful for coordinating the global behaviour. This promotes the learning of motor synergies and multi-joint coordination.

2.4 Neural-driven control

Our motion intention is realised through contractions of our muscles, which are activated through the combination of signals descending from multiple neural sources from our central nervous system (CNS). It is possible to make measurements of the neural activity of the CNS through methods such as electroencephalography, or implantable cortical and spinal electrode interfaces. However, these either provide low-bandwidth information or require highly invasive and risky implantations (Farina et al. 2013b).

Instead, we can study how the CNS activity is mapped to the electrical behaviour of our muscles. By analysing how the electric potential fields evolve in space and time across the muscle, we can get an indirect sense of the underlying motor command (Merletti et al. 2019). This process can be quantified through electromyography (EMG) acquisition (De Luca 1997). Since myoelectric signals hold specific information about user intent prior to the onset of movement (Wentink et al. 2014), and are naturally involved with volitional human motor control, EMG is a common modality to drive P&O control strategies (Tucker et al. 2015; Fougner et al. 2012; Asif et al. 2021). The most successful bionic lower limb devices involve invasive neural interface surgery, providing both high quality intramuscular signals and sensory feedback (Kuiken et al. 2009; Urbanchek et al. 2016; Srinivasan et al. 2017). However, the requirement for expert surgery limits the accessibility of these approaches, and users are more willing to adopt easily reversible interventions (Kilgore et al. 2001). In contrast, control schemes using superficial sEMG signals are usually preferred for trialling prototype and proof-of-concept HMIs due to their simple, non-invasive and low-risk use.

Since locomotion involves the coordinated activity of the entire body and not just the legs, biosignals across the body are relevant for studying gait. Therefore, EMG from areas such as the trunk are also relevant for exoskeleton and lower limb prosthesis applications, the users of which may have leg muscle signals with reduced quality (Anam et al. 2012). There is evidence of correlation in trunk muscle activity and gait (Anders et al. 2007; Ceccato et al. 2009), but these signals are less frequently investigated for lower-limb P&O than leg muscles (Lopez-Delis et al. 2015).

A wide variety of approaches to decode motor goals from EMG have been explored for P&O control, ranging from simple threshold-based on-off control, activity classification with statistical features of the EMG, and continuous proportional control with pattern recognition systems (Oskoei et al. 2007b; Jimenez-Fabian et al. 2012; Tucker et al. 2015). Gait and locomotion in particular are highly autocorrelated processes, with muscle activity evolving gradually and cyclically, a key difference from upper limb contexts. This redundancy can be exploited with methods that rely on the temporal information of muscle signals to reduce the otherwise large variability. TCNs are promising models for continuously decoding time sequence signals like sEMG (Bai et al. 2018; Lea et al. 2016), which outperform more computationally demanding recurrent models in many applications. TCNs can learn spatiotemporal signal features and relationships adaptively, meaning they do not require a separate feature extraction stage, unlike support vector machines (SVMs) or motor unit spike train-based approaches (Huang et al. 2011; Shenoy et al. 2008; Kapelner et al. 2017). This architecture has been applied successfully for upper limb classification problems with high-density sEMG (Betthauser et al. 2019; Lea et al. 2017). However, the performance of TCNs on lower limb regression tasks with bipolar electrodes have not yet been assessed in detail. These devices are cheaper and more convenient to fit than their high-density counterparts.

2.5 Simulation tools

A wide range of software is available to simulate the locomotion of virtual humanoids (Figure 2.6). Physics engines allow iteratively simulating physics interactions, but cannot be used on their own to explore virtual wearable robotics. To be able to systematically create and develop virtual scenes of P&O and their users, an integrated development environment is needed (C. K. Liu et al. 2021). Unity was chosen for this purpose for the following reasons (Unity Technologies 2021):

- Large amounts of official and unofficial documentation and guides to familiarise with the framework and implement custom functionality.

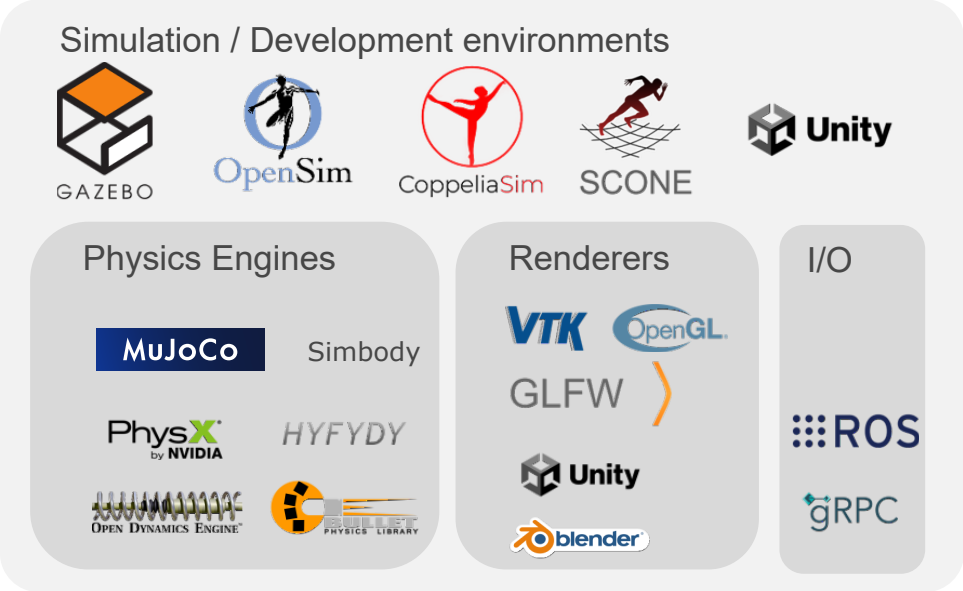


Figure 2.6: Relevant software tools for performing biomechanical simulations, and different roles they play. *Phyiscs engines* calculate and integrate multi-body dynamics over time, enforce constraints and contacts (Erez et al. 2015; Sherman et al. 2011; Geijtenbeek 2021). *Input/Output* protocols allow communication between external control systems and the simulated actuators. *Rendering engines* visualise and display the resulting movement. *Integrated Development Environments* provide tools to edit models and environments, script behaviours and controllers, and run simulations and learning (Geijtenbeek 2019; Unity Technologies 2021; Delp et al. 2007; Rohmer et al. 2013; Koenig et al. 2004).

- High performance simulation, with industry standard tools for running and retargeting humanoid animations.
- Easily extendable and modifiable behaviour via C# scripts in an object oriented way, a programming methodology recommended for biomechanics research for its modularity, extensibility and ease of maintenance (Hicks et al. 2015).
- Native support for importing and training machine learning models, including RL agents, through the ML-Agents package (Juliani et al. 2018).
- Official support and compatibility with third-party software from sEMG to motion capture applications via dedicated Application Programming Interfaces.

Unity on its own, however, may not be appropriate for many biomechanical simulations of virtual wearable robotics and their users:

- It lacks the options for tendon and muscle actuation, relevant to capture characteristics

of motor control relevant for many use cases (Hicks et al. 2015; Kidziński et al. 2020; Caggiano et al. 2022).

- Its articulated generalised coordinate representation of systems (Featherstone 2010), is still under development and less feature complete when compared to Unity's default Cartesian coordinate based one. However, generalised coordinates are more suitable to model biomechanics (Hicks et al. 2015).

- Access to the internal physics pipeline (powered by the PhysX engine) is limited, forcing users to hand calculate important measures such as the inertia matrix, necessary, e.g., for stable proportional-derivative (SPD) control (Tan et al. 2011).

- In PhysX, contact dynamics have low versatility and struggle with efficient and stable modelling of soft constraints. (Chappell et al. 2022).

For these reasons, to model dynamic activities such a gait, Unity must be combined with a physics toolkit other than the default PhysX that comes built-in. One such solution is to use the recently released Unity plug-in for the open-source physics engine MuJoCo, which addresses all of these issues. MuJoCo has precedence for use in biomechanics modelling, and has been shown to replicate results from the gold standard biomechanics simulation toolkit OpenSim (Caggiano et al. 2022; Hicks et al. 2015).

2.5.1 The Modular Agents package

The features available for using MuJoCo in Unity at the start of the project were not sufficient for the purposes of learning motor control tasks with RL, one of the key methods we employed in later chapters. For these reasons, we extended the functionality of the MuJoCo plug-in to support our research goals, and developed an open-source toolkit for training locomotion agents with RL. We documented and released this toolkit as the "Modular Agents" package with tutorials and examples (Hodossy 2024; Hodossy et al. 2023b). These tools formed the basis of the work presented in Chapter 5.

2.6 Virtual character animation

The study of biomechanics, gait analysis in particular, is closely related to the industry of animating characters for visual effects and interactive media. Once a user-determined pattern of motion is generated, the main difference between the two disciplines is whether those motions are rendered on a screen or guide a P&O device. For example, each actuated joint's angle-trajectory can be extracted from an animation, and sent as targets for LL proportional-derivative (Tan et al. 2011) or impedance control (Hogan 1984) of a prosthesis or exoskeleton.

To this day many techniques and tools, like motion capture, are shared between these two fields. Motion synthesis methods originally described with biomechanics and P&O devices in mind, such as locomotion finite-state machines (FSMs), were also adapted to computer graphics (Zeltzer 1982; Tomovic et al. 1966). Applied for locomotion, FSM either refers to an event based decomposition of a single gait cycle into periods of different joint behaviour, or extended to a system for recognising and transitioning between activity modes. These definitions are shared between character animation and P&O device control, and many exoskeletons and prostheses implement a FSM approach to control (Tucker et al. 2015; Unity Technologies 2022; Varol et al. 2009; Beil et al. 2015; Gillies 2009). Computer graphics has seen many refinements and successful variations of these techniques, and whether these could be adopted for assistive robotics should be investigated.

Motion graphs allow for a data-driven approach of constructing detailed locomotion FSMs. They describe a systematic process to determine points in time to allow transitions from one animation to another in a motion capture library, producing a sequence of movements according to desired parameters (Kovar et al. 2008). Based on how a motion graph is configured, it is possible to get more or less transition points. Following this line of logic, we may ask: What if we do not restrict transitioning between animation states at all?

2.6.1 Motion matching

Motion matching is a state-of-the-art way of animating characters in video game production (Buttner 2015; Clavet 2016). It is relatively easy to implement, and yields a robust and responsive controller that is flexible to different gait types and speeds. It requires approximately 5-20 minutes of unstructured motion capture data, which portray a comprehensive collection of expected locomotion types. These include walking at different speeds, veering left or right, pivoting, coming to a stop and starting walking. During preprocessing, each frame is labelled with a vector of features summarising the current pose and direction of travel. During operation, the controller repeatedly searches its library for the best match (in a nearest neighbour sense) to the character's current position and the user's HL intent of locomotion (which is left as a free variable). This is performed in a feature space of kinematics. If that frame is different than the current upcoming one, the controller jumps to it, and plays the animation from that point. Discontinuities are removed with interpolation or inertialisation blending (Bollo 2017; Bergamin et al. 2019). Inverse kinematics based post-processing enforces constraints, such as non-sliding feet. If optimised, motion matching only uses approximately 400 CPU cycles per frame which is feasible for even modest computing platforms (Buttner 2015). To avoid confusion, it should be noted that the term "motion matching" also identifies different concepts in literature, such as participants tracking trajectories (Pham et al. 2015; Verweij et al. 2017).

The motion matching algorithm can be classified as a kinematic controller, since it uses no information about the kinetic properties of the character (e.g. segment masses, joint torques), and does not interact with the environment in a dynamic way. Kinematic controllers are great at producing realistic looking and inconspicuous movement that is highly responsive to user intent. Other modern motion-capture based kinematic controllers include phase-functioned neural networks (Holden et al. 2017) and neural state machines (Starke et al. 2019). They excel at memory efficiency and object interactions respectively, however they both require a lot of manual tuning and involve advanced ML systems.

Both Chapters 5 and 6 use a custom motion matching implementation extensively to generate a kinematic baseline for a humanoid locomotion agent. The HL locomotion intent space in

Chapter 4 was likewise heavily inspired by the trajectory features used to make motion matching controllable. Additional information regarding motion matching and our implementation of it is available in Appendix A.

2.6.2 Closed-loop control

These motion synthesizers are promising starting points for virtual humanoid control policies, but they have significant limitations if used alone. Since they are kinematic as opposed to kinetic, even with appropriate LL control, simply tracking the estimated motion is prone to become unstable and cause the humanoid to lose balance in a dynamic simulation. This is because they have no way to correct for unexpected instability caused by the user's movement, changes in loading, or perturbations like being pushed or tripped (they are essentially open-loop). Another layer of modulation is required on top of the kinematic controller. Conceptually this is similar to how open-loop control originating from interneuronal circuits in the spine (called central pattern generators) are modulated in human motor control (Ivanenko et al. 2013; Lindén et al. 2022). This locomotion tuning is understood to be done via two pathways: with HL volitional control from the brain, and LL reflexive feedback of afferent mechanoreceptors and nerve endings (Dietz 2002). In a shared autonomy scheme as explained by Millan (2010), a HL controller could direct the ensuing locomotion in a coarse manner, while a RL-trained agent could alleviate the LL burden of maintaining balance.

RL is an online ML paradigm that improves a controller acting in an environment by experiencing the impact those actions had on the controller's measure of success. This measure can prioritise things such as maintaining balance or a natural look during locomotion. Recently, RL has been shown to successfully achieve these goals when modulating kinematic controllers in a physics simulation of animated characters (Bergamin et al. 2019; S. Park et al. 2019; Luo et al. 2020; Won et al. 2020). This method is also capable of retraining the controller for characters of different dimensions and DoFs while maintaining performance (Peng et al. 2018; Merel et al. 2017; Z. Wang et al. 2017; Lee et al. 2021), meaning it could potentially adapt the controller for users with different height, weight and neurophysiological condition, a key practical consid-

eration for HMI control (Graimann et al. 2013; Anam et al. 2012). Further more, RL could automatically evolve with the user, maintaining performance with the changing state and needs of the user, a possible solution to this need described by (Tucker et al. 2015). However, many other strategies for physically simulated interactive character animation exist that could fill this role (Geijtenbeek et al. 2012). Alternative methods to determine the human control policies include trajectory optimisation, collocation or neuromechanical control models (Dembia et al. 2020; C.-Y. E. Wang et al. 2005; A. R. Wu et al. 2017).

Chapter 3

Myoelectric Knee Motion Estimation*

To control wearable robotic systems, it is critical to obtain a prediction of the user's motion intent with high accuracy. Surface electromyography (sEMG) recordings have often been used as inputs for these devices, however bipolar sEMG electrodes are highly sensitive to their location. Positional shifts of electrodes after training gait prediction models can therefore result in severe performance degradation.

This study uses high-density sEMG (HD-sEMG) electrodes to simulate various bipolar electrode signals from four leg muscles during steady-state walking. The bipolar signals were ranked based on the consistency of the corresponding sEMG envelope's activity and timing across gait cycles.

The locations were then compared by evaluating the performance of an offline Temporal Convolutional Network (TCN) that mapped sEMG signals to knee angles. The results showed that electrode locations with consistent sEMG envelopes resulted in greater prediction accuracy compared to hand-aligned placements ($p<0.01$). However, performance gains through this process were limited, and did not resolve the position shift issue.

Instead of training a model for a single location, we showed that randomly sampling bipolar combinations across the HD-sEMG grid during training mitigated this effect. Models trained

*This chapter is a reproduction of the work published as "Balint K Hodossy, Annika S Guez, et al. (2024). *Leveraging high-density EMG to investigate bipolar electrode placement for gait prediction model.* In: IEEE Transactions on Human-Machine System", which is a shared first authorship publication with Annika S Guez.

with this method generalised over all positions, and achieved 70% less prediction error than location specific models over the entire area of the grid. Therefore, the use of HD-sEMG grids to build training datasets could enable the development of models robust to spatial variations, and reduce the impact of muscle-specific electrode placement on accuracy.

The methods described in this chapter are broadly relevant for myoelectric intent estimation using a pattern recognition or ML approach. However, we investigate a LL intent representation in this study in the form of estimating desired joint kinematics directly. This output is more closely related to exoskeleton applications, where less autonomy is expected from the device than with robotic prosthetics. For this reason, in this chapter we discuss the relevancy of the methods in the context of exoskeletons. In the following chapter, we will adapt these methods for HL intent estimation, targeting a lower limb prosthesis use case.

3.1 Introduction

Exoskeletons and other wearable robotic devices are designed to guide and support users in a variety of contexts, from rehabilitation to movement augmentation. In most cases, a user-driven human-machine control interface is necessary to ensure active participation from the user(Knaepen et al. 2014). By estimating the active contribution of the subject to the desired motion, an exoskeleton can provide the required level of assistance synchronously with the user's movement - i.e. assist-as-needed (AAN) control (Staudenmann et al. 2010; Caulcrick et al. 2021). The performance of the system therefore depends on the reliability of the signals used as inputs to the model that estimates the intended motion.

SEMG is a commonly used non-invasive sensing modality for motion intent estimation due to inherent association between muscle signals and movement (Oskoei et al. 2007b). A subject's sEMG signals can be used as indicators of voluntary effort (Sheng et al. 2019) and as inputs for gait prediction models (Jimenez-Fabian et al. 2012; Steinhardt et al. 2018), which form key elements in the control schemes of wearable robotics (Tucker et al. 2015).

Despite the extensive use of sEMG in human-machine interfacing and clinical practice (Rodriguez-

Tapia et al. 2020), there is a limited number of studies that assess the impact of electrode shift for lower limb assistive devices (Simao et al. 2019), especially when employing deep learning methods (Xiong et al. 2021). Whilst some studies have investigated the effect of electrode placement on signal characteristics for lower-limbs (Campanini et al. 2006), this has not been done across different signal acquisition and gait conditions, such as wearing an orthosis. Furthermore, changes in electrode location can cause significant distortions in the EMG signal features required for pattern recognition (Subbaswamy et al. 2020), making previously learnt features inapplicable unless the model is retrained or calibrated for the new location (Steinhardt et al. 2018).

Based on the signal's sensitivity to electrode placement, there are two main approaches that have the potential to mitigate this effect and help the model retain high gait prediction accuracy:

1. Determine the electrode location on each muscle that will provide the highest sEMG signal quality;

2. Develop a model that is robust across electrode locations.

While available standards and tutorials (Hermens et al. 1999; Merletti et al. 2019) provide some indication on sEMG electrode placement, these are unspecific and often do not account for the signal variability during dynamic movements (Campanini et al. 2006; Rainoldi et al. 2004; Sacco et al. 2009), or the changes introduced by donning an orthosis with limited DoFs. Furthermore, there is no guarantee that these palpation- and eyesight-based guidelines provide the location of the optimal signal for pattern recognition-based control schemes (Young et al. 2011).

Previous studies have attempted to mitigate the impact of electrode shifts by including additional preemptive steps that recognise and compensate for them. These are either signal registration methods transforming signals to a predetermined expected distribution (X. Zhang et al. 2020), or conversely apply the appropriate prediction model from an ensemble (J. He et al. 2020; J. He et al. 2019). Alternatively, models can be trained on a distribution of locations

simultaneously, either recorded experimentally (Young et al. 2012) or simulated through data augmentation (L. Wu et al. 2020).

In this study, we investigated both methods (selection of electrode location and generalisation of the model by extended training), and compared their usability in practice. For this purpose, we used TCNs as an example of a data-driven model suited to regress time-series data (Lea et al. 2017). This network architecture benefits from more straightforward implementation and much faster training times than its recurrent counterparts, while maintaining comparable performance (Bai et al. 2018). Previously, TCNs have been successfully applied to lower limb exoskeleton control with inertial sensor data (Fang et al. 2020; Kang et al. 2021), and recently for bipolar sEMG based control (J. Liu et al. 2022). Here, single channel bipolar sEMG signals were used to estimate the knee angle in a one-step-ahead offline prediction scheme. Muscle signals were sampled from grids placed on knee and ankle flexor and extensor muscles.

To compare a wide range of bipolar electrode placements, HD-EMG grids were used to simultaneously acquire data from a larger surface area of the skin, which allowed sampling of bipolar signals from various positions and orientations (Huang et al. 2008). This diverse dataset was then used to produce more reliable convolutional models for exoskeleton applications under dynamic conditions, and was also tested when a one DoF passive orthosis was donned. The orthosis condition was not investigated in an attempt to model a patient user. Instead, it was used to examine whether signal artefacts from the DoF restriction or the orthosis' contact with grid impact the performance of our models. While the device used here was passive, noise from these sources is expected to be similar in an actuated rigid exoskeleton as well. In addition to our use of location specific models to evaluate signal quality, we showed that TCNs are capable of learning features suitable for robust prediction across the area of the grid. This is achieved by randomly sampling bipolar electrode location during training. This technique is a promising method for mitigating the effect of electrode shift on deep learning models. Our findings highlight the potential of using the spatial information from HD-EMG data to improve model robustness for bipolar signal driven control applications.

3.2 Materials and methods

3.2.1 Subjects

5 subjects (3 male, 2 female, mean age = 25 ± 4 years) were chosen for this study, with no history of neurological or physiological conditions that could impact their natural gait. Informed consent was obtained for all subjects. The experiments for data acquisition were performed in compliance to ethical documentations approved by the Imperial College Research Ethics Committee (ICREC reference: 21IC7204).

3.2.2 Experimental setup

High-density surface electromyography

For all subjects, HD-EMG signals were recorded independently from the Rectus Femoris (RF), the Biceps Femoris (BF), the Tibialis Anterior (TA) and the Soleus (SO) using 13×5 electrode grids in monopolar derivation (8-mm inter-electrode spacing; model number GR08MM1305). When recording, the relevant grid was connected to a Sessantaquattro acquisition system, and the ground electrode was placed on the same leg's lateral malleolus, using a designated damp ankle strap. All equipment used for the acquisition of the HD-EMG data came from OT Bioelettronica, Italy (OT Bioelettronica 2022).

The positioning of the centres of the electrode grids followed the SENIAM placement guides (Hermens et al. 1999). Throughout the paper, any mention of the "middle" or "standard" sampling will refer to the bipolar placement following SENIAM standards.



Figure 3.1: Experimental setup, showing the HD-EMG grid placements (white rectangles) *(a)*: the RF and TA, *(b)*: the BF and SO, the motion capture marker positions (white circles). *(c)*: Set-up with the passive orthosis strapped to the subject's leg.

A double-sided adhesive foam (model number FOA08MM1305) was placed on the HD-EMG grid, and CC1 AC conductive cream was applied to the grid to ensure satisfactory contact between the electrodes and the skin. Before electrode placement, the area of skin was shaved as needed, mildly abraded with abrasive paste, cleansed with a non-irritating alcoholic solution, and dried.

Motion Capture

The knee joint kinematics were acquired using a Vicon Motion Capture system equipped with ten 120 Hz Vicon Vero v2.2 cameras placed around the treadmill, mounted on the ceiling of the recording space. Trajectory data were labelled and gap-filled in Vicon Nexus.

Rigid body segments were reconstructed using the following set of markers (see Figure 3.1):

1. **Pelvis:** 2 markers placed on the superficial aspects of the left and right anterior superior iliac spines (ASISs), and 1 marker on the sacrum.

2. **Thigh:** 2 markers placed on the femoral epicondyles, and 1 marker placed on the lateral side of the thigh.

3. **Shank:** 2 markers placed on the lateral and medial malleoli, and 1 marker placed on the lateral side of the shank.

Knee flexion was defined according to (Grood et al. 1983) using all above-mentioned segments. The hip joint centre was approximated as a fixed position in the pelvis' coordinate frame, scaled by the ASIS breadth (Vaughan et al. 1992).

When wearing the passive orthosis (see Figure 3.1), the lateral thigh and shank markers were replaced onto the orthotic frame after all other components had been fitted. Both the medial and lateral epicondyle markers were placed along the axis of rotation of the orthosis. The HD-EMG grids were not moved or replaced when donning the knee brace.

System synchronisation

Shared reference pulses, sent from a micro-controller at the start and end of a trial, were recorded by both data acquisition systems. These were then used in the preprocessing phase to methodically align the sEMG and motion capture signals in time based on the timing of the rising edges of the pulses.

3.2.3 Experimental procedure

Each subject was asked to walk on a treadmill at a steady-state speed of 0.8 m/s, selected to be comfortable for all subjects. Two 3-minute trials were recorded for each muscle separately, with rest breaks in-between to ensure muscle fatigue would not impact the recordings. This led to a total of 40 recorded steady-state trials (2 trials $\times$ 4 muscles $\times$ 5 subjects), with approximately 350-400 strides per trial. For the second half of the experimental session, the user was fitted

with a dual-hinged, single DoF passive knee support brace. A single 2-minute trial was recorded for each subject and muscle, leading to a total of 20 trials with the orthosis.

3.2.4 Signal preprocessing

Bipolar signals were extracted from the 64 channel HD-EMG grid by taking the monopolar voltage difference between an electrode and its neighbouring electrodes within a 3-electrode radius and a maximum angle from the longitudinal axis of 45° (see Figure 3.2). This process was repeated for every electrode, and repeated pairings of electrodes were ignored. A total of 360 unique bipolar combinations that would satisfy these requirements for sensor placement contained within the grid were identified.

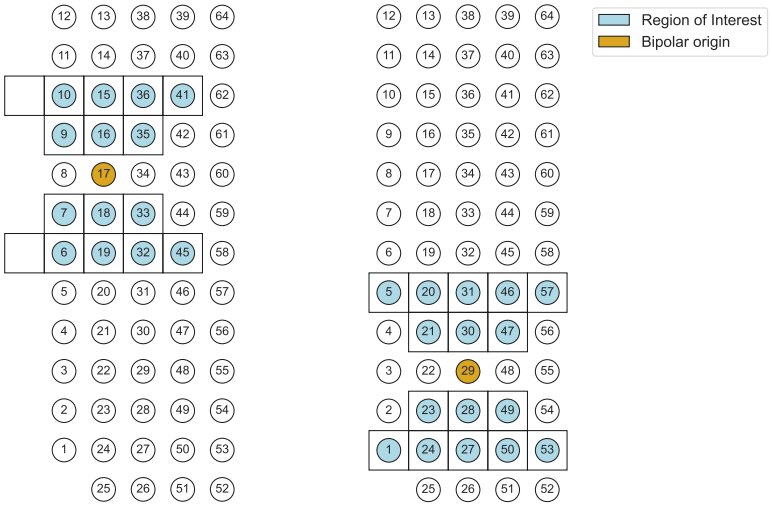


Figure 3.2: The range of neighbouring electrodes considered for each electrode when determining the bipolar combinations. This schematic shows the combinations extracted for two example electrodes. The "origin" electrode (orange) was eligible for pairing with all the electrodes (blue) from its region of interest. This resulted in up to 16 possible bipolar combinations for each electrode.

The motion capture data were up-sampled to the EMG sampling frequency (2000 Hz) using linear interpolation, and the data were cropped to extract the steady-state walking sequence.

The consistency metric analysis was performed based on the sEMG envelope. To obtain it, the signals were digitally band-pass filtered (4th order Butterworth, 20-400 Hz), rectified, and subjected to a 10Hz low-pass filter. Gait cycles (GCs) were segmented using the peak angles of the knee, and temporally normalised so that one GC corresponded to 100% of the time axis.

During pre-processing stages, 2 HD-EMG files were found corrupted and were excluded from further analyses, reducing the dataset from 40 available trials to 38. The remaining data still covered all subjects and muscles investigated due to repeated trials.

3.2.5 Bipolar sampling selection

Quantifying a single bipolar channel's quality by training a location specific TCN model takes several hours. This gives motivation for identifying quick-to-evaluate metrics that indicate suitability for regression. We selected channels to evaluate with TCN performance based on metrics of sEMG envelope consistency. While most studies focus on the location with sEMG maximum amplitude (Lapatki et al. 2010; Burden 2010), this study investigated the consistency of the signal across steady-state gait cycles as a measure of quality (Campanini et al. 2006). Considering how many myoelectric control models are based on pattern recognition mechanisms, it was hypothesised that a consistent signal would lead to more accurate and stable predictions.

The consistency across gait cycles was calculated from the following three measures of signal's envelope:

1. Maximum peak amplitude
2. Integrated area
3. The gait phase of the maximum peak's location

For the first two of these metrics, the inverse of the coefficient of variation (CoV) was used to quantify the consistency of muscle patterns across the gait cycle (a lower CoV implied a higher consistency), expressed as the standard deviation (σ) normalised by the cycle's mean value (μ) and given in percentage form:

$$CoV(\%) = \frac{\sigma}{\mu} \times 100 \tag{3.1}$$

The location of the maximum sEMG peak within the gait cycle was expressed using phasor representation calculated by:

$$\Phi_t = \left(e^{j2\pi/T}\right)^t \qquad t \in [0, T-1] \cap \mathbb{Z} \tag{3.2}$$

where t is the temporal index of the peak within a gait cycle of length T. This allows description of normalised progress in the gait cycle without discontinuity between the start and end. The average vector was calculated across all cycles and the variance metric was calculated using the absolute value of the angle between the peak location vector and the normalised average vector. An illustration of the measures used to quantify consistency is shown on Figure 3.3.

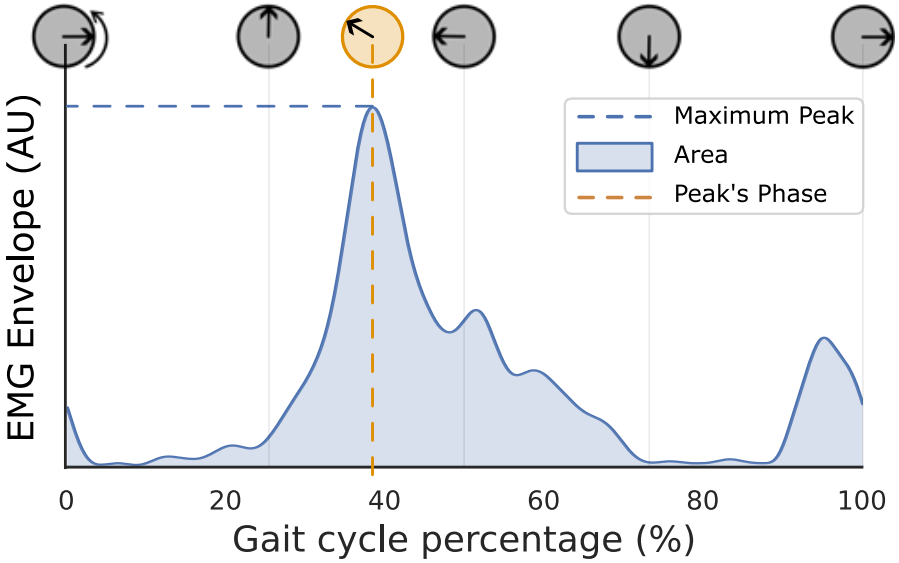


Figure 3.3: Illustration of the three envelope measures within a gait cycle that were used to quantify signal consistency for the RF. Channels are considered more consistent if these measures don't change from cycle to cycle. Phasor representation of the continuous progress along gait cycle is shown above the curve.

The electrode channels were ordered with ascending variation of the three metrics to facilitate the selection of those with the least variation (i.e. the most consistent ones), or conversely the ones with the highest variability (i.e. least consistent). Evaluating all three metrics for our entire dataset took less than 2 minutes. However, the obtained rankings were different across all three consistency metrics.

To resolve this issue, we employed a method that combined the channel selections of all three metrics, referred to as the *Agreement* approach. This procedure consisted of selecting the first

bipolar electrode combination that occurred across all metrics when progressing through their respective ordered consistency rankings. This ensured all three metrics were considered when selecting the electrode pair. On average, a match was found in the first 10% of the ranked combinations.

To determine whether to rely on one of the three rankings or the Agreement approach, preliminary tests were conducted where one TCNs model was trained and evaluated with a 50-50 train-validation split for 3 subjects, across every muscle with each sampling method. From the results shown in Figure 3.4, the Agreement approach was subsequently adopted for all following evaluations as its chosen sampling led to the highest model performance.

A common bipolar electrode size (64 mm^2) was simulated virtually by spatially filtering the HD-EMG grid with a 2x2 averaging window before sampling bipolar signals, reducing the grid size from 13×5 electrodes to 12×4. No significant difference in model output was observed, so remaining experiments were conducted with the HD-EMG signals to allow the inclusion of samplings with smaller inter-electrode distance (IED) as well (8 mm).

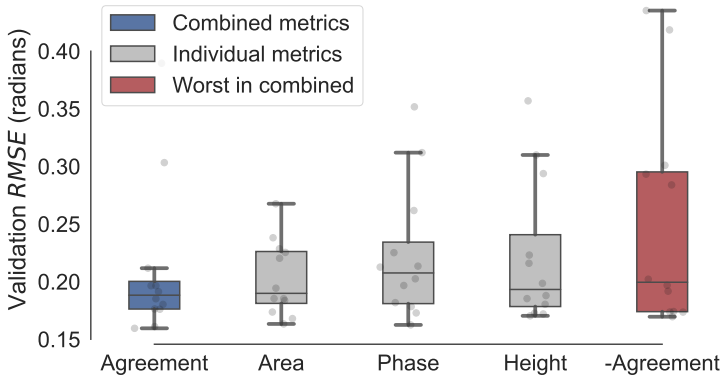


Figure 3.4: Metric selection for further analysis based on model performance on reduced data set. Angle prediction error on validation data is compared for one model trained for each muscle of the first three subjects. The labels for each plot correspond to the metric used to determine the channel locations for that group's TCNs. The "-Agreement" label specifies the least consistent electrode location based on the "Agreement" method.

3.2.6 Knee angle prediction model

The TCNs investigated were composed of two convolutional layers with pooling and dilation layers between them. This was followed by three dense layers with dropout in-between (see Figure 3.5). This architecture is similar to the one in (Zanghieri et al. 2019).

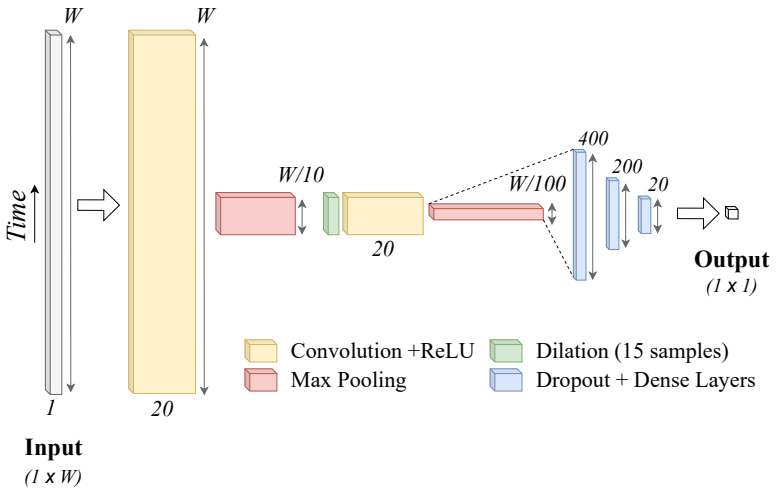


Figure 3.5: Schematic of the 1-dimensional TCN architecture, mapping the windowed sEMG signal $(1 \times W)$ to the user's knee angle (1×1). W stands for the width of the input window, either 1000 or 500 samples.

No envelope or other feature extraction was performed on the muscle signals, and no temporal normalisation was applied. The sEMG input was normalised using the mean and standard deviation values from the entire training set of the given channel. Then, it was split into windows of equal sizes, partly overlapping at a 40 sample stride. Two window sizes were investigated, 500 and 1000 samples. The shorter (250 ms) window is similar to window lengths from TCNs used for EMG processing in upper limb contexts (Betthauser et al. 2019). However, when compared to the upper limb tasks, gait is a more auto-correlated process. Therefore, a longer window (500 ms) was also explored. The windows were paired with the flexion signal in a one-step-ahead regression scheme. An early termination condition was used to stop training after 100 repeats of the training data, or after performance on validation data has not improved or has worsened for the most recent 6 repeats.

Connection between consistency metrics and model performance

root mean square (RMS) error between the knee angle predicted by the model and the angle from the motion capture was calculated. The error of trained models on validation data was compared across different spatial sampling conditions and window sizes (250 ms and 500 ms). A 5-fold split was performed on each training set for cross-validation (Dietterich 1998), performed for each muscle and subject. The splitting of the recordings remained consistent across conditions, and no window was allowed in a fold of cross-validation that had any overlap with other folds. With 5 subjects, 4 muscles, and 5 folds, 100 models were trained and evaluated for each condition on spatial sampling and window size, for a total of 600 models in this comparison. The following bipolar channel selection methods were compared:

- The subject and muscle specific channel with the best cross-metric consistency based on the Agreement approach, as outlined in Subsection 3.2.5.
- The bipolar combination aligned according to the SENIAM placement standard (between electrodes 31 and 33).
- The subject and muscle specific channel with the worst cross-metric consistency based on the Agreement approach, as outlined in Subsection 3.2.5.

Spatially robust bipolar feature learning

Models trained on the electrode combination in the middle (electrodes 31 and 33) were evaluated on the signals of all other valid electrode combinations. Their performance was compared with non-placement specific models, which regressed signals from any valid combination in the grid directly. During the training of these TCNs, the bipolar channel that determines the signal was sampled uniformly from the set of all valid bipolar channels for each 500 ms input window. Therefore, even in this case, only a single channel of sEMG input was used at a time. To achieve consistent performance, the model needs to extract features that were suitable for generalising

over the area covered by the grid, and can not optimise for any single placement. As such, this method can be considered as a technique for implicit regularisation of sEMG models. Lastly, this analysis of model robustness to placement shifts was repeated while wearing a one DoF knee orthosis, to simulate the impact an exoskeleton frame may have on the signal distribution due to restrictions to the range of motion or motion artefacts (Sylos-Labini et al. 2014; Hidler et al. 2005). A causal band-pass filter of 10-400 Hz was used, a common strategy to mitigate the impact of these types of interference (Mello et al. 2007). This pre-processing was also applied for the models without the orthosis in this section, for a fairer comparison.

Performance was evaluated similarly to the previous section. For each training condition (training with a single location, or randomly sample during training) and orthosis condition (donned or doffed) a TCN is trained for each of the 5 subjects, 4 muscles and 5 folds (100 for each condition).

The discussed methods of bipolar signal sampling from grids could potentially be used in clinical applications. Robust features may be learnt from HD-EMG grid data, and then deployed with more convenient to use traditional bipolar sEMG signal acquisition devices. A preliminary data collection using a Delsys Trigno electrode system with two of the subjects was used to trial this approach. TCNs trained with bipolar signals from the HD-EMG grid were applied to the acquired Delsys data (see Figure 3.13).

Deep learning was performed using the Keras framework for TensorFlow (Chollet et al. 2015), on a single computer with a Intel i7 CPU and NVIDIA GeForce 2070 GPU, taking less than 30 minutes to train a model.

3.3 Results

3.3.1 Connection between consistency metrics and model performance

The samplings chosen based on their score from the consistency metrics led to higher performance of the TCN, with a significantly lower validation root mean squared error (RMSE)

Table 3.1: Median validation RMS error and standard deviation (σ) for 100 models trained across subjects and muscles. Values in radian.

Channel Selection	0.25 s	0.5 s
High Consistency	0.067 ± 0.040	0.043 ± 0.022
Middle	0.084 ± 0.042	0.045 ± 0.020
Low Consistency	0.084 ± 0.059	0.045 ± 0.028

compared to the bipolar location based on SENIAM standards and inconsistent samplings (see Table 3.1 which summarises the median and σ statistics). This validates the metrics as indicators of signal quality, suggesting they could be used as a preliminary verification method for sEMG based models, whilst also demonstrating the impact of electrode placement for sEMG pattern recognition.

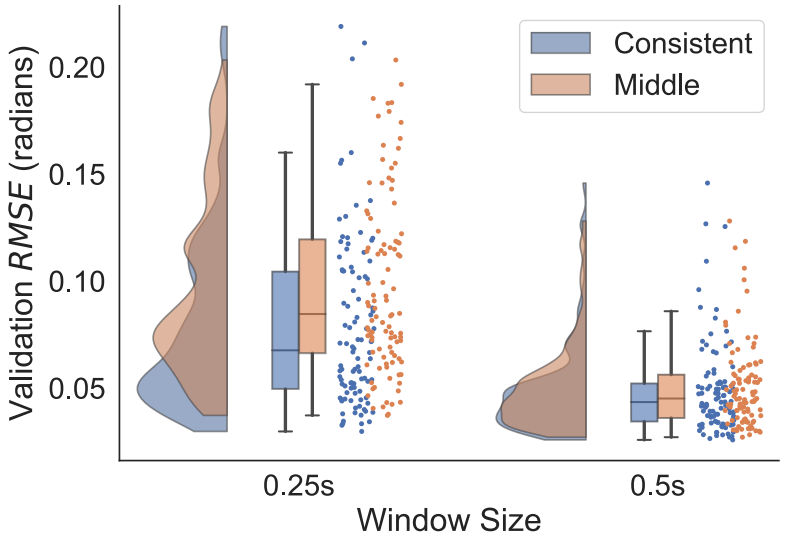


Figure 3.6: RMS error of angle predictions from TCNs trained on the combination with the best consistency score (blue) vs. the middle electrode combination (orange), for two different input window durations. The kernel density, the box plot and the raw data are shown for the different conditions.

Figure 3.6 shows the validation RMS flexion angle prediction error distributions of models of different electrode locations and input window sizes. The use of a larger window size can help mitigate the effect of electrode position (see Figure 3.6). This comes at additional computational

costs, which may be important considerations for real-time applications.

The prediction error was comparable to other gait regressors reported in the literature (Xiong et al. 2021). Due to the non-gaussian distribution, and to take inter-subject variation into account, paired non-parametric tests were used to assess differences and interactions between groups. A Friedman test showed significant differences in distributions, $p << 0.0001$, which was mostly due to differences across window sizes.

When repeating the test within the shorter window groups, differences remained significant ($p = 0.0263$), therefore further inter-group comparisons were made. The hypothesis that samplings with high consistency scores perform better than the sampling from the middle was tested with a one-tailed Wilcoxon signed-rank test. The 'consistent' sampling has a lower mean error of the two with $p = 0.0002$. Similarly, models using the 'consistent' sampling had a lower error than those trained with 'inconsistent' signals ($p = 0.010$).

When repeating the Friedman test within the longer window groups, no significant differences in RMS error was detected ($p = 0.432$).

3.3.2 Effect of the orthosis on the sEMG recorded signals

Even though the bipolar locations with high consistency scores performed better, this optimal placement selected during normal walking did not transfer when donning the orthosis, as the muscle activation patterns showed significant changes in amplitude and spatial distribution across these two conditions.

Figure 3.7 illustrates the change in behaviour of the 3 metrics when comparing the muscle activation patterns across the grid with and without the orthosis. Whilst some features seem to have some spatial similarities across conditions, the amount of muscle activity (shown by the envelope's integrated area) and the peak activation amplitude intensify when wearing the orthosis, and become significant in areas of the muscle that were initially inactive. In addition, the main activation peak takes place earlier in the gait cycle around the belly of the muscle,

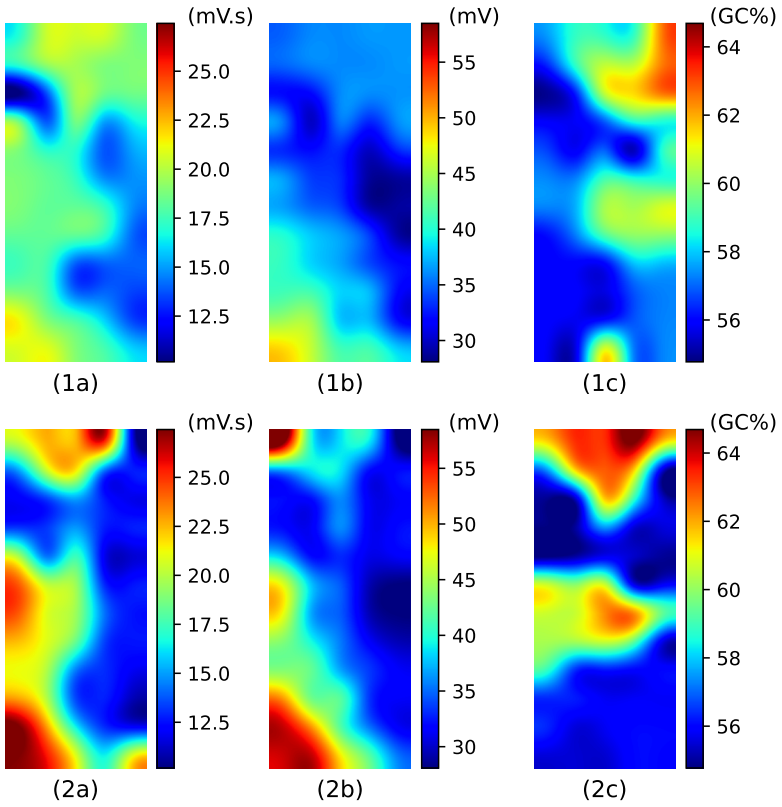


Figure 3.7: Visualisation of the average cross-cycle values of the sEMG envelopes' (a) integrated area, (b) maximum peak amplitude, and (c) maximum peak location in terms of GC percentage (with 0% corresponding to the heel strike) for a subject's SO (1) without and (2) with the orthosis. Each of the graphs show the spatial distribution of the metrics across the HD-EMG grid, and the values shown correspond to the vertical differential between 2 monopolar electrodes (therefore showing a 12×5 grid as opposed to the original 13×5 monopolar grid.

further distorting the pattern the model was trained on. This behaviour is in conflict with the hypothesis that there exists a subject-specific area on each muscle that will consistently provide reliable muscle activation patterns and generalise across conditions such as donning an orthosis.

3.3.3 Spatially robust regression with location sampling

Table 3.2 and Figure 3.10 evaluate the capability of the TCN to generalise over the available electrode combinations on the grid when compared to a single location training. There was a significant drop in validation RMSE when training the model with all the possible bipolar combinations contained within the HD-EMG grid. A Friedman test was used for the four conditions to evaluate the capacity for the model to generalise, and the improvements in performance across all conditions were significant with Wilcoxon signed-rank tests at $p < 0.01$. Figure 3.12 further illustrates the increased robustness of the gait prediction model, with a significantly lower σ across different electrode combinations.

Figure 3.9 shows the four conditions investigated in Figure 3.10, with model performance scores separated based on which muscle they were trained on. The relationships between the median RMS error from different muscle inputs is unchanged between conditions, with models using Soleus signals consistently performing best.

Table 3.2: Median validation RMS radian error and σ for 100 models trained across subjects and muscles, when evaluated on all valid combinations.

Training channels	No Orthosis	Orthosis
Full Grid	0.090 ± 0.030	0.122 ± 0.044
Middle Only	0.309 ± 0.096	0.338 ± 0.109

The average number of epochs until termination increased from 30 to 35. Figure 3.8 illustrates the combination specific performance when not using additional spatial samples compared when the regularisation is applied, visualised on the grid. This shows an increased robustness to longitudinal shifts compared to lateral ones, a pattern observed in all subjects and muscle groups in these results.

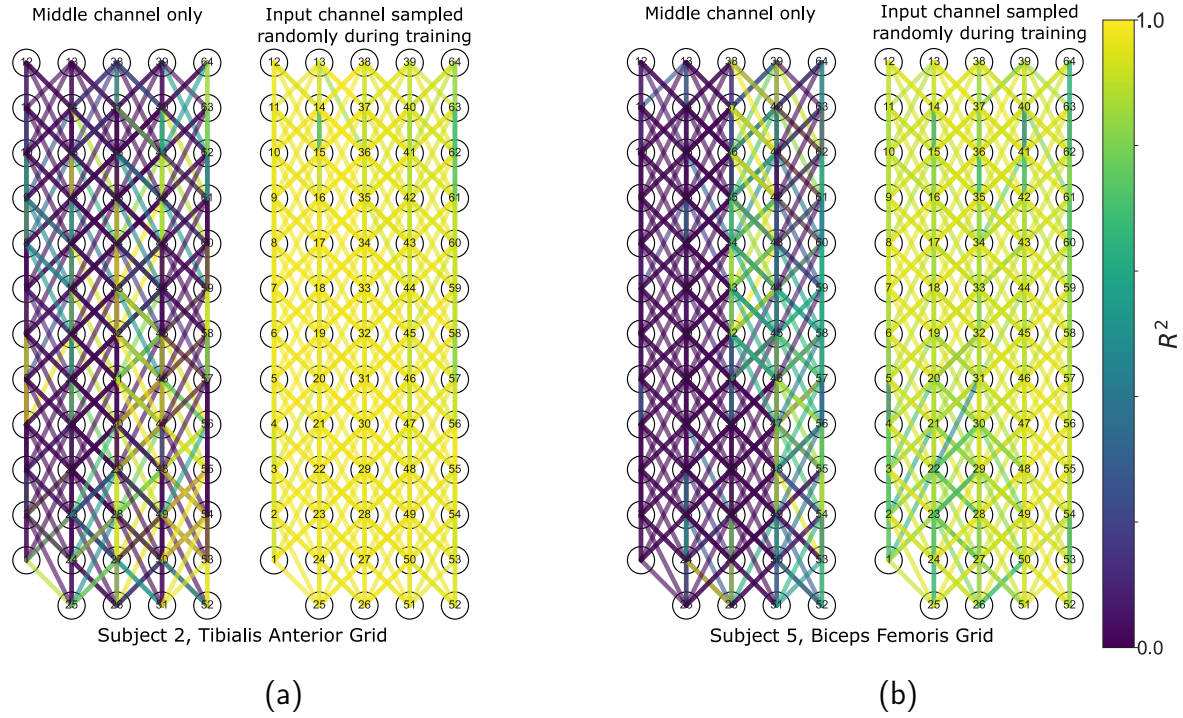


Figure 3.8: R^2 performance of models evaluated on every valid electrode combination on a grid, shown for TA in participant 2 (a), and BF in participant 5 (b). The left grid in each panel shows models that were only trained with the middle combination (31-33). The right grid in each panel shows performance if the input electrode combinations are sampled uniformly during training. R^2 values are visualised with a lower limit of 0. Transparent lines are used, as the region of interest for each electrode contains overlapping sections (See collinear combinations in Figure 3.2).

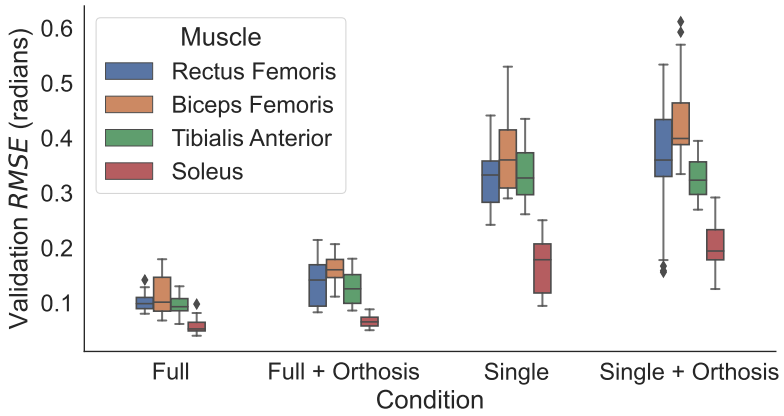


Figure 3.9: Predicted knee angle RMSE of each muscle group across conditions when trained on the entire grid (Full) versus only the SENIAM bipolar combination in the centre of the grid (Single). The condition of wearing the knee orthosis is also shown.

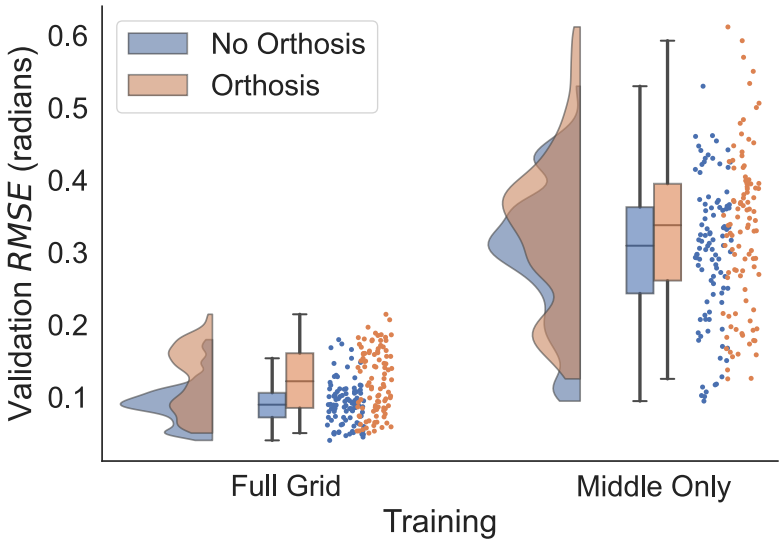


Figure 3.10: Average model performance across all valid bipolar combinations of the grid. Models were trained either with electrode combinations sampled uniformly from the grid, or solely with the middle combination (electrodes 31 and 33). The hue determines whether the model was trained and evaluated with data recorded while wearing a knee orthosis.

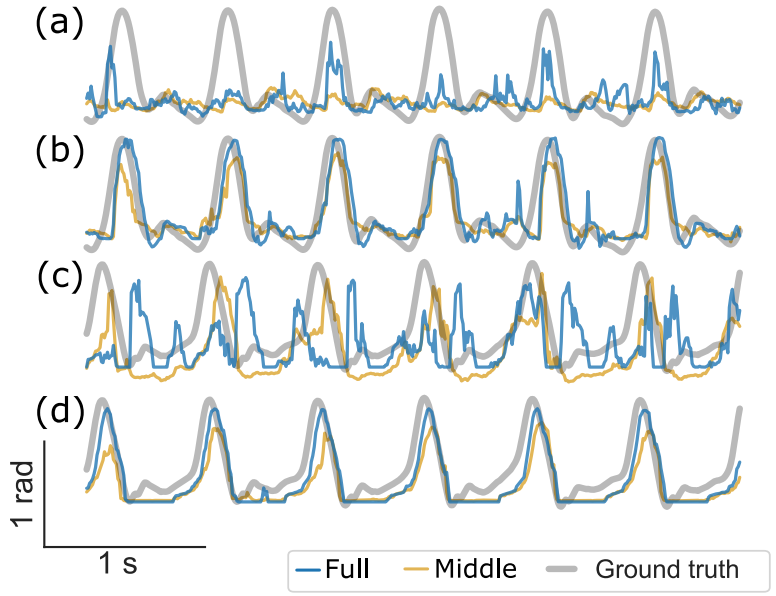


Figure 3.11: Knee flexion predictions with data recorded with Delsys bipolar electrode, with models trained on signal from the HD-EMG grid. Two subjects shown, subject 3 with the RF (a) and TA (b), and subject 1 with the BF (c) and SO (d). "Full" refers sampling the input channel uniformly from the entire grid, "Middle" refers to only training with the middle location. The angle from motion capture is overlaid.

3.4 Discussion

Our chosen consistency metrics were validated as indicators of signal quality, since electrode locations selected through them were associated with higher prediction accuracy, when compared to those hand-aligned based on guidelines. This suggests that consistency metrics could be used for high-level analysis of signal quality, especially in models with short input windows. However, no distinct pattern emerged to reliably identify better areas on the skin surface for sEMG acquisition from a given muscle. Not only was there no clear preferred electrode position on the muscle, but even small electrode shifts caused significant prediction accuracy degradation. Due to the severe impact of electrode shifts, and the inevitable placement discrepancies between sessions from electrode re-application (Murley et al. 2009), relying on optimal electrode placement is unrealistic. We must therefore prioritise the development of spatially robust models to improve gait prediction for exoskeleton applications. As the signal feature landscape is heavily affected by interference from the orthosis (Figure 3.7), including this condition should be a high priority for future studies investigating signal acquisition for neuro-orthoses. Potential causes for this difference may involve the additional stress, shear and relative movement at the orthosis-electrode grid interface. Alternate causes may originate from a change in the muscle activity due to the additional weight or DoF restriction. Lastly, a phase shift may be caused by the replacement of the motion capture markers, necessitated by the orthosis. However, since on Figure 3.7c the effect is non-uniform, this does not appear to be the dominant cause. Mitigating these effects through orthosis design and fitting could be investigated for more reliable signal acquisition.

The use of spatial correlations in sEMG signals can significantly improve robustness (Stango et al. 2015), however, the generalisation of spatially rich HD-EMG data for bipolar applications has not yet been investigated. Increasing the size of the electrode detection surface helps reducing model sensitivity to electrode shifts when perpendicular to the muscle fibres (Young et al. 2011). Hence, it can be hypothesised that increasing the data acquisition area from an electrode grid would improve model robustness in a similar way.

By collecting data from all 360 possible bipolar locations across the model, we widened the

training data distribution to include more possible scenarios that could be encountered due to electrode shifts. A wider range of spatial samplings includes channels where the effect of physiological features, such as muscle cross-talk and innervation zone (IZ) effects, manifest to different extents. As such, this could promote the recognition and reduction of these sources by the model, leading to more stable predictions. Another possibility is that modelling multiple orientations encourages the TCN to favour features that are on a longer time-scale, and are shared across various channels. This group training approach has been met with scepticism in the past, as it was previously associated with longer acquisition and signal processing times (Kyeong et al. 2020). However, these concerns were partly addressed in this study with the simultaneous collection of bipolar signals through the grid and the sampling of electrode location during training. A further limitation of the group training approach is that the additional training information can potentially crowd the feature space and decrease model performance for the original non-shifted location (J. He et al. 2020). Whilst it is true that spatial sampling methods decrease the average maximum performance, this regularisation method can still be a useful tool in training models, especially in scenarios where electrode shifts are inevitable. Within an exoskeleton context, due to its dynamic application to the lower-limbs, model robustness needs to be prioritised, and performance degradation can be moderated by other components in the system such as a higher-level controller or sensor fusion techniques (Tucker et al. 2015; Camargo et al. 2021). Furthermore, using grid recordings to train spatially generalising models could facilitate muscle selection studies by removing the impact of specific electrode placement on signal quality, and allowing for a more general comparison of signal reliability across muscles (see Figure 3.9).

Including bipolar samplings from across the grid does not substantially increase the training time of the predictive network, making this method convenient during both data acquisition and processing. Defining a non-uniform sampling distribution may allow the trade-off of generalisation and maximum performance to be fine-tuned. The observation that longitudinal shifts lead to less decline in performance is congruent with existing work (Young et al. 2012). Movements perpendicular to muscle fibres correspond to locations influenced by different neighbouring muscles, or to areas with different neuro-physiological characteristics, such as the anterior aspect of

the tibia. Figure 3.12a exhibits a heteroskedastic distribution throughout the gait cycle, which further illustrates the non-linearity of the signal variation across channels. However, our approach enabled the TCN to be robust to these transformations and phase shifts, as illustrated in Figure 3.12b.

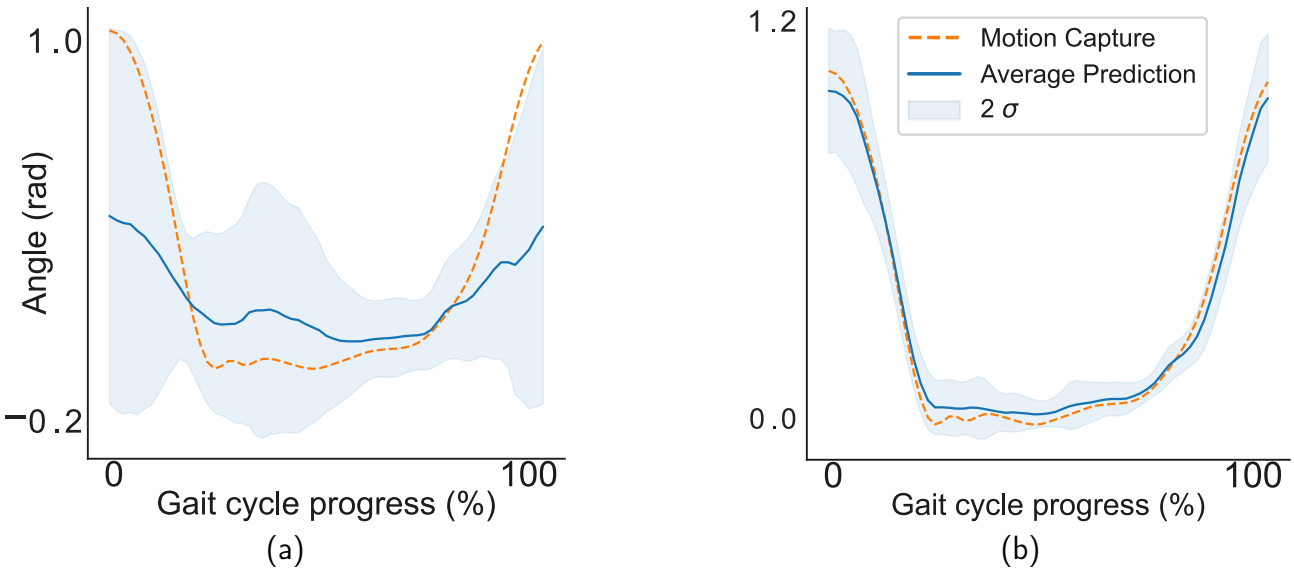


Figure 3.12: Standard deviation of the average gait cycle across electrode combinations for the BF models of participant 4. (a) shows the average prediction of a model trained with just the middle sampling, showing that the uncertainty is phase dependent. (b) shows the average prediction from a model trained with combinations across the grid.

Despite many differences, such as electrode size, original sampling frequency, and the use of electrode gel in the grid, there is a potential for transferability of information across acquisition systems. This would enable rapid and spatially rich data collection from grids utilised for robust bipolar applications. The shift in the profiles of the Delsys signals relative to the simulated bipolar signals from HD-EMG, and the failure of the model transfer in the case of the thigh muscles would suggest that this regularisation method would likely have to be combined with domain adaptation techniques to fine tune predictions to various expected conditions (Ameri et al. 2020). Furthermore, this method does not necessarily mitigate the requirement to re-calibrate networks due to long-term changes in the neural interface, or application to different subjects. The combined application of the different frameworks addressing these key issues should be investigated (Kim et al. 2019).

The approach of deploying pre-trained models may not be as effective in cases where the

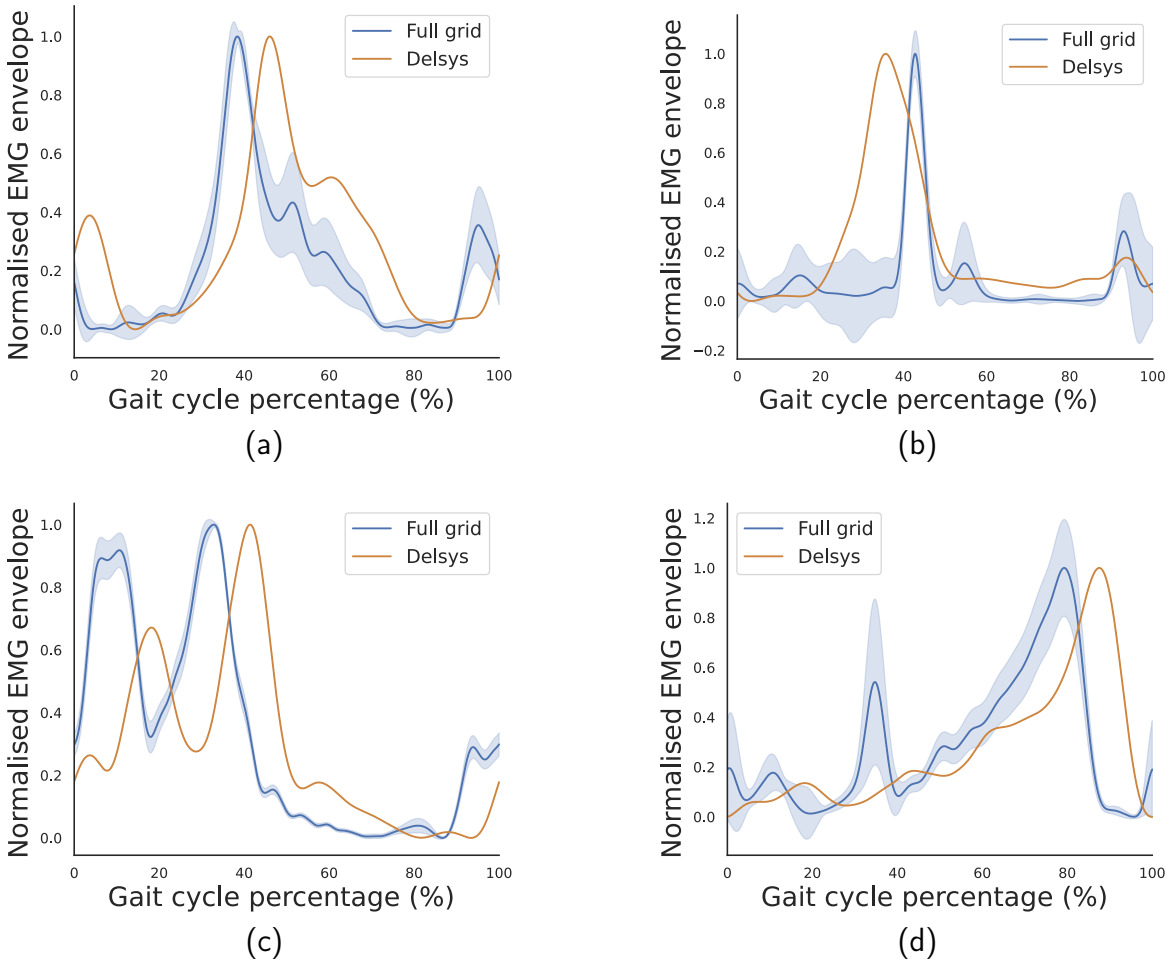


Figure 3.13: Comparison of the average bipolar sEMG signal envelopes across the HD-EMG grid with the bipolar signals from the Delsys system for *(a)*: the RF, *(b)*: the BF, *(c)*: the TA, and *(d)*: the SO. The HD-EMG error band corresponds to the spatial variation of the signal.

signal is subjected to external disturbances not accounted for during training; wearing an orthosis, changing speed, electromagnetic interference, electrode re-application, sweating or fatigue (Oskoei et al. 2007b; Venugopal et al. 2014). To account for the signal alterations caused by external factors, data augmentation approaches, including non-linear transformations of the training set, are promising. This would require in-depth analysis of the impact of different environmental factors to reproduce them into the training distribution, which could be accelerated by the acquisition of spatially robust features.

Steady-state level ground walking was selected as target motion as it facilitated the efficient collection of many repetitions of the same movement. However, the intended movement is not without variation. A degree of variability is present in the dependent variable due to small changes in gait patterns adopted by the participants over the course of a trial. A well performing model must be able to track these effects. This is a limitation of the described consistency metrics, as they don't distinguish between meaningful, informative variability and noise. A non-steady-state locomotion setting, including transitions such as turning or coming to a stop would be suitable to explore the intent estimation aspect further.

Increasing the size of the input window led to an increase in performance and a decrease in the sensitivity to specific electrode placements (Figure 3.6). However, longer input intervals imply greater computational cost both during training and inference, making the optimal window size vary depending on the requirements of a given application. The impact of this parameter on the model's sensitivity to intent changes in particular should be examined. Lastly, in addition to the window size, the influence of other TCN hyper-parameters on prediction stability should be investigated.

3.5 Conclusion

The investigated consistency metrics were validated as indicators of signal quality, with the selected samples leading to higher output accuracy from the TCN. This correlation could be used to identify electrode placements that will perform better in data-driven models. However co-

variate shifts due to changing conditions and variability in inter-session electrode re-placement makes optimising for a single, high signal quality location unreliable. The use of HD-EMG grids for data acquisition with the aim of bipolar applications enables the inclusion of valuable spatial information that improves model robustness to electrode shifts. While there appears to be transferability between HD-EMG and bipolar sEMG signals, future works should focus on domain transfer methods to seamlessly reproduce bipolar signals from HD-EMG grids, and construct data augmentation strategies based on expected distributional shifts for all deployment environments. Current work involves applying these findings to lower-limb exoskeletons in order to achieve robust AAN control.

Acknowledgment

Special thanks to Professor Ferdinando Rodriguez y Baena, director of the Hamlyn Centre for Robotic Surgery, for granting us access to his lab and equipment, and our volunteer test subjects.

Chapter 4

High-level Locomotion Intent Estimation

4.1 Introduction

We demonstrated in the previous chapter that TCNs are powerful tools for motion intent estimation, provided their potential to overfit is addressed during their training. However, the LL joint kinematics we extracted with them may not be appropriate to drive wearable robotics on their own. The predicted kinematics will be highly specific to an individual's motion patterns. Even more so, many essential aspects of locomotion, such as modulation of the gait patterns to perform turns are de-emphasised by focusing solely on joint angles.

Abstract HL parametrised intent representations reduce the complexity requirements for lower-level controllers, and have regularising effects (Merel et al. 2019). A further benefit is the modularity of the interface between controller levels (Tucker et al. 2015). An intent estimator can be paired up with different motion synthesizers fine-tuned for users and vice versa. Lastly, if the intent estimator stage avoids predicting joint angles directly, biases and errors in the prediction do not necessarily correspond to biomechanically impossible or unsafe states. User safety being inherent this way is an essential property expected in P&O control (Tucker et al. 2015; Anam et al. 2012). In this chapter we wish to extend our use of TCNs to abstract, higher-level representations of locomotion goals, which can then be used to explore intent-driven prosthetic control policies in the following chapters.

How we decide to describe and quantify our motion intent is determined by which characteristics we deem worthwhile to capture, and which ones are best relegated to lower-level systems to handle. It is an act of compressing and encoding the motion to a format from which it may be restored accurately. As such, different representations will be suitable for different locomotion settings and for different devices, and will consist of more or less information and variability. A key reason why many robotic LL control strategies can not be applied for wearable systems, is that the device does not have agency over the human body, nor does it have a perfect understanding of its configuration (Martin et al. 2017). The role of the HL controller is to extract useful estimates of the current and future state of the user.

Many authors distinguish between high and mid level controllers. The line between these is often blurry, and where it is drawn is left up to their designer. For the purposes of this chapter, we will be treating controllers as HL if their output does not directly modulate an actuation system, but rather conditions a lower level control system that still maintains a level of sensing, autonomy and decision making independent of the higher-level (i.e., the output of the modulated lower level system provides meaningful additional information not present in the higher level).

4.1.1 High-level locomotion intent representations

A wide range of options are available in the literature for representing HL locomotion intent, each containing an exponentially growing tree of design choices within. Discussing them all in the detail they deserve is outside the aims of this chapter, which instead focuses on a single intent representation appropriate for our purposes. However, as context we provide a non-exhaustive summary of the main alternatives identified, ordered roughly from most to least abstract:

Locomotion mode classification:

The most common approach is to define a set of locomotion modes, out of which only one at a time may be active. Common categories include standing, walking, stair ascent, stair descent, slope climbing, sitting down, standing up. The currently active state is maintained over any number of gait cycles or modulations of the gait cycle, until the environmental or user determined factors change. Commonly this approach utilizes a classifier that selects an active state based on sensor information, based on which the parameters of lower level controllers (such as the impedance of actuated joints) are chosen. The transitions between the state may be constrained by a set of rules, forming a FSM (Quintero et al. 2011; Peeraer et al. 1990).

This approach offers many benefits, mainly its complete modularity. The LL controllers may be treated and solved for separate from each other, and no communication is necessary between the HL and LL controllers. These classifiers are often combined with gait phase estimators, which keep track of the cyclic progress with the class of motion currently executed (Tucker et al. 2015; Holden et al. 2017). Blending between the discrete classes of motion, while sharing a continuous phase variable is one solution to transition between the discrete locomotion modes (Starke et al. 2019).

The main limitation of this approach is that the needs of the assisted human and their environment are not always stereotypical, and may not fit well into one of the existing categories. For example, there are many steps and stairs of various dimensions that need to be adapted to to assist with ascension and descension. To tackle them robustly either additional mid level control layers must be implemented to extract the necessary information from the environment, or the set of classes must be expanded to include variations on each class. Once combinations and transitions of classes start to also be accounted for, the controller can balloon into hundreds of potential classes; our entire motion behaviour is challenging to be described as a FSM (Buttner 2015).

Walking path trajectory control:

One can summarise locomotion in terms of its total outcome, the displacement. The centre of mass (CoM)'s (or its approximation's) position is usually the target for this approach. This can go beyond simply defining a target location and facing direction at point B relative to a current position at point A (Mombaur et al. 2010); a horizon of future desired Cartesian space coordinates may instead be used to describe in greater detail how the CoM evolution takes place. This is relevant as locomotion is not a conservative/path-independent process; key outcomes such as stability, energy use or comfort depend on the path taken to travel between two points. In wearable robotics settings the control space usually only involves a subset of all possible trajectories, or are described in a highly parametrised fashion, e.g., by only defining the current desired horizontal walking velocity (Kidziński et al. 2020; Karulkar et al. 2021). Including a predictive window of upcoming future velocities allows capturing other gait features such as progress in approaching turns or stops. Anticipating changes in intent could allow wearable robotic devices to adjust their behaviour in preparation of an event, or to overcome delays due to filtering and computational processes. In related applications outside wearable systems, it has been applied to investigate robot path planning (alone or when accompanying a human) (Mombaur et al. 2010; Maroger et al. 2021), avoiding collisions with humans (Q. Wang et al. 2010), and surveillance use cases (Rudenko et al. 2020).

This approach conceals most of the ambulation aspect of locomotion (which relates to the movement of the lower limbs), removing concerns such as planned foot placement or gait phase out of consideration. However, it is still much more versatile than the locomotion mode classification. It captures behaviours such as standing, walking, running and transitions in between, as well as turns and strafing. This versatility is demonstrated through its success when used in motion matching (Clavet 2016; Bergamin et al. 2019).

If vertical displacement is also considered, then stair and ramp navigation can also be encoded using this approach. Since each trajectory point may continuously vary, no discrete variations are needed, and it is possible to smoothly interpolate between control states. This property makes it possible for lower level systems that incorporate function estimation to reuse existing

knowledge between similar control states, tackling the potential problems when switching from a discrete control space to a continuous one.

The level of precision necessitated in the locomotion task plays a role in the shaping of the trajectory. Similar to Fitt's law being applied in manipulation task, similar concepts are sometimes used in locomotion, such as the law of turning (Zhai et al. 2003).

Foot placement and zero moment point (ZMP) control:

The primary dynamic interaction with the environment during locomotion is the contact with the ground and the resulting reaction force. This force is what drives the human body forward in the end. Treating the human body as a mobile point mass like in the previous method may miss key information where the details of this interaction may heavily change the outcome of the movement. Avoiding tripping on steps, navigating uneven terrain and maintaining a suitable margin of stability all depend on where are our feet placed. Biomechanical models, such as inverted pendulum models can then be used to extrapolate dynamics of the human body system. For example, the ZMP can be identified. Constraining the ZMP in controller outputs can provide guarantees on the stability on the behaviour. The actuators of the wearable (and its LL controllers) can then be modulated to drive the modelled system towards states deemed optimal during development while still achieving the same HL goals described in the walking path method (which could be estimated in parallel or reconstructed from the foot placement).

4.1.2 Input modalities

As discussed in Chapter 1, all of these representations can be combined in a modular fashion with various modalities of input. Early systems relied on direct volitional control via switch-interfaces (e.g. buttons on a wrist-pad) or user-initiated shifts in the CoM (Yan et al. 2015). Myoelectric, kinematic inertial sensors can also be applied for this purpose (Tucker et al. 2015; Wentink et al. 2014). An alternative view of hierarchical control strategies is treating HL controllers as independent input processing steps, which are then combined in the LL in a

multimodal late-fusion scheme (Pawłowski et al. 2023). As such, the input signals should be selected to suit the responsibilities of a given HL controller. Due to their involvement in turning activities, we collected sEMG signals from the ankle inverter/everter muscles and trunk muscles (Hase et al. 1999; Anders et al. 2007; Ceccato et al. 2009). While sEMG is an important signal source for responsive controllers, additionally we also investigated body posture as an input which we hypothesise can provide stability to the predictions when combined with the otherwise variable sEMG signal.

4.1.3 Chosen representation

We select the horizontal walking path trajectory as our HL intent format. This is for three main reasons. First, its versatility to represent a wide range of non-stationarities in gait. Second, it is readily available to be synthesised in simulation, as described below in Section 4.4.4. Lastly, it has a close relationship with motion matching, a method that will be applied in Chapters 5 and 6.

One way to represent the walking path trajectory is through a cost function, via which any trajectory may be evaluated. Then the optimal trajectory with the lowest cost may be searched for, subject to the parameters of the cost function (Maroger et al. 2021; Mombaur et al. 2010). Mombaur et al. (2010) has identified through inverse optimal control that human walking path trajectory's cost function can be described through terms on reducing time taken to reach the goal, terms constraining the linear and angular accelerations and terms defining final position and orientation. In contrast, terms constraining the velocity, jerk and energy expenditure seem to be less fundamental. They also remark a strong preference to forward motion, a quasi-non-holonomic quality of trajectories, i.e., we prioritise moving towards where we face and only gradually changing our facing direction, instead of freely moving in any direction (e.g., sideways and backward).

We selected a critically damped in velocity forward model for our trajectory, which is a second-order model optimal in terms of time for reaching target velocity and facing direction without overshoot (which would necessitate additional periods of acceleration). It is subject to a fixed

stiffness parameter which allows tuning the trade-off between time to convergence and maximum acceleration. Its formulation will be detailed in Section 4.2. The output space is 6 dimensional, regressing from over 8000 values of EMG and posture. The estimator must leverage the redundancy available in the input space to make low-variability predictions with heavily reduced dimensionality.

Both bipolar sEMG and body posture derived input variables were investigated, as well as their multimodal combination. These were regressed to the walking path trajectory using the TCN-based architectures presented in the previous chapter. In this study we did not consider the effect of electrode placement and signal consistency like earlier, rather we focused on contributing a proof-of-concept continuous locomotion intent estimator.

4.2 Methods

4.2.1 Walking path trajectory

The strategy for constructing a HL locomotion intent estimator is illustrated on Figure 4.1. It outlines the requirement of collecting motion capture and other signal modalities, such as sEMG. Once these synchronised data are recorded, the motion capture can be preprocessed to extract walking path trajectories. New collection of motion capture data collection was required, as available datasets at the start of the project did not have concurrent myoelectric recordings that were needed to establish HMI based control, or did not feature turns or transitions between activity modes (CMU 2003; Müller et al. 2007; Ofli et al. 2013; Moissenet et al. 2019).

From the perspective of a wearable robotic device, the global position and facing direction of the human is irrelevant. Instead, motion intent should be described from a local reference frame. We quantify the position of the local frame $\mathbf{p}^L \in \mathbb{R}^2$ throughout the recording, defined as the midpoint of marker positions of the pelvis segment from the 2.5 version of the Conventional Gait Model (Leboeuf et al. 2019). This position is then projected onto the horizontal plane. We also define the orientation of the local frame with three unit vectors:

1. The direction pointing from the right ASIS to the left one, projected onto the horizontal plane.

2. The normal of the walking surface.

3. A mutually orthogonal vector.

Taking vectors 1 and 3 forms a rotation matrix $\mathbf{R}^L$ on the 2-D horizontal plane. For each time point t_0, we create a walking path trajectory tensor $\mathbf{P}^G$ up to a horizon of $H = 1$ second. This collects how the global horizontal position of the local reference frame's origin changes throughout the trajectory.

$$\mathbf{P}^G(t_0) = \begin{bmatrix} \mathbf{p^G}(t_0) \\ \mathbf{p^G}(t_0 + \Delta t) \\ \mathbf{p^G}(t_0 + 2\Delta t) \\ \mathbf{p^G}(t_0 + 3\Delta t) \\ \vdots \\ \mathbf{p^G}(t_0 + H) \end{bmatrix}^\mathsf{T}, \tag{4.1}$$

where H is a multiple of Δt. The starting time of the given trajectory is denoted by t_0.

$\mathbf{p^G}(t_0)$ is then subtracted from each position vector element in this tensor and then multiplied by $\left(\mathbf{R}^L\right)^{-1}$. This forms a local walking path trajectory $\mathbf{P}^L$, as described from the pelvis' orientation at its start. The 2-D local position within this 1 second horizon from a given starting time t_0 is denoted as $\mathbf{p}_h^L(t_0)$, where $h \in (0, 1]$ is the progress along the trajectory, a value of 1 indicating the endpoint of the horizon. $h = 0$ can be excluded from all further discussion to its triviality as $\mathbf{p}_0^L(t_0) := 0$. An illustration of this trajectory in the context of a human walking is available on Figure 4.12.

During a locomotion trial, each instance of $\mathbf{P}^L(t_0)$ can be paired up with a window of muscle activity:

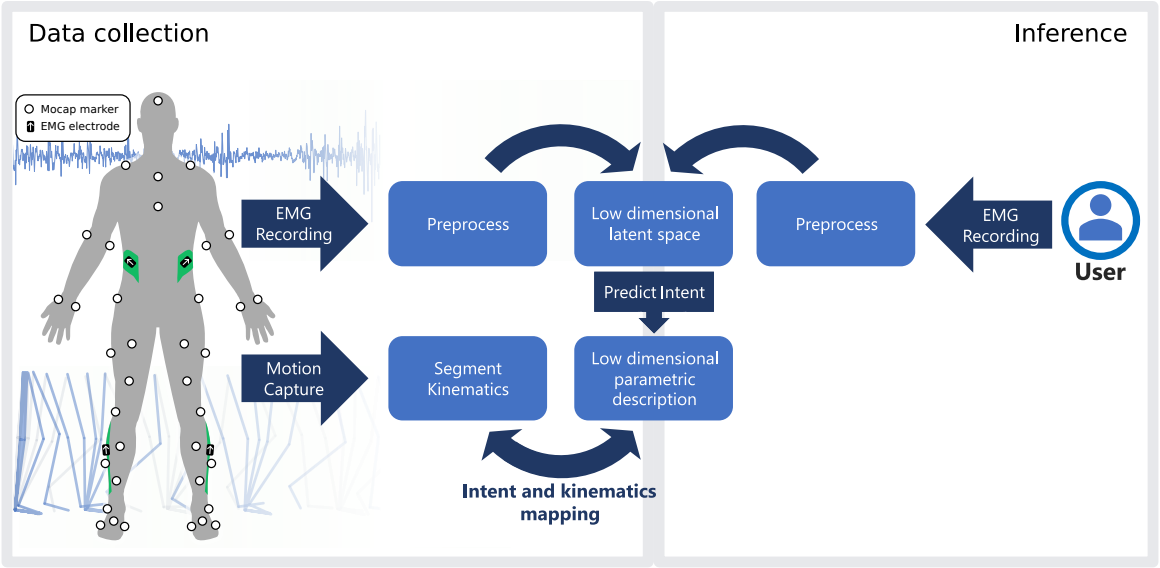


Figure 4.1: Concept for collecting the data set to construct a neural locomotion controller. Motion data can be paired with biosignals like sEMG and body posture during demonstrations of walking. Feature extraction can be performed on both motion modalities to focus the data to represent desired aspects of the behaviour. In the case of sEMG this is performed via TCNs, while for the motion capture the parameters of the critically damped trajectory model play this role ($\mathbf{v}_0$, $\mathbf{v}_d$ and $\mathbf{a}_0$). A mapping can be learnt from these paired embeddings. Once an intent estimator is trained, it may be applied on novel biosignal data coming from the same subject, to forecast their desired walking path. The predicted intent may then also be mapped back to full body kinematics via a method such as motion matching (covered in Appendix A). Anterior view of subject show, for full marker and electrode placements see Appendix C.

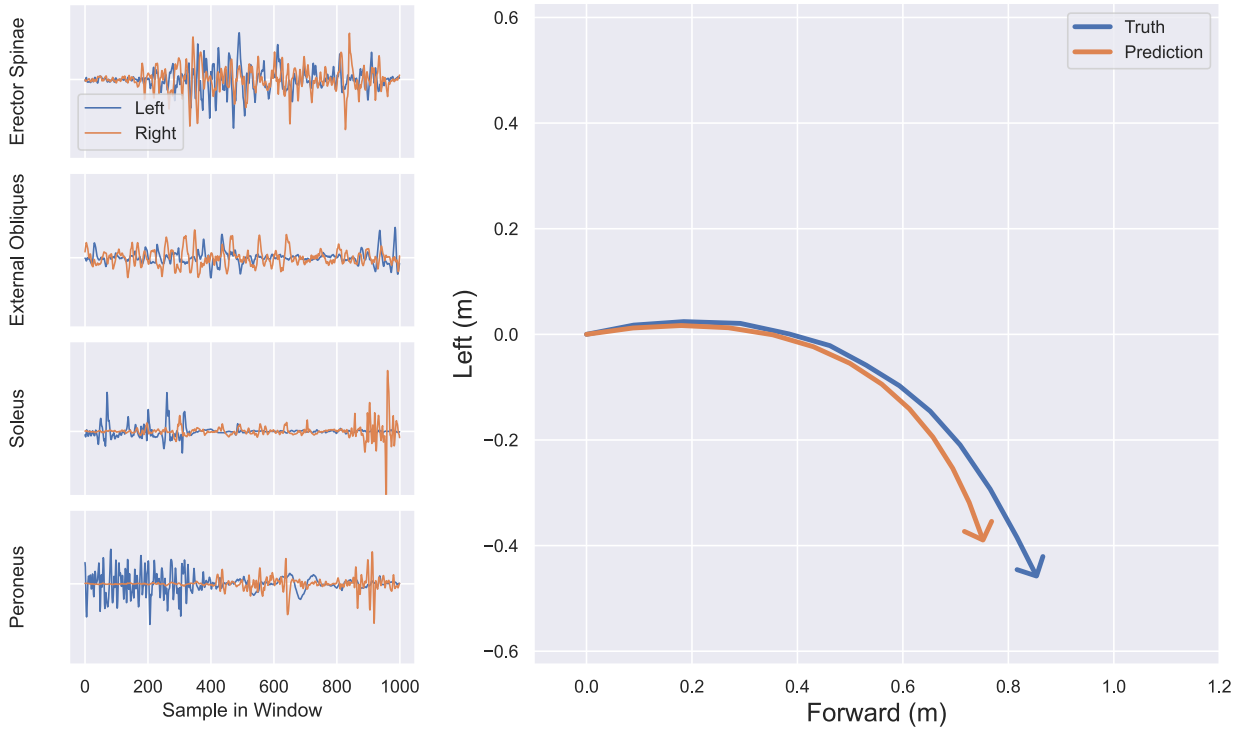


Figure 4.2: 0.5 second window of 8 channels of bipolar sEMG, processed by the TCN and regressed to the walking path trajectory, represented in the local pelvis reference frame. Each trajectory has a corresponding 8 channel window of sEMG, in a sliding fashion. The right column shows the outcome variables: 1 second long horizon of the pelvis' position, relative to the pelvis' position and orientation at the trajectory's start. The event shown is a 90° turn to the right.

$$\mathbf{m}(t_0 - w) \qquad w \in (0, W], \tag{4.2}$$

$$\mathbf{m}(\mathbf{t_0}) \in \mathbb{R}^C, \tag{4.3}$$

where W is the size of the window of EMG in time and C is the number of EMG channels used, and w is the index of the sample in the window.

Figure 4.2 shows a frame of the recorded data set. The role of the intent estimator is to predict the pelvis positional trajectory using a 0.5 second long window of sEMG signals. Windows were taken with a 20 samples stride, leading to a large degree of overlap. Taking windows at smaller strides increases the size of the data set but due to the statistical dependence between overlapping windows, this approach has diminishing returns.

Based on the findings of Chapter 3, we once again adopt a sEMG window size of 0.5 seconds, which is twice as long as usually indicated for upper limb myoelectric interfaces (Smith et al. 2011). This is justified due to the longer time scale events happen during locomotion when compared to frequent and diverse changes in manipulation. The main issue of longer windows is the additional delay they introduce when estimating window-wide metrics (Smith et al. 2011). However, in a TCN the effective delay is automatically tuned, as different weights for different segments of the window are learnt.

4.2.2 Temporal Convolutional Network

We hypothesised that the architecture of the knee angle estimator model from the previous chapter will generalise to this HL intent estimation task too. We used a TCN setup that was nearly identical to the one used in Chapter 3 (see Figure 3.5). A key difference was that in this case multiple simultaneous bipolar sEMG channels were collected and processed using channel-wise separate convolutional layers, instead of the last chapter's single channel input. Note that while this task can be thought of as a sequence-to-sequence problem, a purely convolutional network is disadvantaged at this particular task. Due to the equivariance property of these layers, periodic inputs will produce periodic outputs. This is of course a desired property in the feature extraction (encoding) stage. However the walking path trajectory has significant aperiodic components (e.g., the case of steady state walking). For this reason, the dense layers after the encoding stage are also kept from the last Chapter. This way time dependent properties can be learnt to modulate and offset the predicted trajectory samples. Like before, we train subject specific models.

The experiment concept was trialled first in a pilot study with a single participant, with whom a proof-of-concept data set was recorded. We describe this pilot, summarising its results and shortfalls to provide context for the changes we implemented afterwards in a six subject follow-up study.

4.2.3 Pilot study

Due to its novelty, we tested the feasibility of the sEMG driven walking path trajectory estimator with a pilot study involving a single female subject recruited from the Imperial College Department of Bioengineering. In this pilot, only trunk muscle signals were collected bilaterally at two heights of the Erector Spinae muscle groups and abdominal obliques for 8 bipolar channels using a Delsys Trigno system. A reduced motion capture marker set was used, only capturing the pelvis and thorax 6 DoFs kinematics and 3 DoFs foot kinematics. Thorax marker set was adapted from (Boser et al. 2018), and electrode placement was adapted from (Anders et al. 2007). This recording session took place at the Musculoskeletal Biomechanics laboratory of the Michael Uren Biomedical Engineering Research Hub, London.

The pilot regressor's output was 6 dimensional: it predicted $\mathbf{p}^L_{1/3}$, $\mathbf{p}^L_{2/3}$ and $\mathbf{p}^L_1$ concatenated. A key shortfall of the pilot model was its use of non-causal processing, where the trajectory predictions were performed at the start of the window of EMG, instead of a causal situation where only past input signals are used. This was necessary to reach a coefficient of determination larger than 0.5.

Contrasting with the causal window in Equation 4.2, the non-causal window used in the pilot can be denoted as such:

$$\mathbf{m}(t_0 + w) \qquad w \in [0, W), \tag{4.4}$$

$$\mathbf{m}(\mathbf{t_0}) \in \mathbb{R}^C, \tag{4.5}$$

5-fold cross-validation was used on the angular walk subset of the data to quantify performance. The cross-validation splits were made between separate recording files and not all samples. This distinction is made due to the overlapping nature of the windows. A split like this guarantees no sample from a window in the validation set is present in any window of the training set.

Time series of the predictions, as well as further discussion on the results from the pilot study

Sample (s)	RMSE (m)
1/3	0.048 ± 0.0413
2/3	0.098 ± 0.083
1	0.179 ± 0.142
Overall	0.108 ± 0.112

Table 4.1: The non-causal pilot model's euclidean distance error and its standard deviation, averaged across the five cross-validation folds. Rows represent the position of the sample in the trajectory horizon, in terms of time. 37188 frames were used in total to quantify performance.

is available in Appendix B.

The angle of the position vector pointing to the pelvis position at 1 second in the horizon was calculated for both truth and prediction. Average correlation coefficient in these angles was estimated at 0.6. This measure oversimplifies a lot of details regarding gait, but it quantitatively shows that even though non-causal windowing was used, there was still plenty of room for improvement.

Nonetheless, the pilot study's results indicated that there was useful information in trunk muscle sEMG for estimating intended gait path, not only in terms of gait onset and stopping, but also in direction and extent of turning. We proposed and implemented a series of changes to be made for the follow-up study, which we hypothesised could lead to improvements.

4.2.4 Follow-up study

We summarise the most important flaws of the pilot model:

- Trajectory estimates in the lateral direction were more variable and less accurate when compared the anterior direction.
- When testing in causal regression schemes, turning direction could no longer be reliably extracted.
- The intent estimator's predictions were at only three points along trajectory: $\mathbf{p}^L_{1/3}$, $\mathbf{p}^L_{2/3}$ and $\mathbf{p}^L_1$. This made extracting kinematic features like the current or future walking

velocities imprecise.

To address these limitations, the following changes to the experimental setup, data collection and model structure were proposed for the follow-up study:

- We extended the trial protocol, with increased number of repetitions for a larger dataset per subject.
- We moved 4 out of 8 electrodes to the lower limb, for a greater diversity in signals.
- We switched to a full body motion capture markerset.
- We increased the resolution of the output trajectory, and introduced an alternate, damped trajectory model based regression via a custom deep-learning neural network (NN) layer.

Data acquisition

Prior to performing the experiments for data acquisition, approval covering the protocol was received from the Imperial College Research Ethics Committee (ICREC reference: 22IC7765). Participants were recruited using convenience sampling, advertised through email and word of mouth. The following inclusion criteria was used:

- Between 18 and 60 years old.
- Speak and understand English.
- Able to walk without assistance and assistive devices.

The following exclusion criteria was used:

- Known second or third trimester pregnancy.
- Previous diagnosed neurological, musculoskeletal or physiological condition that affects gait.

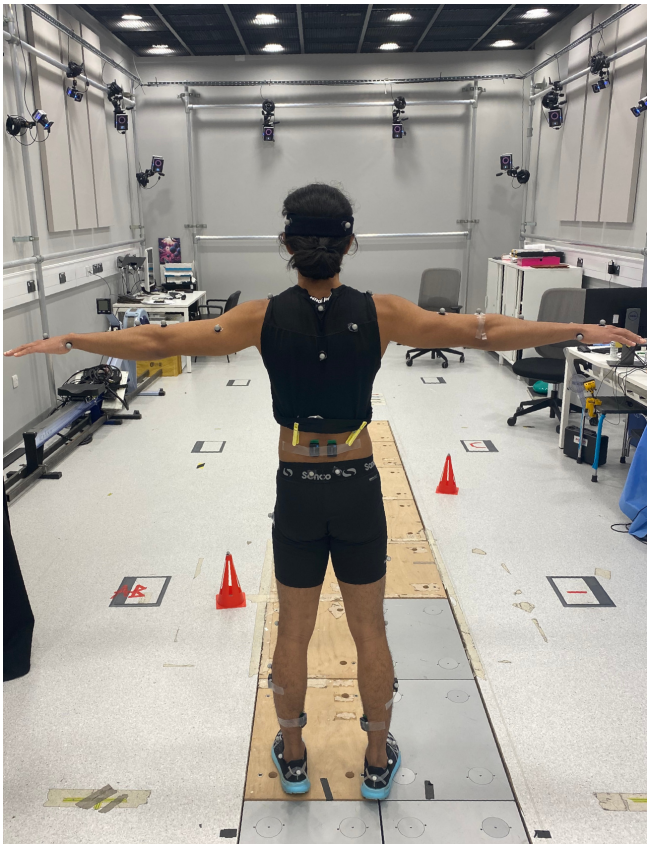


Figure 4.3: Experimental setup with one of the subjects at the Imperial College Biodynamics lab, showing motion capture system, the Conventional Gait Model markerset (Leboeuf et al. 2019), floor markers and bipolar sEMG electrodes.

- Obesity, as defined by the National Health Service (2019).
- Moderate to intense exercise 4 hours prior to experiment start.

Pregnancy and obesity are exclusion criteria to mitigate occlusion of hip markers and to aid sEMG signal quality (Kuiken et al. 2003; Peters et al. 2001). Exercise is limited as sweat and fatigue negatively impact myoelectric signal quality (Oskoei et al. 2007a). All six participants in the follow-up study were male, with a median age of 23 ($\pm$1.2).

Motion capture was recorded using a Vicon Vero system at 200 Hz. Simultaneous bipolar sEMG was acquired with a Delsys Trigno electrodes at 2 kHz, synchronised internally in Vicon Nexus (Vicon Motion Systems Ltd 2020). The 2.5 version of the Conventional Gait Model was used as markerset (Leboeuf et al. 2019), and the SENIAM guidelines were followed for electrode placement (Hermens et al. 1999). Marker trajectories were gap-filled in Vicon Nexus. The experiments took place at the Imperial College Biodynamics lab.

To record a diverse set of level-ground locomotion behaviours efficiently, a motion capture protocol similar to ones used when collecting motion matching data sets was used. This was constructed based on guidance from Clavet (2016), Buttner (2015) and Zadziuk (2016). We labelled the primary locomotion trial type we collected as "Angular Turns". It consisted of visiting the vertices of a trapezoidal path indicated in the motion capture volume through marks on the surface and visibility cones. By walking through this path, turning around and returning, we recorded a variety of turns at different angles succinctly. An excerpt of the protocol document is shown below for reference (Figure 4.4).

Appendix C contains the rest of the experimental protocol used to construct the data set for the intent estimator. The motion trials making up the experiment aim to cover a wide range of non-steady state level ground locomotion behaviours, focusing on turns of various angles, and transitions between them. In addition to turns at discrete points in time (Hase et al. 1999), human turning strategies include gradual, arced turns at various radii (Orendurff et al. 2006), therefore our demonstrations include examples of both. We aimed to have at least 10 trials recorded in the "Angular Turns" category per subject, but it was possible to record 16 trials for

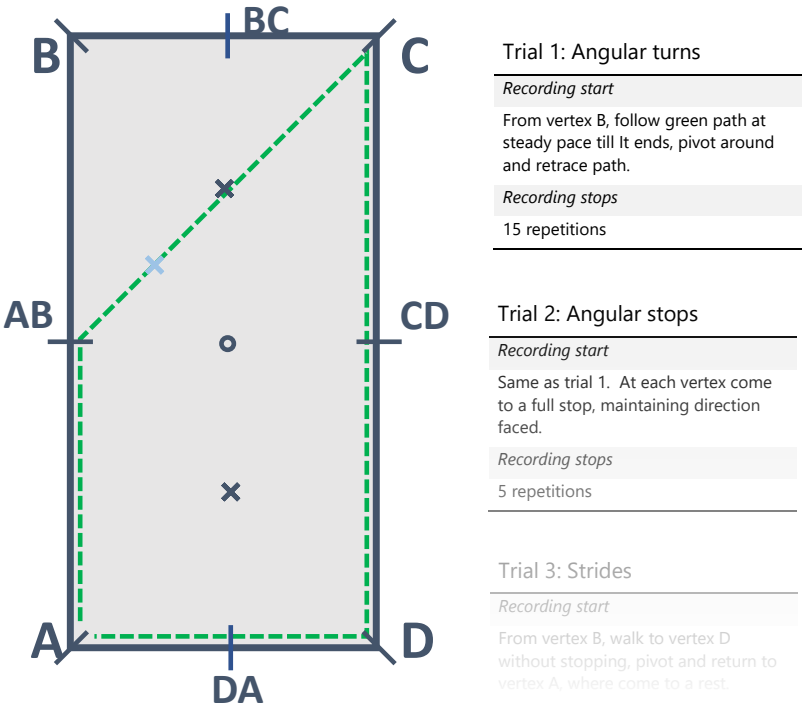


Figure 4.4: Top down schematics of the surface markings applied on the motion capture volume, using which the motion trials could be guided. The A-D distance was 2.3 m, the A-B distance was 7 m, and A-AB was 4.7 m. To the right we show 2 out of the seven motion trial type descriptions. For the full list see Appendix C.

Subject 2, and 20 trials for Subjects 3-6 thanks to increased efficiency in the data acquisition as the study progressed.

Input modalities

We moved 4 out of 8 electrodes to the lower limb, targeting the medial Soleus and Peroneus muscle group of the ankle bilaterally. Ankle stabilising muscles such as these play a key role in many turning strategies (Hase et al. 1999). We also observed greater differentiation of activity between sides than from the trunk. On Figure 4.5 we can see that while the Erector Spinae activity increases for each step (twice per gait cycle, following each heel strike bilaterally), the ankle muscles only have a burst of EMG during their stance phase, once per gait cycle. We also note that the lower limb signals have greater amplitudes, and less corruption for signals such the electrocardiogram (Zhou et al. 2007).

Through collecting full body kinematics, measures such as body posture also available to us for

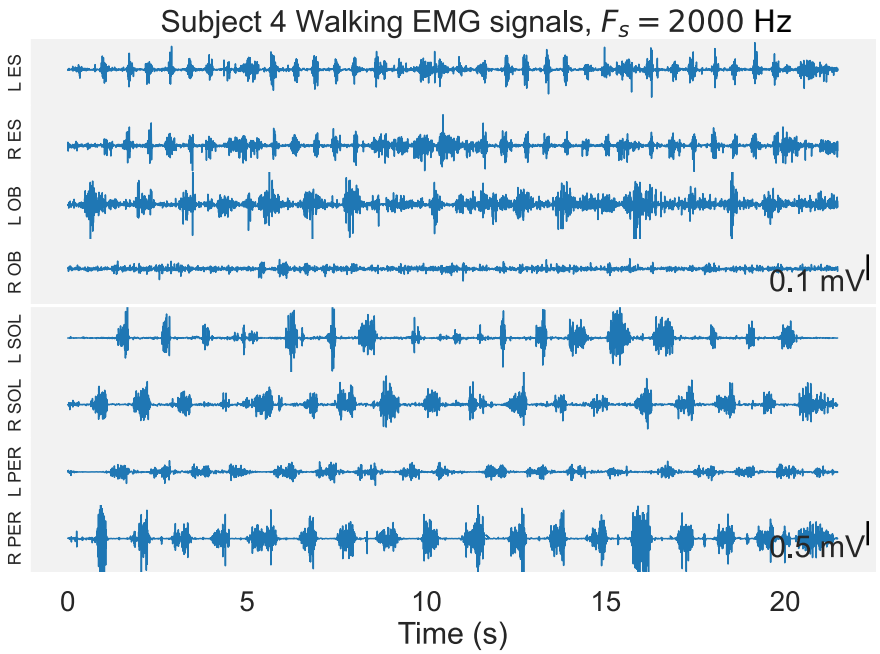


Figure 4.5: Bipolar sEMG signals recorded during one of the angular turns trials. The (L)eft and (R)ight side's Erector Spinae (ES) and oblique (OB) muscles' EMG is shown in the top four rows, the Soleus (SOL) and Peroneus (PER) muscles' signals shown in the bottom four.

use in intent prediction. With this, we can synthesise the signals of wearable motion sensors placed on arbitrary locations on the body even after the experiments are concluded. The reduced variability of these signals may help improve the stability of our predictions in contrast to EMG only models.

The orientation of a body can be represented in many formats. Different formats may aid or hinder the learning of motion concepts with NNs. The most common representations include Euler angles, quaternions and rotation matrices (Pavllo et al. 2020). Euler angles, being unbounded, having non-unique representations and suffering from singularities are only appropriate for rotations in narrow ranges of motion. Quaternions are bounded and always stable, but due to antipodal quaternions (-1 multiples) also mapping to the same rotation, additional care must be taken to provide consistent input to learning models. Rotation matrices are readily interpretable, unique, bounded and without singularities, however they use the most values to represent rotations (9 against the 4 in quaternions). Luckily, due to the orthogonality constraint of these matrices, the third column is redundant. Since NNs do not use these rotation

matrices directly with linear algebra, but rather process and encode the information contained within we do not need to maintain this. Simply using two of the three columns of the rotation matrix, we can get all of the benefits with only 6 values used. Using the first two columns of the rotation is sometimes referred to as the normal-tangent representation (Peng et al. 2021), and is our choice for the pose input of our walking path trajectory estimator. We use the orientation of the thorax and head in the local frame, forming a $2 \times 6 = 12$ dimensional input vector.

Critically damped trajectory model

If we wish to extract higher order measures (velocities and accelerations) from the estimated trajectory after it is predicted, we must increase its temporal resolution (decrease Δt in Equation 4.1). This involves a wider densely connected output layer, with additional computational costs, and inaccuracies from numerical differentiation. We may instead opt for the inverse approach; predict the higher order dynamics, and integrate them to produce a trajectory at an arbitrary temporal resolution. Through this we can also enforce properties we desire our trajectory to have: guaranteed smoothness and minimal overshoot.

We parameterise the trajectory in terms of starting velocity $\mathbf{v}_0$, starting acceleration $\mathbf{a}_0 = \dot{\mathbf{v}}_0$ and equilibrium target velocity $\mathbf{v}_d$. These parameters were described in the local reference frame and assumed to be constant throughout an instance of a walking path trajectory. A critically damped system was used in the local forward (anterior) and left (lateral) dimensions, to smoothly and continuously interpolate between the two velocities. The difference between the two then decays at a rate determined by the eigenvalue λ of the characteristic equation.

$$\begin{aligned} \frac{d^2\mathbf{v}}{d\tau^2} &= 2\lambda\frac{d\mathbf{v}}{d\tau} - \lambda^2(\mathbf{v} - \mathbf{v}_d) \\ \mathbf{v}(\tau) &= \mathbf{v}_d + (\mathbf{c_1} + \mathbf{c_2}\tau)e^{\lambda\tau} \end{aligned} \tag{4.6}$$

Where $\mathbf{c_1} = \mathbf{v}_0 - \mathbf{v}_d$, $\mathbf{c_2} = \dot{\mathbf{v}}_0 - \lambda\mathbf{c_1}$, $\lambda < 0$ and τ is the time within the trajectory. To get the

parametrised future spatial trajectory $\mathbf{P}_\theta^L$ from this, the velocity in equation 4.6 is integrated with respect to τ. The solution is given in equation 4.7:

$$\begin{aligned} \mathbf{p}(\tau) =& \mathbf{v}_d\tau + \frac{\mathbf{c_3} + (\mathbf{c_1}\lambda + \mathbf{c_2}\lambda\tau - \mathbf{c_2})e^{\lambda\tau}}{\lambda^2} \\ \mathbf{P}_\theta^L \leftarrow& \mathbf{p}(\tau) \forall \tau \in (0, H], \end{aligned} \tag{4.7}$$

where $\mathbf{c_3} = \mathbf{c_2} - \mathbf{c_1}\lambda$.

By predicting $\mathbf{v}_0$, $\mathbf{a}_0$ and $\mathbf{v}_d$ our network only needs a 6-D output for an arbitrary number of samples on this trajectory. An essential benefit of this model is that since it involves simple arithmetic operations, the partial derivatives of the trajectory positions with respect to the parameters are available through the chain rule. This means the gradients of the errors in the trajectory may be propagated through them. This is a key difference of between this approach and iterative optimal control methods for trajectory generation.

We implement a custom deep-learning layer in Tensorflow to perform this process, and facilitate training TCN-based trajectory parameter estimators using motion capture data via backpropagation. The eigenvalue is not inferred per sample, instead it is a learnable parameter like other weights in the estimator, trained globally per subject through gradient descent.

Parameter fitting

As a third condition, we investigated the case where we extract the trajectory parameters from the motion capture dataset through optimisation as a preprocessing steps. We could then calculate the losses of the EMG/body position regressor directly in this parameter space. We performed bilevel optimisation. For a given value of λ, we used the trust region reflective minimisation, a newton-like least-squares method to find optimal fits for the trajectories $\mathbf{P}^L(t)$ (Branch et al. 1999). This method is suitable due to it being able to take advantage of the analytical jacobian matrix, and define constraints for the optimisation. The jacobian can

be defined with regards the loss $\mathbf{L}$, the sum of squared distances between the predicted and experimental local trajectories (kept separate for the two dimensions) $\mathbf{L} = \mathbf{E}\mathbf{E}^{\mathsf{T}}, \quad \mathbf{E} = \mathbf{P}^{L} - \mathbf{P}_{\theta}^{L}$:

$$\begin{aligned} \mathbf{J} &= \begin{bmatrix} \frac{\partial \mathbf{L}}{\partial \mathbf{v}_0} & \frac{\partial \mathbf{L}}{\partial \mathbf{a}_0} & \frac{\partial \mathbf{L}}{\partial \mathbf{v}_d} \end{bmatrix} \\ &= -2 \sum_{\tau=\Delta t}^{H} \mathbf{e}_{\tau} \begin{bmatrix} -\frac{e^{\lambda}(2-\lambda\tau)-2}{\lambda} & -\frac{1+(\lambda\tau-1)e^{\lambda}}{\lambda^2} & -\tau + \frac{e^{\lambda}(2-\lambda\tau)-2}{\lambda} \end{bmatrix}, \end{aligned} \tag{4.8}$$

where $\mathbf{e}_{\tau} \in \mathbf{E} \;\; \forall \tau \in (0, H]$.

To prevent converging to minima in areas of non-plausible parameters, we defined the following limits (with units of m/s and m/s^2 for velocities and the acceleration respectively):

$$\mathbf{v}_0 \in \begin{bmatrix} [-0.3, 1.6] \\ [-1.6, 1.6] \end{bmatrix}, \quad \mathbf{a}_0 \in \begin{bmatrix} [-4, 4] \\ [-4, 4] \end{bmatrix}, \quad \mathbf{v}_d \in \begin{bmatrix} [-0.5, 1.7] \\ [-1.7, 1.7], \end{bmatrix}$$

Where the top element of each vector of limits shows the limits in the forward direction, and the bottom element shows the limits in the left-right direction.

For each subject, we randomly selected an "Angular Turns" trial, then subsampled its available trajectories by a factor of 2 to reduce processing times. We then quantified the mean euclidean distance between estimated and experimental trajectory points across the sampled trajectories, evaluating the fit quality for the given λ. The motivation for the search for a suitable lambda per subject is demonstrated with Figures 4.6-4.8. If an inappropriate value is chosen, the lower level optimisation will converge to non-realistic values either with $\mathbf{v}_d$ or $\mathbf{a}_0$ ($\mathbf{v}_0$ is seemingly less sensitive).

It is computationally prohibitive to perform this operation for a high resolution of λ values. Additionally, there is variability and uncertainty in the fit quality measure. Bayesian optimisation is a well-suited technique to handle both of these issues, with which we can efficiently predict the optimal eigenvalues. Figure 4.9 shows the result of this process, using a Gaussian process implementation from the scikit-learn module (Pedregosa et al. 2011). Once a suitable

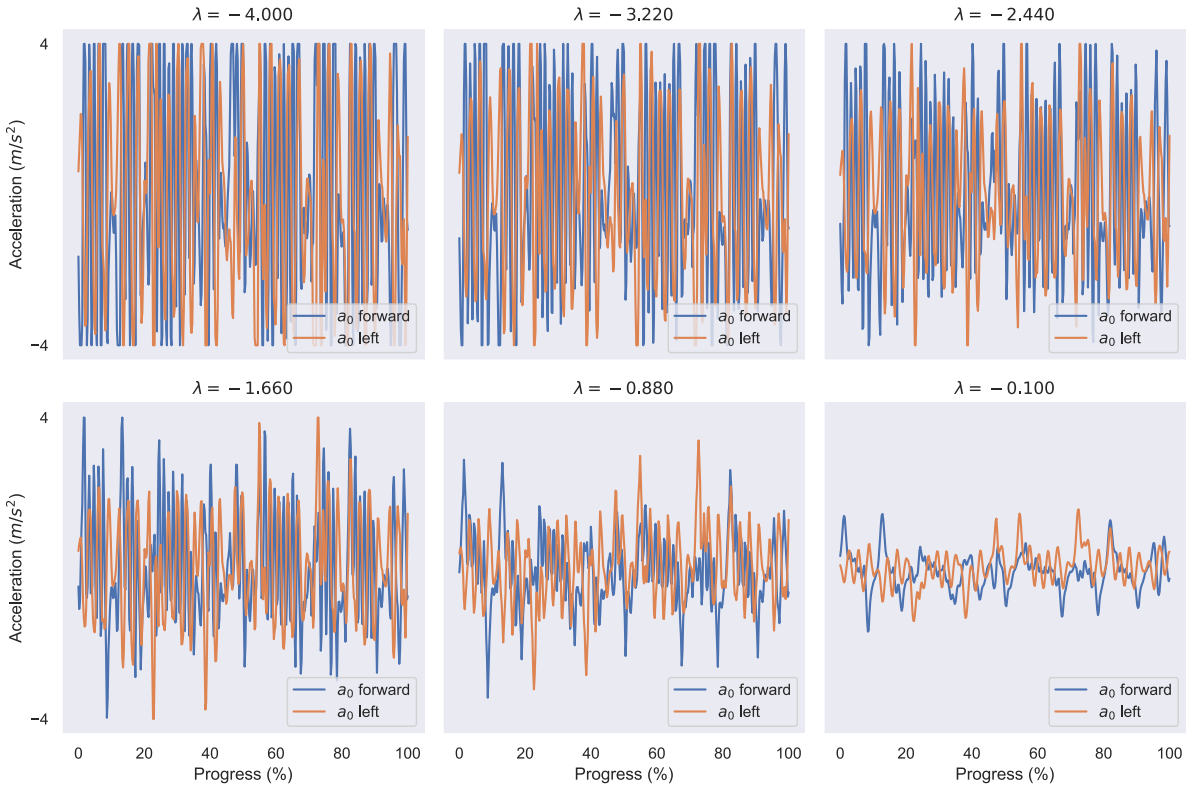


Figure 4.6: The effect of the fixed eigenvalue λ on the velocity momentum/starting acceleration a_0, determined in the inner optimisation level. Values plotted against progress along one of the angular turn locomotion trials. Too high magnitude λ is mitigated by extreme values of a_0 to prevent too sharp turns, exceeding plausible values.

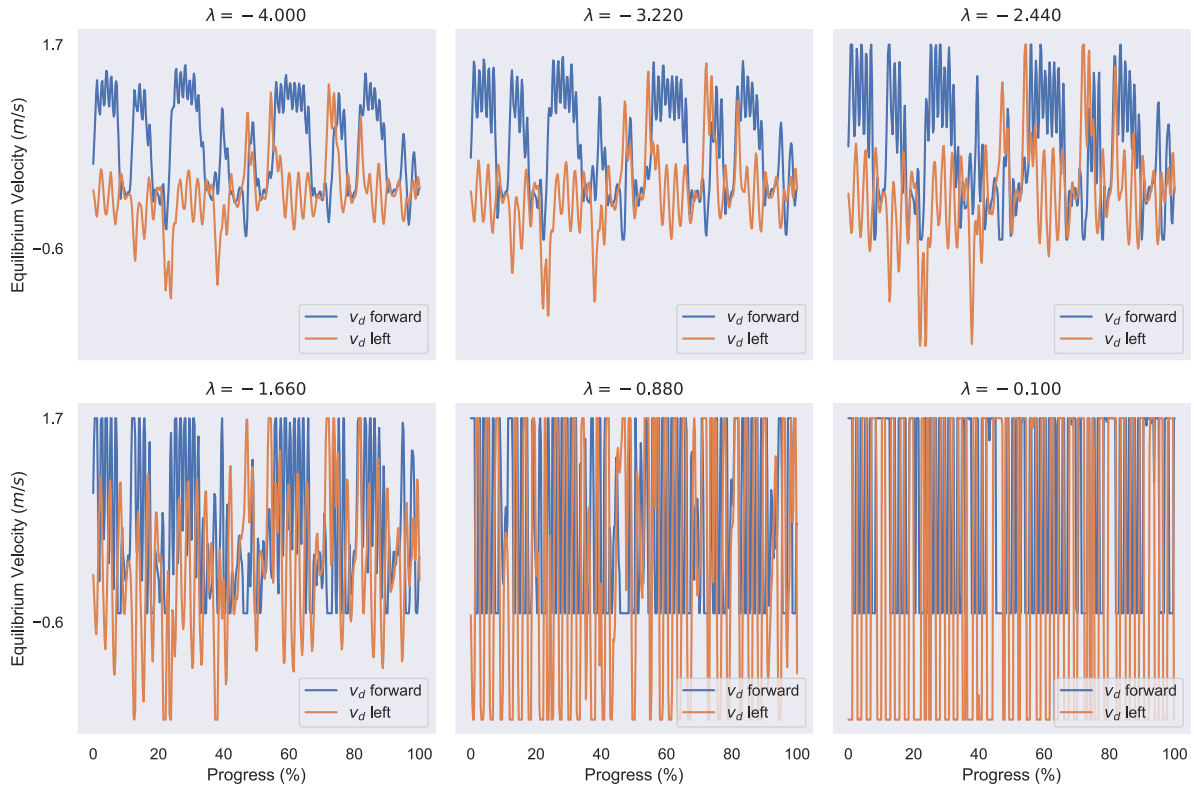


Figure 4.7: The effect of the fixed eigenvalue λ on the equilibrium/desired velocity v_d, determined in the inner optimisation level. With appropriate λ the fitted time series are continuous despite no constraint placed on this property. Too low magnitude λ is compensated by extreme values of v_d to reach the needed velocity by the end of the 1 second trajectory, exceeding plausible levels.

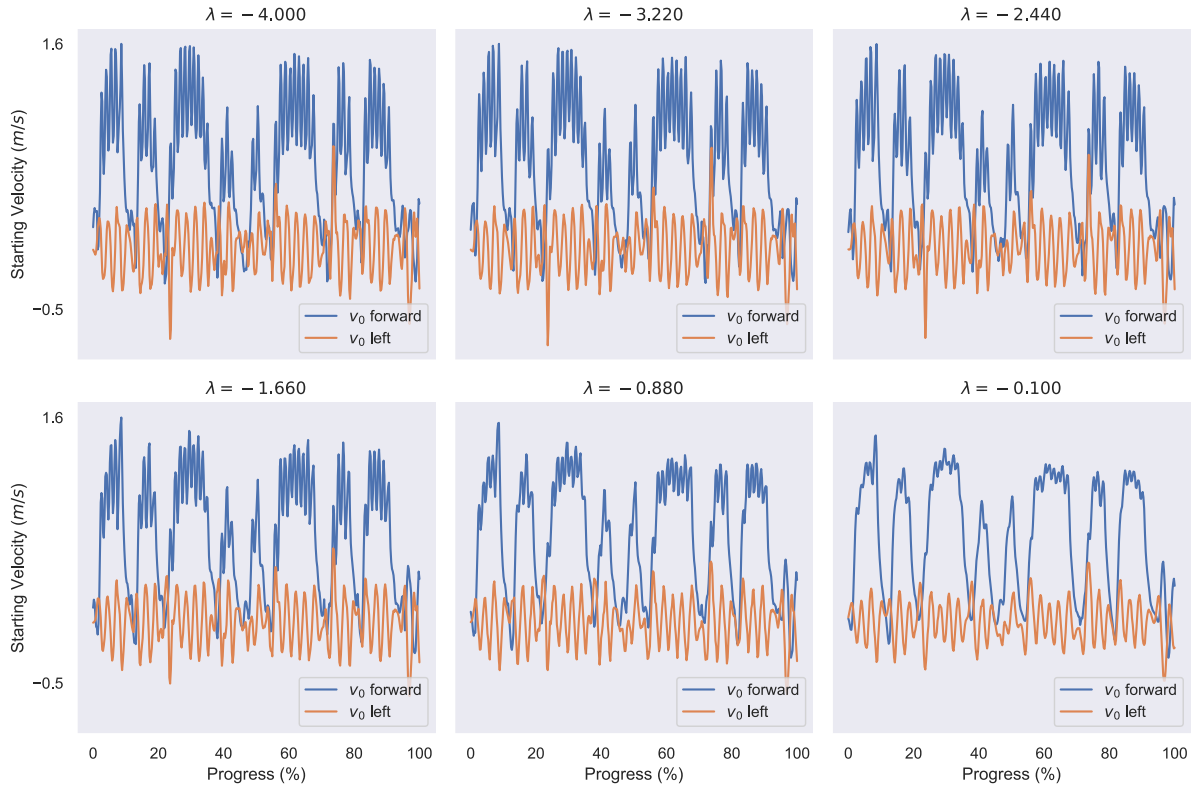


Figure 4.8: The effect of the fixed eigenvalue λ on the estimated starting velocity v_0, determined in the inner optimisation level. This parameter is less sensitive to λ than v_d and a_0. The model can take advantage of higher stiffness (and therefore faster convergence to v_d) by permitting greater variation in v_0.

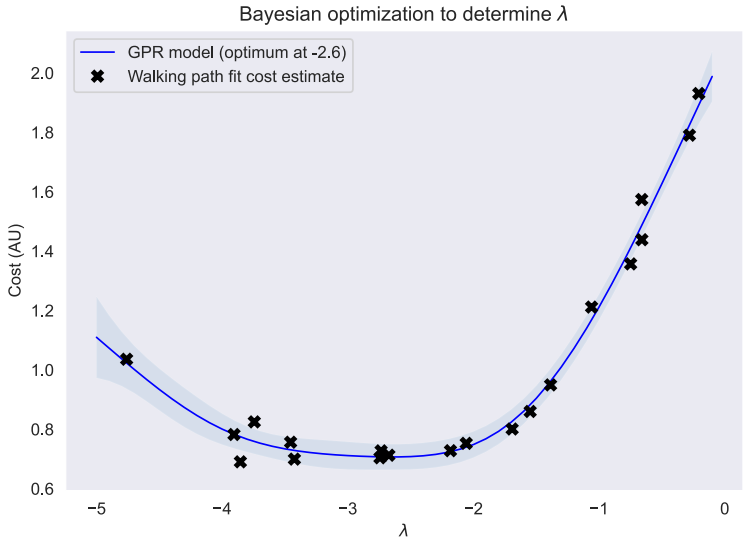


Figure 4.9: Residual of the inner optimisation level across randomly sampled trajectories from one of subject 3's locomotion trials, conditioned on the eigenvalue λ. Finding the λ with the lowest expected residual is the goal of the outer optimisation level. A Matérn kernel was used, with a noise level estimated empirically by evaluating the loss function 10 times on the same value of λ (we selected -2.7). The length scale and output scale of the optimisation were adapted to maximise the marginal likelihood. 95% confidence interval shown

eigenvalue is found, we keep it fixed and use it to extract the remaining parameters for all trials, for a given subject.

4.2.5 Conditions compared

The three variations of the TCN-based walking path trajectory estimator are illustrated in Figure 4.10. The difference between them is the result of whether there is a constraint placed on the shape of the predicted trajectory, and in what space is the loss calculated for gradient descent:

Method 1: Once the input signals are encoded, they are processed by a densly connected feed-forward NN. The output regresses directly to the x and y coordinates of the walking path trajectory's discretized samples (in the local pelvis reference frame). There is no hard-coded constraint on the shape of the predicted trajectory. We perform this at a higher resolution than in the pilot, with P^L having 12 elements along the time axis.

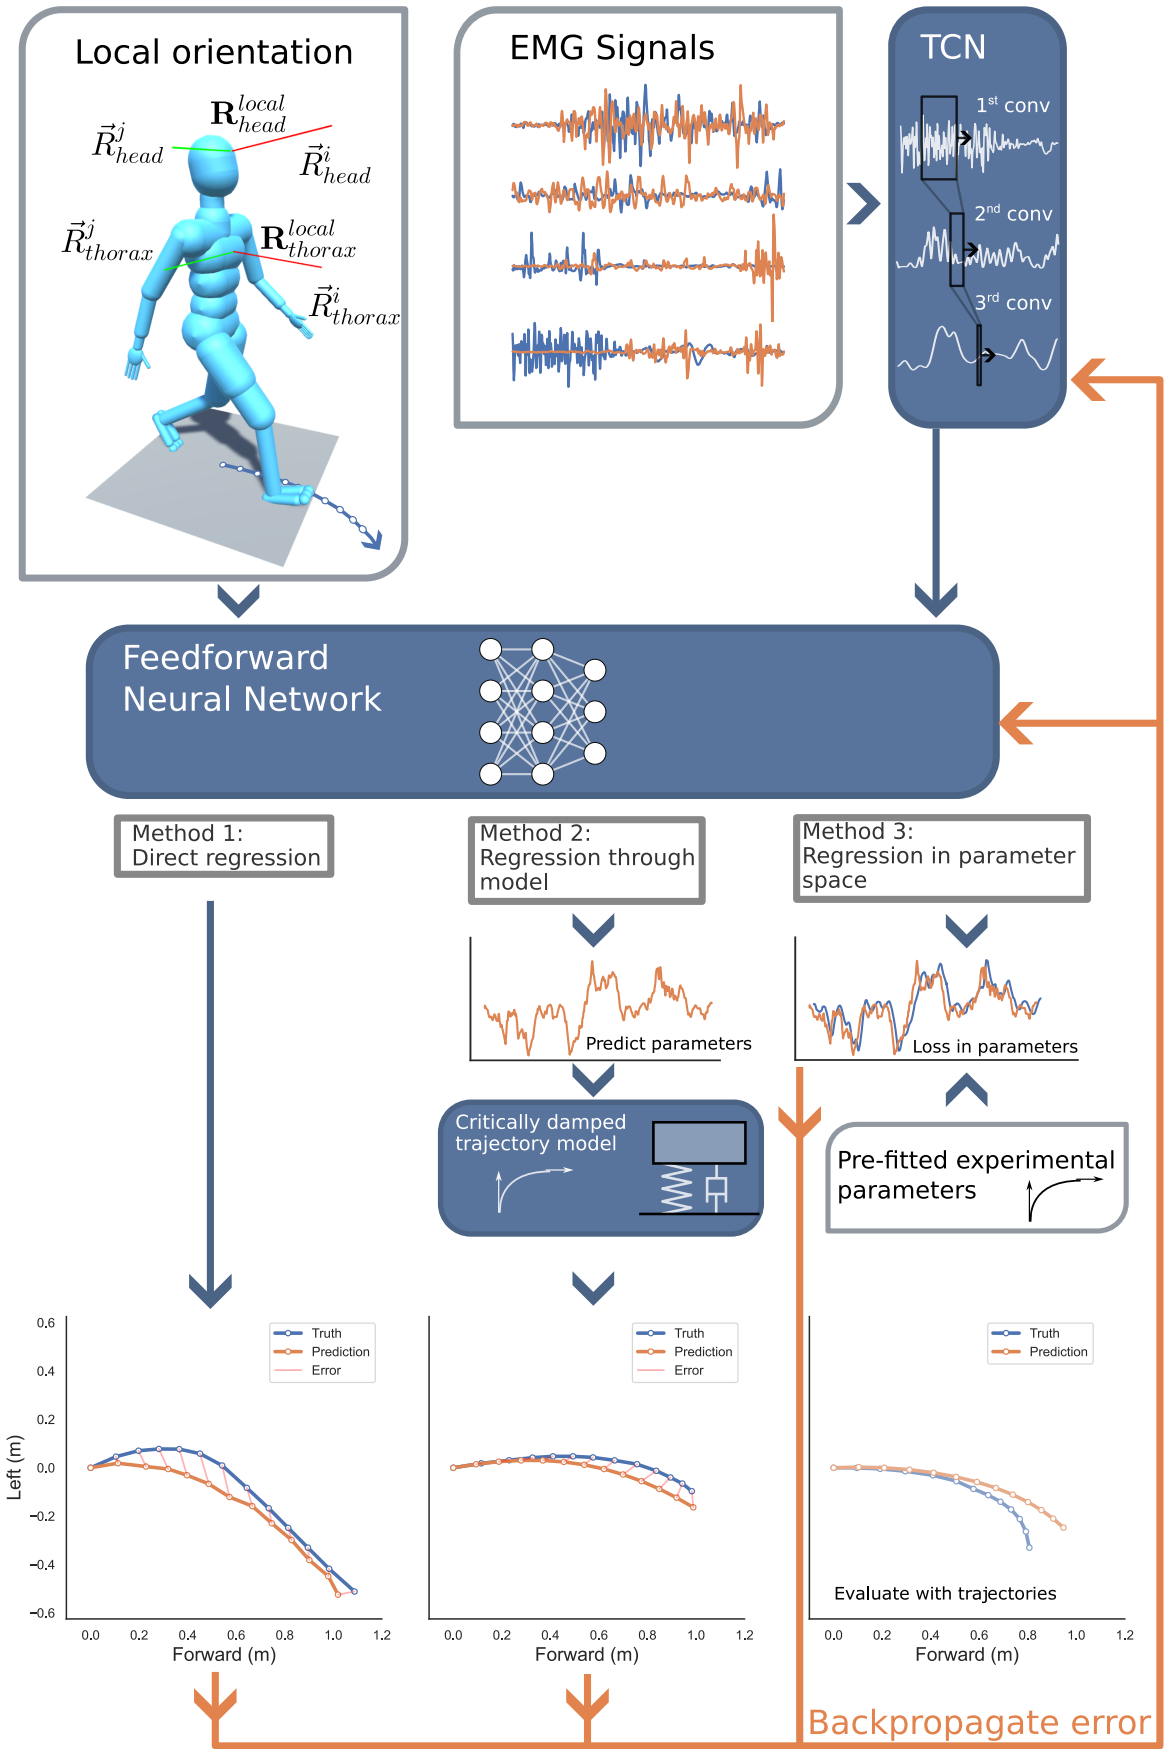


Figure 4.10: Three variant models to regress biosignals to walking path trajectories. All three may include either or both posture and EMG signals. We use an illustration of a virtual humanoid, but these models are applied on experimentally recorded data.

Method 2: The output of the dense network are interpreted as the six parameters of the critically damped trajectory model 4.7, which are then expanded to a full trajectory through a custom gradient-propagating layer with a single learnable parameter (the eigenvalue of 4.7). The shape of the output is therefore constrained to be fully smooth and the resolution of the inference may be changed after training. The losses are calculated in the trajectory and backpropagated throughout the network. To quantify losses we sample the analytical trajectory at the same resolution as Method 1.

Method 3: The output of the dense network are again interpreted as the damped model parameters. They are not processed further. Instead, as a preliminary step, the "ground truth" parameters are determined as in a preprocessing step 4.2.4, and the losses are determined in parameter space instead. The parameters may still be converted to trajectories for evaluation purposes, but not during learning.

For all three network conditions, we will quantify results with three input conditions: sEMG only, body posture only and multimodal inputs. We perform 5-fold cross-validation on the "Angular Turns" dataset, as well as a condition where we train on all motion trial types and evaluate on validation "Angular Turns" files, also with 5-fold cross-validation.

4.3 Results

4.3.1 Eigenvalue extraction

The eigenvalues extracted through bilevel optimisation are reported in Table 4.2. Since the damped trajectory layer's eigenvalue is not fixed and is updated through propagated gradient descent, it automatically learns a λ estimate as well, attracted to a local optimum for all trajectory files (instead of a single one used for the bilevel optimisation).

Table 4.2: Model-in-the-loop λ learned in 5-fold cross-validation (Method 2) and lowest cost λ found in bilevel optimisation (Method 3).

Method	Subject 1	Subject 2	Subject 3	Subject 4	Subject 5	Subject 6
Method 2	-3.86 ± 0.36	-3.64 ± 0.38	-3.68 ± 0.46	-4.39 ± 0.31	-3.52 ± 0.30	-3.74 ± 0.62
Method 3	-2.39	-2.97	-2.60	-2.72	-2.06	-2.95

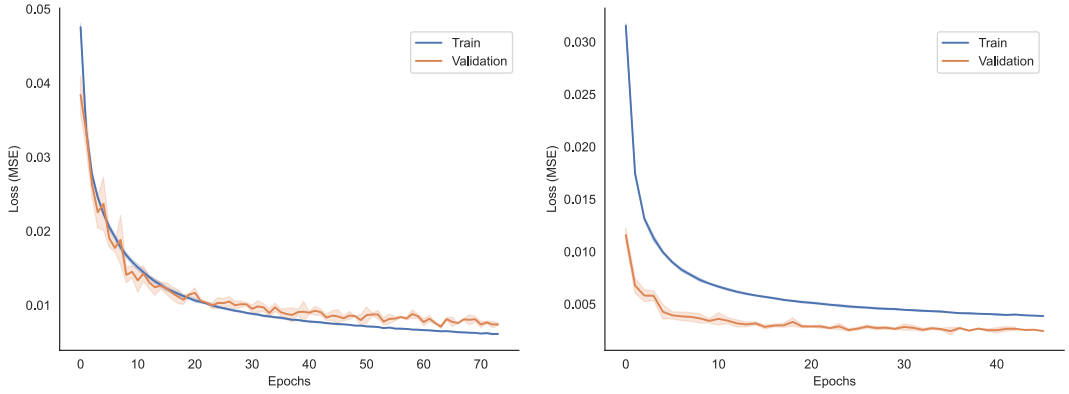


Figure 4.11: Training and validation learning curve examples for subjects 1 and 6. One standard deviation area from 5-fold cross-validation shown around curves.

4.3.2 Walking path estimator

In addition to reporting the RMSE of the overall high-dimensional trajectory P^L, we also present the coefficient of determination. Lastly we also provide the coefficient of determination separated for each time point along the discretised trajectory. This can highlight how the reliability of the prediction changes as we forecast increasingly farther into the future, normalised for the amount of variability at each time point. On Figure 4.11 we show an example learning curve for one of the subjects, demonstrating the regularisation effect of the dropout layers in our network keeping the train and validation performance close. We observe the importance of the training dataset size: Subject 6 had twice the amount of trials recorded compared to Subject 1, and the same amount of regularisation leads to less of an overfit on the training data (ratio of training and validation performance).

Figure 4.12 shows an example time series of the trajectory prediction on validation data. We

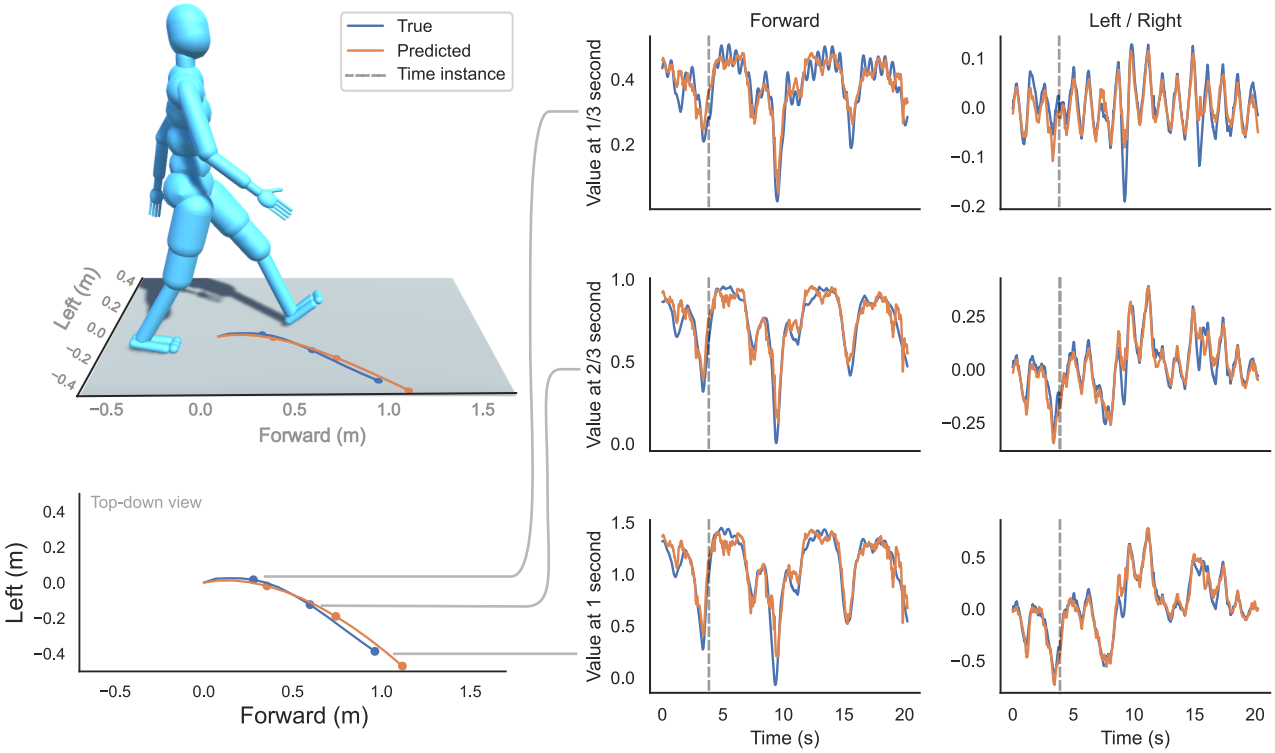


Figure 4.12: Time series visualisation of the inference of the trajectory, plotted separately for the local forward and left/right directions at 1/3, 2/3 and 1s along the future trajectory horizon. A dashed line indicates on the time series results the time instance from which the trajectory on the left is shown. A virtual humanoid is shown as an illustration in the top left to aid interpretation, but these results are experimental with human subjects. This example is for subject 5 on multimodal validation data using the curve model, specialized for the angular-turns dataset. Errors in the prediction are correlated across trajectory points and depend on the type of motion performed (higher during turning). The inference also exhibits low-level constant biases, such as underestimating the magnitude of lateral motion.

find good tracking along both local spatial axes.

Quantified results comparing our conditions and input modalities are shown on Figure 4.13. In particular, separating the r^2 metric along the trajectory leads to some insight between the effects of choices in the regression method and input modalities. Due to the small variance of early trajectory samples, even small mistakes can lead to a low r^2 score, with the inverse being true for the late trajectory. This is why the performance does not necessarily as we progress further in the trajectory decrease when quantified with r^2. Method 3 favours predictions earlier along the curve. This can be explained by Methods 1 and 2 being incentivised for long-term accuracy due to the magnitudes of the errors being greater at the end of the trajectory, whereas in Method 3 the error in the parameters is not accumulated along the trajectory. EMG modalities favour

short term prediction, while the body posture contains more information on the long-term evolution of the walking path. Their combination gives good tracking in the beginning and maintains high performance.

Table 4.3: RMSE (m) of the 5-fold cross-validation results shown in Figure 4.13.

Input	Method 1	Method 2	Method 3
EMG only	0.109 ± 0.021	0.110 ± 0.021	0.115 ± 0.022
Pose only	0.108 ± 0.020	0.109 ± 0.021	0.113 ± 0.020
EMG+Pose	0.067 ± 0.016	0.068 ± 0.014	0.084 ± 0.023

Table 4.4: r^2 of the 5-fold cross-validation results shown in Figure 4.13.

Input	Method 1	Method 2	Method 3
EMG only	0.719 ± 0.077	0.714 ± 0.076	0.685 ± 0.081
Pose only	0.723 ± 0.080	0.718 ± 0.084	0.699 ± 0.069
EMG+Pose	0.893 ± 0.035	0.890 ± 0.037	0.834 ± 0.061

4.4 Discussion

Extending the training dataset and adjusting the sEMG electrode locations made forecasting the walking path trajectory with causal processing feasible. In this section we'll discuss details of the network architecture, implications of the results and limitations of our approach.

4.4.1 Temporal Convolutional Network

A key detail of the TCN network that has been altered since Chapter 3 is that now multiple input bipolar channels are used simultaneously. However, unlike using multiple inputs from e.g., a HD-EMG grid of electrodes, these channels do not have a specific spatially correlated layout. As such, despite the input data forming 2-D arrays, there is no motivation to use

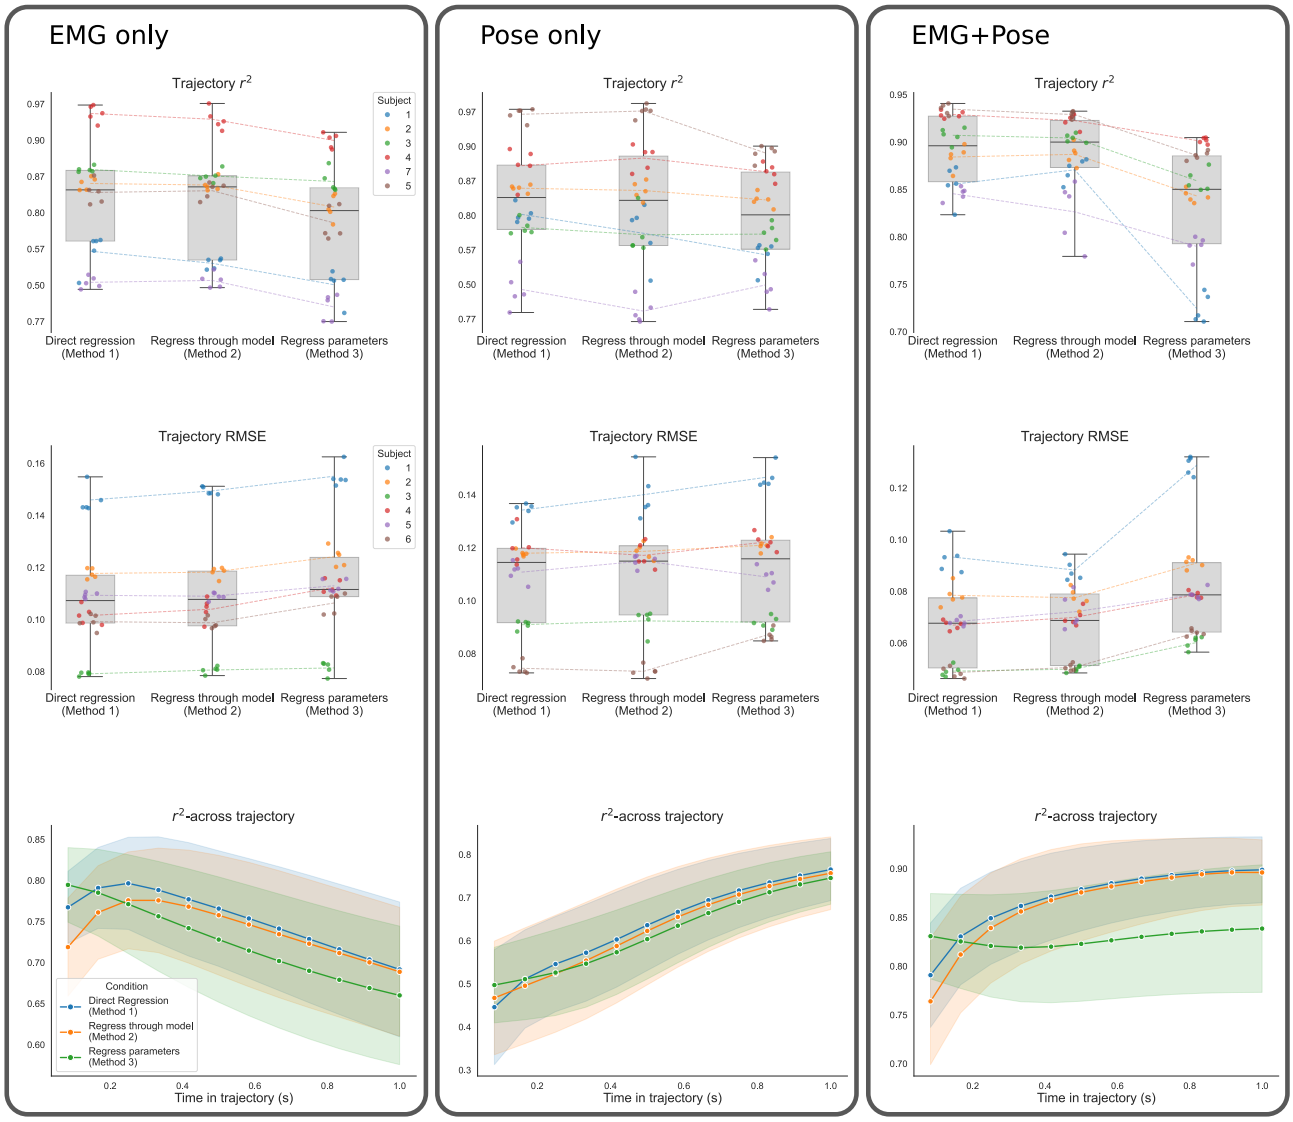


Figure 4.13: Results comparing 5-fold cross-validation performance for the three Methods and three input conditions (organised per column), trained on all motion trial types and evaluated with reserved validation data from the "Angular turns" trials. We show all folds in each category, folds from a given subject labelled via colour, as well lines connecting the mean performance per subject to highlight cross-subject trends. *Top row:* Overall r^2 of the entire trajectory per subject. Estimating the mean trajectory from the dataset (which is not zero due to the non-holonomic nature of walking) would yield an r^2 of zero. *Middle row:* RMSE across the entire trajectory, with units of m. This measure is sensitive to the walking speed of the subject. *Bottom row:* Mean r^2 performance for the three Methods and three input conditions, quantified for the samples along the trajectory separately. Two σ from the 5-fold cross-validation of all subjects shown around each curve. It should be noted that most of the σ arises from inter-subject variability.

2-D convolutions. We could have used 1-D convolutional filters spanning all the channels and sweeping across the time axis. However, we found during experimentation with the pilot dataset that it is more appropriate to use channel-wise separate, single-channel 1-D filters in our convolutional layers. This way the processing of each bipolar electrode is kept separate in the TCN, before mixing them later on in the dense network. A possible explanation to why this improves performance in this data set is the large variability of sEMG time series. Any single sample from one channel has low additional information about the simultaneous sample from other channels. Hence larger scale temporal features are first learned and extracted before mixing channels.

The three methods investigated perform comparably in terms of our metrics. Having access to the trajectory parameters increases the interpretability of these models, making them less of a black-box, and giving us the chance to use the low-dimensional parameter space embedding for other uses (e.g. sending the current and desired velocity estimates to a wearable robotic device to indicate the user's intent instead of the full trajectory). Using the critically damped layer in the training process could be considered a form of very simple physics-informed ML, and captures our goal of explaining experimental data using our chosen representation in a way that is unified with the rest of the learning pipeline. It also avoids the computation heavy bilevel optimisation preprocessing steps required for Method 3.

There are further architectural choices that could be considered that satisfy desired properties in the output trajectory. Additional de-convolutional (or transposed convolutional) layers can follow the output of a dense network, instead of a trajectory model. These can upsample the dense network's activity, producing a smooth autocorrelated trajectory with high frequency samples at reduced computational costs, when compared to simply having a large dense output layer.

A benefit of the non-recurrent nature of the model is that it lacks the ability of learning the trials by rote, which could negatively impact a model's true responsiveness to user input. However, when appropriately controlled for, recurrent structures such as long-short term memory (LSTM) blocks could help stabilise the prediction and improve performance. We experimented with

attention mechanisms and transformer layers as an extension of our TCN, but saw no benefit in terms of performance, coming at additional computational costs. Nonetheless, we do not rule out their use for high-level locomotion intent estimation in the future. Similarly, comparisons with non-NN estimators should also be performed, such as Kalman filters or Gaussian Process Regressors. By the subject consciously shifting their CoM and body posture, volitional control schemes may also be established (Tucker et al. 2015). After appropriate user training, a method like that could also be successful, at the cost of additional cognitive load on the subject.

In our implementation of 5-fold cross-validation we made sure that splits were made in such a way that no windows in separate folds have any overlapping samples of sEMG. Naively shuffling and splitting windows into 5 parts would lead to information leaks from the validation, We observed an overestimation of r^2 values in preliminary tests for the EMG only model approximately by 0.19 if inappropriate folds were used.

We observe that in early points of the trajectory, there are large oscillations in the lateral direction, caused by the pelvis' back-and-forth sway during stepping (Figure 4.12). The intent estimators' ability to track this motion is more pronounced when sEMG data is available (early trajectory performance in Figure 4.13). We believe this is thanks to inclusion of the leg electrodes, the activity of which are heavily correlated with the timing of steps and therefore the gait phase (Figure 4.5). While our intent estimator explicitly predicts outcome variables related to the HL locomotion intent, the trained networks also implicitly estimate lower level concepts such as gait events and gait phase. By separating the r^2 metric along the forward (anterior-posterior) and left-right (mediolateral) directions, we can observe that including pose features helps estimating turning direction in the late trajectory (Table 4.5).

The λ values for the damped trajectory model learned with Method 2 were increased in magnitude for all subjects compared to the values determined through least-squares optimisation. A key difference is that in our bilevel optimisation we selected a single "Angular Turns" trial to determine a suitable λ, whereas in Method 2 this parameter is learned in the gradient-descent on all training data and trial types. Since in Figures 4.6-4.8 we have seen the influence a fixed λ has on the optimal parameters, we examined the parameter output of the networks trained

Table 4.5: Separating the r^2 along the anterior-posterior (AP) and mediolateral (ML) axes and calculating their ratio when using Method 2, for the EMG-only and the pose-only models. Mean values and one standard deviation reported. These results indicate contributions to inferring turning direction from the two input modalities.

Input	Whole trajectory r^2_{AP}/r^2_{ML}	Early trajectory r^2_{AP}/r^2_{ML}
EMG-only	1.492 ± 0.599	1.453 ± 0.452
Pose-only	1.034 ± 0.224	1.639 ± 0.734

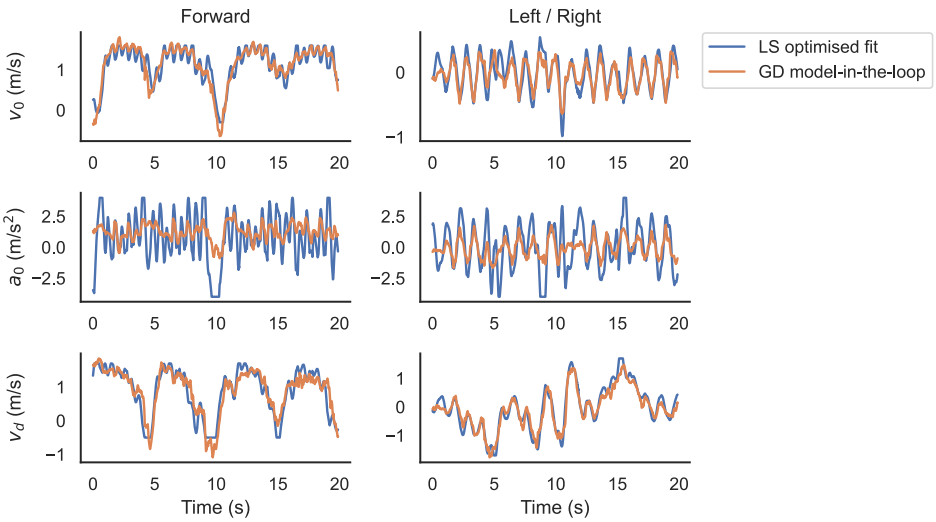


Figure 4.14: The activations of the second to last layer of the network trained in Method 2 can be interpreted as the parameters of the damped trajectory model. Here we compare on one of the "Angular Turns" validation trials of Subject 4 the parameters inferred by the network from EMG and posture, compared to directly fitting the parameters to the motion capture trajectory using least-squares optimisation.

with Method 2. Figure 4.14 shows this for one of the subjects. Even though higher absolute λ values were expected to lead to greater magnitude acceleration estimated (Figure 4.6), we observe much more attenuated estimations of a_0. However, both v_0 and v_d time series are comparable to the values optimised directly on the ground truth trajectories.

4.4.2 Limitations

The limitations of ML EMG regressors should be reiterated. Since the data set was recorded with one subject in one session, the model is expected to overfit to the specific subject and their

momentary level of tiredness/sweating, the particular electrode placement and orientation and the noise environment of the laboratory. These limitations are challenging to fully characterise and mitigate, due to the requirement to extensively expand the data set, but approaches like the one we demonstrate in Chapter 3 could be pursued for this task as well. There is also potential to overfit to specific gait tasks included in the training set (for an example with the pilot dataset, see appendix Figure B.2). We attempted to mitigate this by incorporating a range of different turning behaviours, but the performance drop on the freeform task, when not included in the training, indicates that this is a significant concern. Lastly, the deep learning approach does not discriminate between sources of information in the signal. It is possible (and if reliably beneficial, expected) that it learns to exploit noise signals, digitisation artefacts, or side-effects of preprocessing steps to increase its accuracy on the data set. This can impact the accuracy when these unintentional signals are missing or changed. Band-pass filtering the signal is expected to remove some of these misleading effects, but these concerns can only be addressed by trialling the system online, with the real-time signals of a user and outside lab environments. We causally band-pass filtered (as in Chapter 3) our sEMG signals, but performed some comparisons with simply normalised sEMG. No significant differences were observed in the performance of the two approaches.

Limitations in the dataset

We used the orientation of the head and torso in the local frame as part of some of our models. This was determined using photogrammetric motion capture, but could have also been extracted using IMU-based or other wearable solutions. However, both the variance and bias of wearable motion capture systems is higher than that of optical ones (Xiang et al. 2022), a matter we have not investigated. As such, future investigation with signals from wearable embedded systems, or sensitivity analysis to additional noise should be performed.

All of the subjects were male in the follow-up study. Due to the differences in the sEMG signal properties between sexes (Pradhan et al. 2020), the present findings need to be repeated with female participants as well.

We relied on sEMG from the lower limb bilaterally. However, with a person with transtibial limb difference, the muscle strength and the signal quality may be compromised. We ran a preliminary test on the EMG only models trained using Method 2 and removing the sEMG signals from the right leg. We observed a drop in performance from 0.714 ± 0.076 to 0.681 ± 0.081 in the r^2 metric. In practice, prosthesis users do have meaningful sEMG activity on their affected side (Chen et al. 2014; Kannape et al. 2014), however, the suitability of collecting lower leg muscle signals for walking path trajectory estimation of users with lower limb difference should be investigated in the future.

4.4.3 Assumption of nonholonomy

In our representation of HL locomotion intent, we quantified the behaviour in terms of current and desired horizontal walking velocity, from the current facing orientation of pelvis. We included no variable to specify future facing orientation throughout the trajectory. Essentially, we assume that where you face with your body will be closely tied to which direction you move, which is most often the case in generic locomotion settings (Mombaur et al. 2010). This is sometimes described as a nonholonomic representation of locomotion (as discussed by (Mombaur et al. 2010)). Nonholonomy implies that the series of state changes is needed to be able to fully deduce the system configuration (since we do not specify facing direction, we can only recover it by keeping track of past trajectories). However, more nuanced motion behaviours such as sidesteps, strafing or walking backwards require a holonomic state format, where e.g., each point in the trajectory is paired with a facing direction unit vector (Zadziuk 2016). The motion trials we recorded were chosen to approximately satisfy the assumptions of our simplified nonholonomic representation.

To challenge our own assumptions, and highlight potential limitations, we used the ground truth motion trajectories from one of our trapezoidal path trials (see Appendix C) in a motion matching framework (see Appendix A). We substituted the missing facing direction needed in the holonomic motion matching process by the normalised velocity along the trajectory. This way we can examine the full body motion as reconstructed from our encoded format. Figure

4.15 shows the resulting walking path.

The integration was tested on a recording of two clockwise turns, followed by a 180° turn, then three counterclockwise turns, along a trapezoidal path (see Appendix C). The following observations can be made:

- Due to the open-loop nature of the motion synthesis and the distributional shift between the motion data set of the two systems, errors add up and the pathway diverges and does not return to the starting point.
- The motion synthesizer correctly identifies turn directions. Magnitude of turn angle could also be retained.
- Notable exception is seen at the 180° turn. In natural locomotion we begin turning while still moving in the original direction, which is missed in our representation, leading to an underestimation of turn angle.

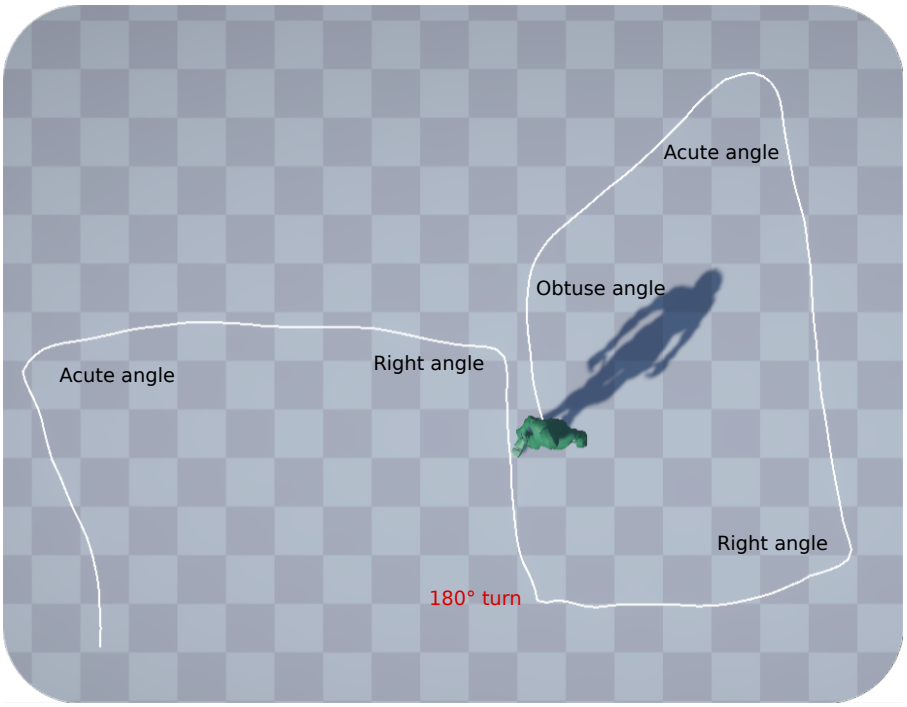


Figure 4.15: Resulting path from synthesised motion from trajectory horizon from recorded motion capture, i.e. the ground truth of the intent estimator. The 180° turn is incorrectly generated as coming to a stop, then a 90° counterclockwise turn.

The fact that our intent representation does not full reconstruct the full body motion does not imply that it is unsuitable for HL control. However, if in future studies nuanced locomotion

patterns like sidestepping during P&O use are also a priority, holonomic locomotion intent spaces may be a more appropriate choice.

4.4.4 Intent synthesis in simulation

One of the main reason hierarchical control strategies are appreciated is their modularity. Our intent representation, the horizontal walking path trajectory, is sufficiently abstract to be able to be easily synthesised when evaluating other parts of the control scheme. In simulation we can use Equations 4.6 and 4.7 to generate smoothly changing walking path trajectories as determined by a free variable $\mathbf{v}_d$. The remaining parameters $\mathbf{v}_0$ and $\mathbf{a}_0$ can be updated analytically using the critically damped system (and not by tracking the virtual kinematics via finite differences, which can lead to numerical instabilities). These parameters and the forecasted walking path associated with them can be used to drive motion synthesis (e.g. via motion matching (Clavet 2016)), or be used to control intent driven virtual assistive robotics. In the context of full body motion synthesis, the trajectories can be driven by sampling target directions and speeds of walking (Figure 4.16) or target locations that the humanoid needs to reach (Figure 4.17).

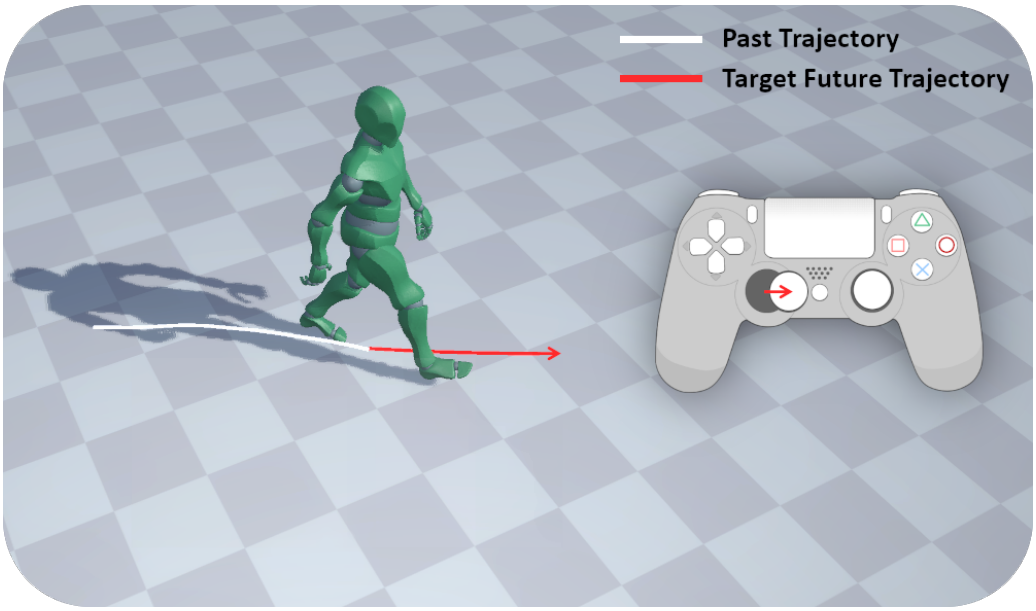


Figure 4.16: Target velocity as high-level control driven trajectory generation. Could be conceptualised as 2-D analogue joystick interface.

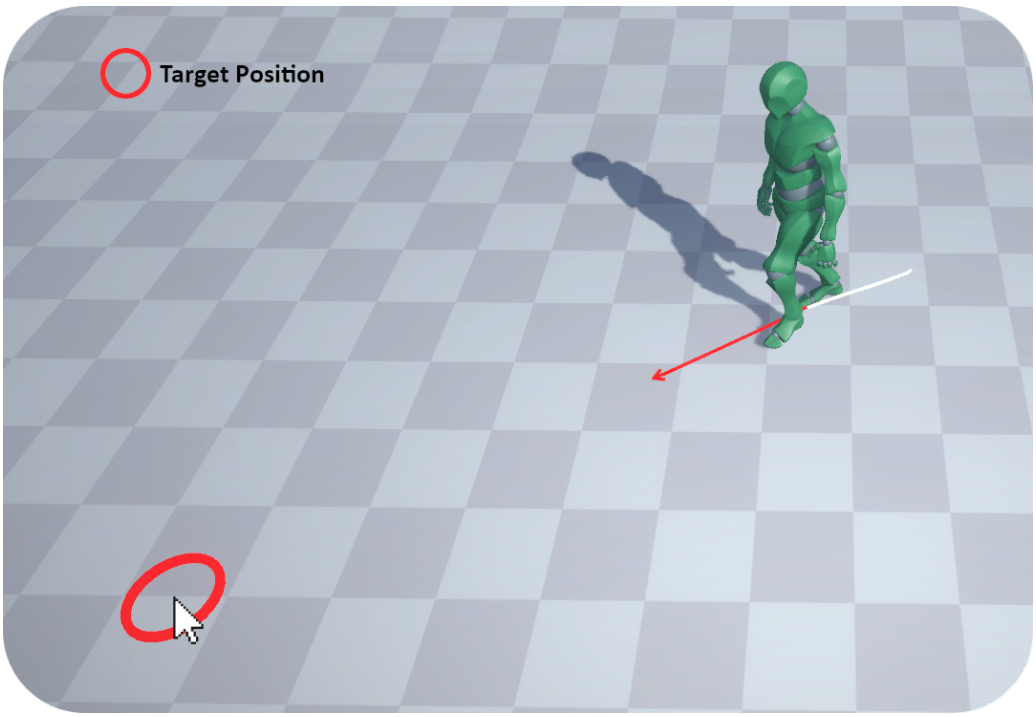


Figure 4.17: Target position as high-level control, as determined by the intersection of a ray cast from the camera based on the cursor's position and the ground plane. Target velocity can then be oriented towards target position, its magnitude is clamped to a configurable maximum, otherwise it equals the velocity that would reduce the distance to zero in one second.

4.5 Conclusion

In this study we collected a novel kinematic and electromyographic motion dataset with diverse non-steady state locomotion involving various turning methods. We demonstrated that TCN architectures can process bipolar trunk and lower limb sEMG and estimated turning direction and walking velocities in a continuous space representing HL locomotion intent. By introducing a critically damped trajectory model into the estimator, we made it more interpretable in terms of gait parameters like the current and desired horizontal walking velocity. The presented results indicate that estimating the walking path trajectory is feasible from EMG and body posture, and may be a suitable method to perform HL intent estimation in hierarchical control strategies of powered devices assisting the function of the lower limb.

Acknowledgments

We would like to thank Professor Allison McGregor for providing access to the Biodynamics Lab and its equipment for our motion capture experiments, which were supported by Dr Matthew Banger as technician.

Chapter 5

Locomotion Synthesis for Simulated Prosthesis Use*

Virtual environments provide a safe and accessible way to test innovative technologies for controlling wearable robotic devices. However, to simulate devices that support walking, such as powered prosthetic legs, it is not enough to model the hardware without its user. Predictive locomotion synthesizers can generate the movements of a virtual user, with whom the simulated device can be trained or evaluated. We implemented a Deep Reinforcement Learning based motion controller in the MuJoCo physics engine, where autonomy over the humanoid model was shared between the simulated user and the control policy of an active prosthesis. Despite not optimising the controller to match experimental dynamics, realistic torque profiles and ground reaction force curves were produced by the agent. A data-driven and continuous representation of user intent was used to simulate a Human Machine Interface that controlled a transtibial prosthesis in a non-steady state walking setting. The continuous intent representation was shown to mitigate the need for compensatory gait patterns from their virtual users and halve the rate of tripping. Co-adaptation was identified as a potential challenge for training human-in-the-loop prosthesis control policies. The proposed framework outlines a way to explore the

*This chapter is a reproduction of the work published as "Balint K Hodossy and Dario Farina (2023). *Shared Autonomy Locomotion Synthesis With a Virtual Powered Prosthetic Ankle.* In: IEEE Transactions on Neural Systems and Rehabilitation Engineering 31, pp. 4738–4748"

complex design space of robot-assisted gait, promoting the transfer of the next generation of intent driven controllers from the lab to real-life scenarios.

5.1 Introduction

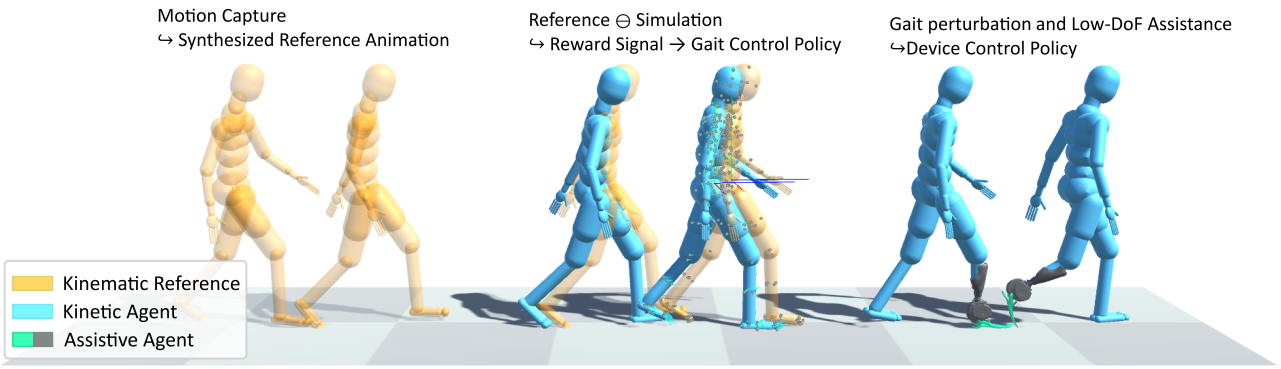


Figure 5.1: Developing a virtual lower limb device testbed. A predictive walking policy was learnt, incorporating information from experimental data-driven references. A gait pathology was then simulated. Addressing the induced biomechanical deficit was the goal of the device's agent. Additional conditions and tasks can be introduced to assess the versatility of the controllers. Prosthesis mesh used from the Open Source Limb project (Azocar et al. 2020).

Commercially available powered lower limb P&O have yet to realise their potential impact (Gehlhar et al. 2023; Tucker et al. 2015). Price, weight and user perception are all key factors in determining whether these devices are a worthwhile intervention for their users (Hill et al. 2017).

Moreover, appropriate strategies are needed to synchronise the upcoming generated motion with the user's motor intent. Many wearable robotic devices are limited by unintuitive control interfaces, low number of DoFs control and reliance on signals acquired as a result of a movement already in progress (Baud et al. 2021; Tucker et al. 2015). This is especially true for non-steady-state locomotion (e.g. turning, changing walking speed or starting stair ascension) and for partial assistance systems (Tucker et al. 2015). The design, development and evaluation of novel control strategies are hindered by limited access to hardware and participants, as well as by the risks inherent to testing incomplete device controllers. The physical test-beds and frameworks, which can be used for validated, reproducible and safe tests require complex and unique equipment (Torricelli et al. 2018). Time-efficient iterations on hardware and controller

designs can be facilitated by device emulation hardware (Caputo et al. 2014). However, these may introduce restrictions on the range of test environments and locomotion tasks due to the mobility constraints of the emulation platform.

In the context of upper limb device design, virtual environments are a well established approach to alleviate the aforementioned issues (Armiger et al. 2011). However, there are additional challenges when this approach is applied to locomotion tasks. To provide the kinematic and kinetic context that is necessary for the operation of a simulated lower limb device, the user's movements need to be synthesised as well. It is insufficient to solely use inverse dynamics for this purpose. Indeed, motion trajectories reconstructed from experiments quickly become inaccurate with the occurrence of forces from a virtual device, which leads to instability if deviations from the prerecorded states are permitted. Instead, predictive forward simulations can generate stable walking policies, modelling key aspects of the simulated user's motor control (De Groote et al. 2021), for example, the ability to:

- React and recover from disturbances and to take advantage of the assistance from the wearable device.
- Generate movement conditioned on a modelled gait pathology, and produce or learn compensatory movements.
- Perform long and short-term motion planning during non-steady-state locomotion tasks.

There are various methods for constructing gait policies with these characteristics, such as heuristic reflex-rule-based systems (A. R. Wu et al. 2017), trajectory optimisation (Falisse et al. 2019; Tassa et al. 2012), evolutionary strategies (Conti et al. 2018), supervised learning (Fussell et al. 2021; Merel et al. 2018) and reinforcement learning (Peng et al. 2021; Bergamin et al. 2019; Lee et al. 2019). Summaries on related work are available from the neuromechanical (De Groote et al. 2021; Song et al. 2021) and computer graphics perspectives (Geijtenbeek et al. 2012; Mourot et al. 2022). deep reinforcement learning (DRL) in particular has led to solutions that generalise well to multiple locomotion tasks simultaneously with realistic motion (Peng et

al. 2021), while also reproducing key biomechanical aspects of gait (Lee et al. 2019). In contrast to model predictive control methods, it requires an additional learning phase. However, once trained, DRL can perform inference using less resources than collocation methods. Furthermore, there have been promising examples of sim2real capability using DRL (Akkaya et al. 2019; Höfer et al. 2021; Peng et al. 2020), an essential property when simulating robotics (C. K. Liu et al. 2021). If a kinematic reference animation is available, learning time can be significantly reduced by using a motion tracking approach (Peng et al. 2021; Bergamin et al. 2019; Lee et al. 2019). Musculoskeletal actuation may be essential when investigating orthotic devices (Farris et al. 2014). However, direct torque actuation is multiple times faster to train in comparison (Lee et al. 2019), and may be appropriate for simulating users of prosthetic systems, which do not directly control the same DoFs as the device.

Simulating human locomotion along with the behaviour of active assistive devices is a promising way to introduce an inner design loop to the development of controllers. Signal modalities, model hyperparameters and calibration algorithms can be first investigated this way, before applying the gained insights to the real-life robot-human system. One benefit is the potential to use access to hardware, end-user subjects and testing equipment more efficiently. Alternatively, parameter ranges determined in simulations could be used as starting points when fitting models to real users through adaptive (J. Zhang et al. 2017) or manual methods (Simon et al. 2014). Finally, the device calibration and parameter tuning process, which is one of the most challenging aspects of current powered P&O (Gehlhar et al. 2023), can be improved with insights from simulation.

Simulated gait policies have been previously proposed as suitable test environments for lower limb assistive devices (J. Park et al. 2022; De Groote et al. 2021), and have been applied to orthotic (Lim et al. 2017; J. I. Han et al. 2022) and prosthetic (Falisse et al. 2019; Lee et al. 2019; De Vree et al. 2021) systems. However, existing examples primarily use passive devices during mostly steady-state locomotion tasks. Here, we explored a model of a PD-controlled unilateral powered ankle prosthesis, while its simulated user walked on a level surface with frequent turns and stops. Motion tracking DRL gait policies were learned based on reference kinematic animations generated with motion-matching (Bergamin et al. 2019; Clavet 2016),

which provided a data-driven representation not only of the user's desired movements, but also of their high-level abstract intent. Following this phase, agency over the user's below knee control signals on one side was assigned to a second control policy to mimic a transtibial prosthesis. The device control policy was also trained with DRL. We compared a prosthesis controller that only had inputs from implicit sensors (Tucker et al. 2015; Tschiedel et al. 2020) with one that additionally received a representation of the simulated user's locomotion intent. This low-dimensional abstract intent serves as a surrogate signal of a neural interface in a late-fusion setting. In addition to reinforcement learning-based controllers tuned automatically using proximal policy optimisation (PPO), we also reimplemented a FSM device controller (Simon et al. 2014). The gait phase estimated by this rule-based system progressed naturally, indicating that the kinematic and kinetic context of the virtual device is plausible. The parameters of the FSM were tuned manually through a slider interface. An overview of the main stages of constructing the device controller testbed are illustrated in Figure 5.1.

Main contributions

In summary, the following are the main contributions of this paper:

- Modelling of prosthesis use in a non-steady-state locomotion scenario involving stops and turns, with a virtual user that can react to perturbations or intent changes.
- Implementation of a motion tracking gait controller from the field of character animation, and demonstration that it generates a plausible dynamic context for a transtibial prosthesis.
- Modelling the co-adaptation setting where both the user policy and device controller could be learned through DRL simultaneously, but as separate agents.
- Demonstration that the desired horizontal walking velocity is a suitable control signal for a prosthetic ankle and reduces the need for compensatory gait patterns, compared to an ankle without inputs controlled by the user.

5.2 Materials and Methods

The following section details the key subsystems used to simulate the dynamic human-prosthesis model. First, the motion-matching approach for generating kinematic reference motions is explained. Second, the DRL environment for learning gait without a prosthesis is presented, followed by the changes introduced to model prosthesis use.

5.2.1 Locomotion intent synthesis

The high-level intent driving the motion synthesis was characterised as the desired horizontal two-dimensional velocity vector of the pelvis, ranging in amplitude from 0 to a moderate speed of 1.4 $\frac{\text{m}}{\text{s}}$ (Murtagh et al. 2021). The process for generating the desired walking velocity during gait policy learning will be detailed in Section 5.2.2.

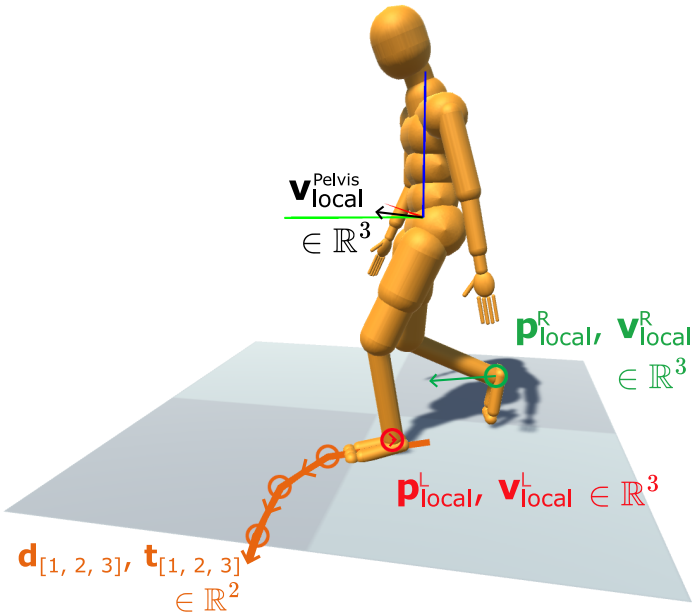


Figure 5.2: Illustration of the features used in the nearest neighbour search for the motion-matching system, and the local orientation frame they were quantified in. The orange trajectory consists of the position and normalised velocity of the pelvis center of mass, projected on the horizontal plane at 1/3, 2/3 and 1 1 s in the future. This curve represents the upcoming walking path in the motion capture data set.

To convert this abstract locomotive goal to specific mid-level motion trajectories, the high-

level intent was used to extrapolate a 1 s long, critically damped walking path, which was used to drive a motion-matching system (Clavet 2016). Motion-matching allows efficient usage of limited motion capture data sets by generating transitions between disparate parts of them. These transitions are key in augmenting the training data to include a diverse range of non-steady-state walking including turns and stops. The motion capture recordings of (Harvey et al. 2020) were used as the data set in this study. They consist of unsegmented clips of diverse locomotion. A "meta-data" feature vector was calculated for each frame of these recordings, as described in (Bergamin et al. 2019). This vector consists of the following elements (illustrated in Figure 5.2):

- Velocity of the pelvis $\left(\mathbf{v}_{local}^{Pelvis} \in \mathbb{R}^3\right)$.
- Position and velocity of the feet $\left(\mathbf{p}_{local}^{[L,R]}, \mathbf{v}_{local}^{[L,R]} \in \mathbb{R}^3\right)$.
- Position and normalised velocity (walking direction) of the pelvis projected to the horizontal plane at 1/3, 2/3 and 1 second in the future $\left(\mathbf{t}_{[1,2,3]}, \mathbf{d}_{[1,2,3]} \in \mathbb{R}^2\right)$. These are referred to as the trajectory and direction features respectively.

This gives a total of 27 dimensions for motion-matching, which were all normalised to zero mean and unit standard deviation. These features were described from a semi-local frame of reference, which was located at the pelvis, with one axis aligned with the global vertical and another with pelvis forward directions and was assigned zero global velocity relative to the ground. Therefore, only the directions of velocities described in this local frame were influenced by pelvis kinematics, but not their magnitudes. Motion-matching velocity features were estimated using a first order Savitzky-Golay filter. To adjust the relative importance of the elements in this feature vector they can be scaled by a set of weights. These weights determine the trade-off between responsiveness to matching the desired trajectory and the smoothness of motion. A factor of 6 was used for velocity, 3 for position, 4 for trajectory and 2 for direction features, set through manual tuning.

During every 10^{th} frame of motion synthesis, the meta-data vector of the current motion frame

was collected. Then, its walking trajectory and direction features were replaced by the critically damped walking path, as determined by the artificial high-level intent. This modified vector was then compared with all other meta-data vectors in the data set, and the kinematic motion continued from the frame that corresponded to the nearest neighbour of the current feature vector. Due to the inclusion of a single locomotion style and task (level ground walking with turns and stops), the data set was small enough that parallelisation or K-D tree-based methods yielded no performance gains over a linear search for the nearest neighbour match (Clavet 2016). "Inertialization" was applied on the kinematic animation targets to smooth discontinuities (Bollo 2016), a blending technique inspired by zero-jerk trajectory control principles (Flash et al. 1985). The walking generated by motion-matching is kinematic in the sense that it cannot describe interactions between the humanoid, the ground and other elements in the scene (such as a virtual prosthesis). Indeed, tracking the output of motion-matching with a pure PD controller will almost immediately lead to tripping and falling. However, this movement provides a plausible first guess for target poses of a dynamically simulated humanoid controlled by a PD-DRL hybrid system, described below and illustrated in Figure 5.3. This combination of motion-matching with DRL was first proposed in (Bergamin et al. 2019).

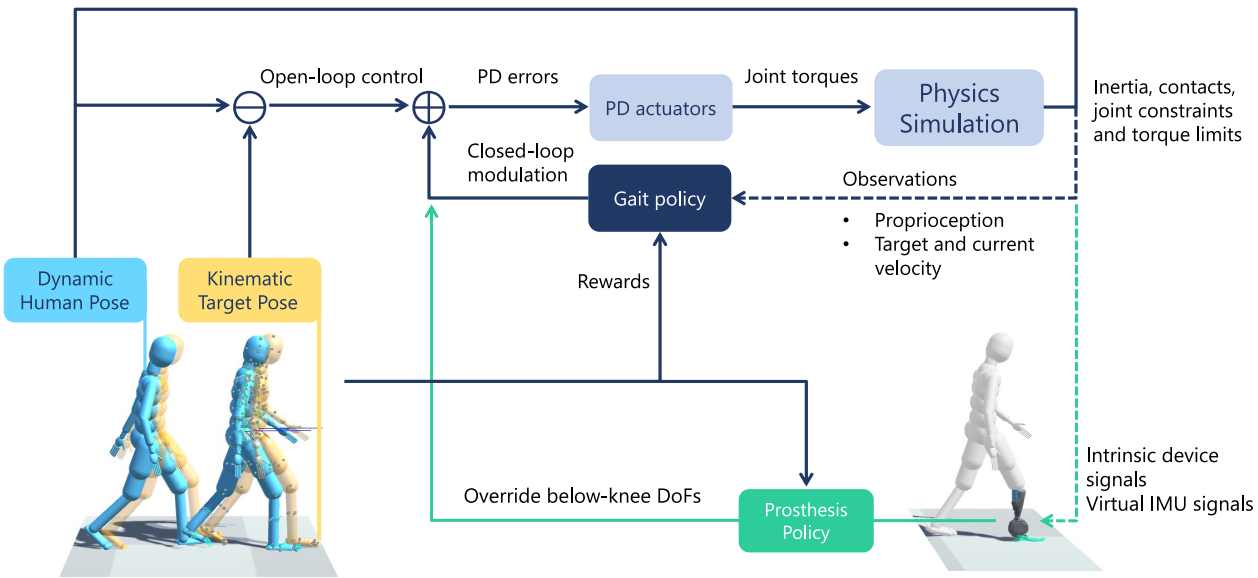


Figure 5.3: The control problem solved using DRL. The target kinematic state from motion-matching is modulated by the gait policy, which receives observations from the simulation state. These targets are then converted into joint torques through stable PD. The simulation and reference states are then updated, and the process repeats. In prosthesis use conditions, below knee DoFs are controlled by a separate prosthesis policy, which receives its observations. Both gait and prosthesis policies are trained using the same reward signal.

5.2.2 Gait policy

The "CMU humanoid" model described in (Tunyasuvunakool et al. 2020) was simulated in the MuJoCo physics engine (Todorov et al. 2012) at 125 Hz. The torso, lower back, neck, shoulder and hip joints were replaced in the model with spherical joints instead of serial hinge joints. A 2 DoFs ankle model was used, and all joints distal from the elbow were removed. This resulted in a total of 43 DoFs. The clavicle body segments, modelled as 2 DoF joints were fully passive, leading to 39 DoFs to be controlled by the gait policy. The agent received no external forces or torques aiding its balance. Scene construction, visualisation and task logic was performed in the Unity engine, using the ML-Agents package for communication with a learning framework implemented in PyTorch (Juliani et al. 2018).

Reinforcement Learning

The gait control problem can be represented as a Markov decision process. Then, the probability of transitioning to a specific physics simulation state at the next time step is wholly determined by the current state ($\mathbf{s}_t$), the system dynamics and the actions ($\mathbf{a}_t$) taken by the decision-making agents in the system. If we don't assume *a priori* what the ideal actions are, we can instead quantify desirable properties for the controlled system with a reward function $r_t(\mathbf{s}_t, \mathbf{s}_{t+1}, \mathbf{a}_t)$. DRL is a family of optimisation algorithms that use past experiences of states, actions and rewards to maximise the weighted sum of rewards gained within an episode of learning. The "deep" part of DRL refers to the use of deep learning function estimators in the agent's behaviour. In PPO (Schulman et al. 2017), this is achieved through the use of an actor and critic system. The critic is iteratively updated to predict the future weighted sum of rewards that will be collected after taking a given action at a given state. The actor maps states to actions through a stochastic policy. The policy ($\pi(\mathbf{a}_t|\mathbf{o}_t)$) takes a subset of the state vector (called observations $\mathbf{o}_t$) as inputs, and outputs the probability distribution for each possible action to take. The parameters of the policy determine how likely it is to select an action conditioned on the state. When an action results in more/less reward than anticipated by the critic for a given observation, these parameters are adjusted by gradient ascent to make

that action more/less likely in the future. PPO was chosen as the learning algorithm due to its robustness with respect to hyperparameter choice, its suitability for continuous action spaces and previous successes in applying it for locomotion (Peng et al. 2018; Bergamin et al. 2019; Lee et al. 2019).

The policy trained by PPO was a feedforward neural network with 3 hidden layers, 512 units each with swish activation (Ramachandran et al. 2017). The same architecture was used for the critic network. A visual graph of the gait policy's role in the simulation is shown in Figure 5.3. Training was performed on 6 parallel environments running on the same system. Further hyperparameters of the learning environment are detailed in the configuration file included in Appendix G.

Observations, actions, rewards

A vector of observations was sampled from the simulation environment at each control step queried at 60 Hz, based on the feature set used in (Bergamin et al. 2019). Inside this observation vector, 6-dimensional kinematic information (position and linear velocity) was collected of the following body segments from the pelvis' coordinate frame:

- Left and right feet (12 dimensions)
- Left and right forearms (12 dimensions)
- Upper back (6 dimensions)
- Head (6 dimensions)

The difference between the desired and actual kinematics of these body segments was also provided as an observation (36 dimensions), as it was previously found to speed up the learning process (Bergamin et al. 2019). This was concatenated with the desired and actual centre of mass velocity, as well as their difference (3×3 dimensions). The high-level walking intent in the form of the desired horizontal walking velocity used in the motion matching process was

provided, along with its difference from the actual horizontal centre of mass velocity (2×2 dimensions). Lastly, the agent's previous actions (described below) were also provided as observations (39 dimensions). In total, the observation vector was 124 dimensional. This feature set (Bergamin et al. 2019) is different from many other implementations due to a lack of reliance on a phase variable (Peng et al. 2018; J. Park et al. 2022; Lee et al. 2021), which is not straightforward to define during continuous but non-steady-state motions. The observation vector was concatenated with the previous decision step's observations before using them as inputs for the policy network.

Given this input, the policy outputs an action vector at each control step. Exponential smoothing was applied to this output with a smoothing factor of 0.9 (Bergamin et al. 2019), and was assumed to be constant between control steps. The actions modulated the open-loop poses of the motion-matched animation. They were not interpreted as velocity targets (Bergamin et al. 2019) or position targets (Peng et al. 2018) from which PD errors were then calculated. Instead, the action vector was added directly to the PD error vector calculated as the difference between the open-loop reference and the current pose. While this is equivalent to pose modulation in the case of hinge joints, it reduces the amount of computation necessary for spherical joints. This is because their error signal has a smaller dimensionality (3D) than their quaternion-based (4D) positional descriptions. Torques generated by Stable PD (Tan et al. 2011) actuators were limited under 190 N m, constraining the output to reasonable values for gait (Winter 1984) and improving simulation stability. A position gain of 510 $\frac{\text{N·m}}{\text{rad}}$ and a velocity gain of 5.2 $\frac{\text{N·m·s}}{\text{rad}}$ was used, roughly approximating the human upper limit of stiffness in the knee and hip (Pfeifer et al. 2014; Shamaei et al. 2013). By modulating the error signal the agent can mimic other stiffness and damping parameters. However, particularly in early stages of learning, a high stiffness starting point was found to be helpful in learning to avoid collapsing. A copy of the simulated humanoid followed the motion-matching animation, enforced through equality constraints on position and orientation. In this way, joint-space state and higher-order kinematics could then be queried from this "puppeteered" character for reference observations and actions.

The reward collected by the agent was calculated based on the kinematic differences between the reference and simulated body, consisting of four terms summed together (Equation 5.1).

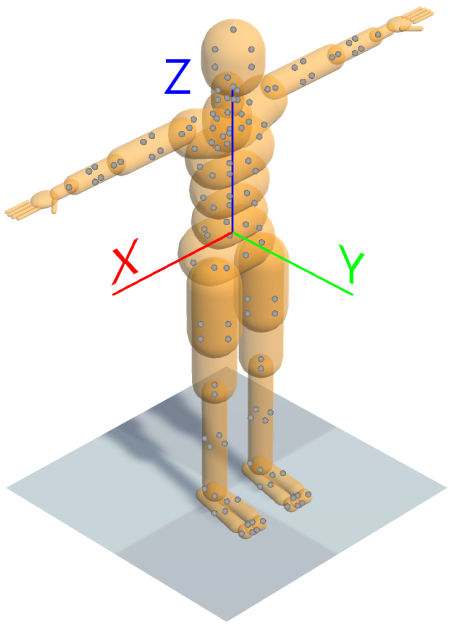


Figure 5.4: The local reference frame and the bounding box surface points used in the reward calculation. Segments distal to the forearm were not included in the reward calculation.

$$r_t = e_{\text{fall}} \left(r_p + r_v + r_{local} + r_{v_{CM}} \right) \tag{5.1}$$

The first two were calculated using the "surface point" method introduced in (Bergamin et al. 2019).

The *position reward* was defined as the distance between the body segments of the dynamic humanoid and its kinematic reference (Equation 5.2).

$$r_p = \exp\left(\frac{-7.3}{N_{\text{segments}}} \sum_{i=1}^{6} \sum_{j=1}^{N_{\text{segments}}} \left\| \hat{\mathbf{p}}_{ij} - \mathbf{p}_{ij} \right\|_2^2 \right) \tag{5.2}$$

Where N_{segments} is the total number of tracked segments. $\hat{\mathbf{p}}_{ij}$ denotes the i^{th} centre point among the j^{th} reference segment's bounding box faces. The square distance is then calculated as the positional difference from the corresponding point on the dynamic humanoid ($\mathbf{p}_{ij}$).

A squared distance measure rather than the L2 norm was used in the position, velocity and local pose rewards, based on the implementation of (Booth et al. 2020), which led to faster

learning during preliminary results. The surface points' positions were resolved in the pelvis' reference frame, illustrated on Figure 5.4. Like (Bergamin et al. 2019), the dynamic humanoid's local reference frame was also considered to originate from its pelvis, but its orientation matches the reference frame's. This implicitly penalises facing the wrong direction with the dynamic humanoid.

The second term, the *velocity reward*, was used to match the linear velocities of these points (Equation 5.3).

$$r_v = \exp\left(\frac{-1}{N_{\text{segments}}} \sum_{i=1}^{6} \sum_{j=1}^{N_{\text{segments}}} \|\hat{\mathbf{v}}_{ij} - \mathbf{v}_{ij}\|_2^2\right) \tag{5.3}$$

The linear velocities of these points were also affected by the angular velocity of their parent segments. While the direction of the velocity vectors was also resolved in the local reference frame, the pelvis's global velocity was not subtracted from their values. This strengthens the requirement to match the overall velocity of the locomotion.

The local rotation of each segment with respect to its parent segment determines the third term in the reward (Equation 5.4).

$$r_{local} = \exp\left(\frac{-6.5}{N_{\text{segments}}} \sum_{j=1}^{N_{\text{segments}}} \|\hat{\mathbf{a}}_j \ominus \mathbf{a}_j\|_q^2\right) \tag{5.4}$$

where $\|\hat{\mathbf{a}}_j \ominus \mathbf{a}_j\|_q$ is the angle magnitude of the angle-axis decomposition of the quaternion difference of the two rotations.

The last reward term was provided based on matching the velocities of the two centres of mass (Equation 5.5).

$$r_{v_{CM}} = \exp\left(-\|\hat{\mathbf{v}}_{CM} - \mathbf{v}_{CM}\|_2^2\right) \tag{5.5}$$

Finally, these reward signals were combined with a scaling factor based on the difference of the vertical location of the two head segments (Bergamin et al. 2019) (Equation 5.6). This prioritises learning how to avoid falling and tripping first, before reward from motion tracking

can be gained.

$$e_{\text{fall}} = \max\left(\min\left(1.3 - 1.4\left\|\hat{\mathbf{h}}_{CM_{head}} - \mathbf{h}_{CM_{head}}\right\|_2, 1\right), 0\right) \tag{5.6}$$

This reward was collected at every time step t when the agent takes an action. If an episode lasts more than 15 s, or the agent's rewards in a step fell below a near-zero threshold, the episode was terminated, and the simulated agent was reinitialised to the reference's state. As the reference animation was not restarted, episodes started at different states of walking. This approach can be thought of as a simplified version of Exploring Starts (Sutton et al. 2018), or Reference State Initialisation (Peng et al. 2018).

Locomotion task

The learning process took place in an $8\times8\,\text{m}^2$ area. The motion-matching animation was generated independently of the dynamic simulation (i.e. the state and actions of the agent did not influence the animation). This kinematic animation moved between target locations, smoothly transitioning between turning and straight walking (see Figure 5.5).

When the character was closer than 0.8 m to a target location, a new one was generated that was at least 2.4 m away, sampled from a uniform distribution within the learning area. A wait period was introduced between target locations with a 30% probability, sampled uniformly from the range [6-9] s. During a wait period, the animation comes to a stop, and stands still with a desired horizontal velocity of $0\,\frac{\text{m}}{\text{s}}$, before receiving the next target location. Once the agent learnt to handle turns and the policy's performance converged to a static level ($\sim$10 million decision steps), additional perturbations were introduced every 1.2 s, in a uniformly sampled direction, applied for 0.24 s at a uniformly sampled body segment, with a force sampled uniformly from the range of [50-150] N. This type of perturbation can force the agent to explore recovery strategies in situations it would not encounter in its regular environment, but will once a change is introduced (such as a gait pathology) (Akkaya et al. 2019).

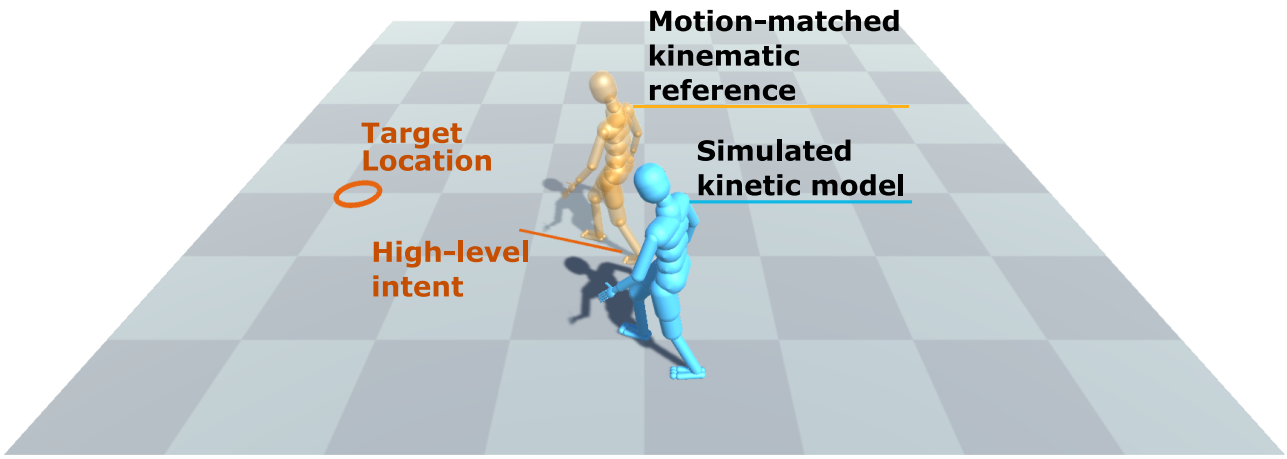


Figure 5.5: Snapshot of the non-steady-state locomotion learning environment. The vector pointing from the reference to the target location is the high-level intent provided to the gait synthesis agent. Once a target location was reached, a new one was generated which may necessitate turning.

5.2.3 Virtual device

The use of a transtibial prosthesis was modelled by reassigning control over below knee DoFs unilaterally to a separate agent (Kidziński et al. 2020; De Vree et al. 2021). Dividing autonomy over the humanoid character, the locomotion and assistive agent then share the collaborative goal of restoring the character's original gait. This approach does not reproduce the behaviour or limitations of any one specific prosthetic device and assumes perfect weight matching and mechanical interface conditions, simplifying an already complex design space. Both flexion and adduction were controlled by the assistive agent, in contrast to more common flexion-only designs.

Similarly to the gait policy, PPO was used to train the prosthesis policy with the same network structure. However, only 128 units were used per hidden layer, and the assistive agent was queried only at 30 Hz.

Observations, actions, rewards

There is an important distinction to make in the types of observations used as inputs to the device policy when compared to the gait policy. The gait policy may use all privileged information available from the simulation. In contrast, the kinematic poses from the motion-matching

process should not be used by the assistive agent. Instead, inputs to the device's control should be signals available through plausible instrumentation of real hardware. The simulated powered ankle device uses virtual accelerometer ($\in \mathbb{R}^3$) and gyroscope ($\in \mathbb{R}^3$) measurements from virtual IMUs placed on both the shank and foot, as well as virtual encoder signals for the ankle joint angle ($\in \mathbb{R}$) and angular velocity ($\in \mathbb{R}$) for a 14 dimensional observation vector in total.

One of the main benefits of simulated locomotion environments is the ability to test novel control schemes. Conditioning behaviour on high-level intent is one such improvement proposed for improving performance in non-steady-state and irregular environments (Tucker et al. 2015). Since biosignals commonly used in making this type of control schemes anticipatory and intuitive are not available in a straightforward way in simulation, an intermediate representation of the intent is necessary. In a real-life system, this representation would need to be estimated with human-machine interfaces (e.g., with electromyography). The intermediate signal could then be provided to the worn device as a compressed form of its user's intent. If the intermediate representation is chosen so that it is available during simulation, then it can be used as a surrogate input for controllers conditioned on high-level intent. As the desired walking velocity was used to drive the locomotion synthesis based on the experimentally determined relationship between movement and this abstract intent (enforced through motion matching), it is a natural representation to use for this purpose.

The actions of the device agent were interpreted as PD parameters of stiffness, damping and joint angles. No reference motion was used to influence this target for the device agent. Similarly to the gait policy, a high stiffness and damping starting point was used for the controller ($35\ \frac{\mathrm{N{\cdot}m}}{\mathrm{rad}}$ and $10\ \frac{\mathrm{N{\cdot}m{\cdot}s}}{\mathrm{rad}}$).

This horizontal velocity ($\in \mathbb{R}^2$) is a continuous high-level control signal and is conceptually between activity recognition and direct volitional control methods (Tucker et al. 2015). The prosthesis policy was conditioned on this intent by manipulating the policy network parameters through a hypernetwork (Ha et al. 2016).

5.2.4 Simulation conditions

In early tests with simulated characters using a passive prosthesis, it was confirmed that robust locomotion policies can be learned with passive devices. However, this was at the cost of compensatory movements (Lee et al. 2019). Therefore, conditions were included where the gait policy was frozen after training without a prosthesis, then the device was introduced and only the prosthesis policy was trained. Once a gait policy was learned without a prosthesis (stopping training after ∼20 million decision steps), it was used as a starting point for these conditions after it. In these cases, compensatory changes to the gait were prevented, and the assistive agent had the responsibility to restore stable locomotion. When the gait policy was no longer trained, the action with the maximum likelihood was selected for it deterministically (the prosthesis policy remained stochastic). All conditions involved walking with turns and stops:

1. No prosthesis, perturbations added.
2. Active, non-intent-driven prosthesis, gait policy pre-trained.
3. Active, intent-driven prosthesis, gait policy pre-trained.
4. Active, non-intent-driven prosthesis, gait policy pre-trained and frozen.
5. Active, intent-driven prosthesis, gait policy pre-trained and frozen.

5.3 Results

5.3.1 Gait policy

The learned policy generates robust walking between arbitrary landmarks, while being underactuated (no external forces applied to center of mass to help balancing). It was able to synthesise turning at various degrees and recover from losing balance (Figure 5.6). A video

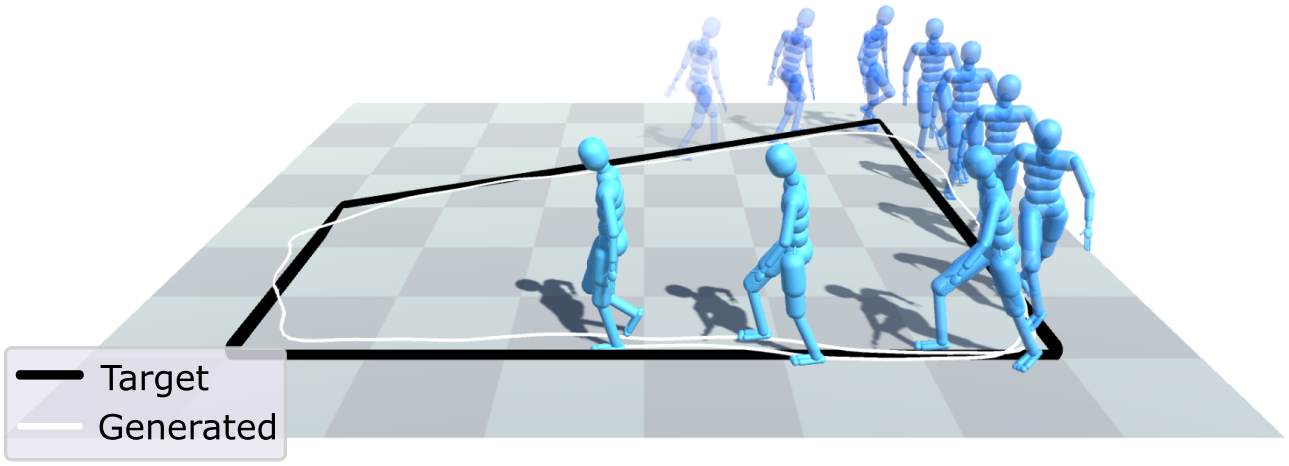


Figure 5.6: Path tracking ability of the gait policy. "Stroboscopic" visual trail of the agent's poses made by following the learned gait policy recorded at 1 Hz. The agent was capable of performing continuous circuits without falling or deviating from the path.

demonstration of the motion-matched reference and the synthesised dynamic gait is available as supplementary material[†].

While the DRL reward implemented only constrained kinematic properties of the motion synthesis, important characteristics of gait dynamics were reproduced by the final policy. This includes bimodal peaks of the ground reaction force (Figure 5.7) (Winter 1984), and ankle dynamics (Figure 5.8).

The peaks of the ground reaction force were between 1 and 1.2 times the body weight, following normative data (Chao et al. 1983). Moreover, there were contact forces (above 5% body weight) with both feet in the simulated walking during 21% of the gait cycle. This matched the expected ratio for double and single support within the gait cycle (Kharb et al. 2011). The peak ankle moment was 1.56 $\frac{\mathrm{N\,m}}{\mathrm{kg}}$, which was also in accordance with experimental results (Moisio et al. 2003).

To verify the suitability of the learned locomotion synthesizer to test prosthesis controllers, the Finite State Machine based controller of the Open-Source Leg (Azocar et al. 2020; Simon et al. 2014) was reproduced and tuned manually for the virtual user. This rule-based controller transitions between stance and swing states, as determined by the load on the prosthesis, and changes its behaviour based on knee velocity and ankle angle. Gait phase detection progressed

[†]DOI: 10.1109/TNSRE.2023.3336713/mm1

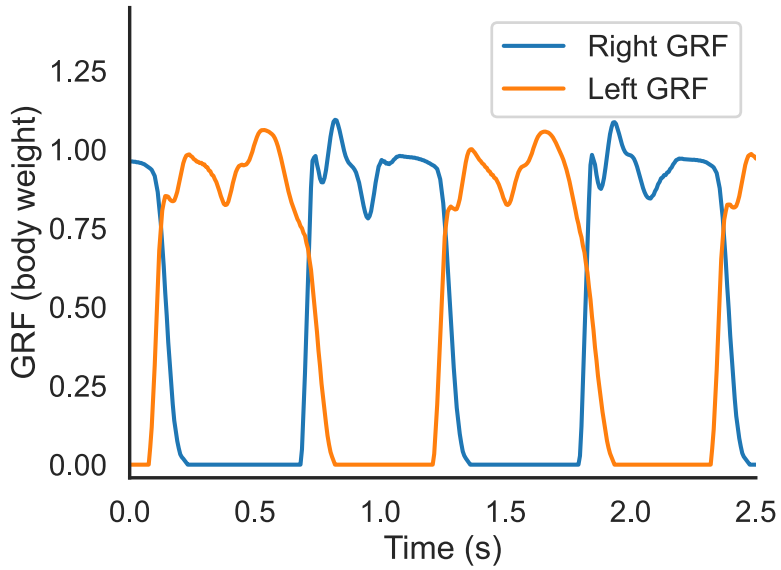


Figure 5.7: Ground reaction force during walking as measured by virtual sensors placed on both feet of the character. Recorded at 1.3 $\frac{\text{m}}{\text{s}}$ straight walking.

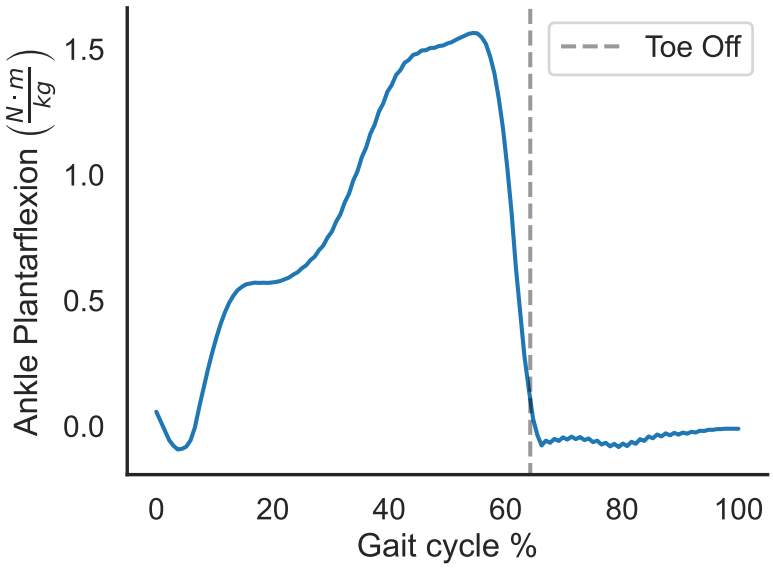


Figure 5.8: Plantar flexion torque in one of the ankles during simulated gait. Early stance dorsiflexive and late stance plantar-flexive torque can be observed (Winter 1984).

naturally through early stance, late stance, early swing and late swing states (Figure 5.9), and the virtual device was able to restore walking. Transition to swing phases happened at 60% of the gait cycle (tracked between heel strikes), a value matching the expected timing in gait (Kharb et al. 2011). This FSM controller was primarily used as a validation of the locomotion synthesizer. Its natural progression through estimated gait phases indicated that the synthesized locomotion provided a reasonable kinematic and kinetic context to the virtual prosthesis. In further results the FSM controller was no longer used, instead a DRL policy controlled the prosthetic ankle.

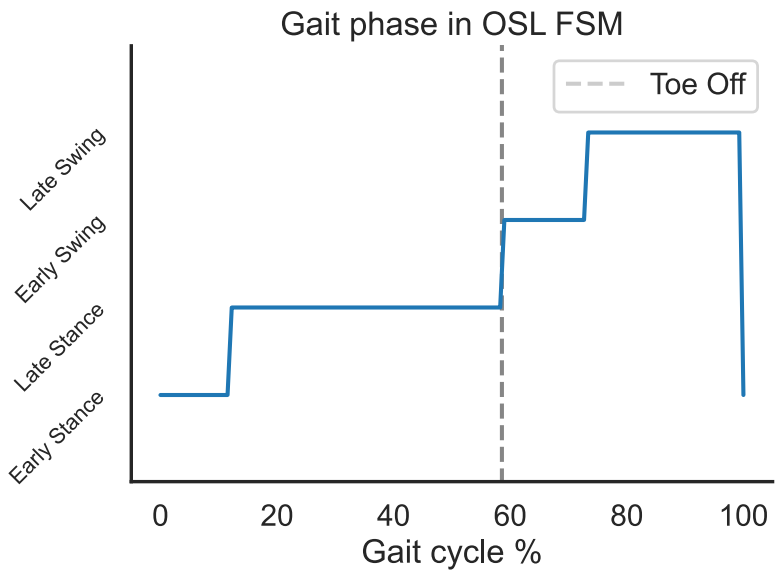


Figure 5.9: Phase profile of the rule-based classifier used in a traditional Finite State Machine controller reproduced from (Simon et al. 2014), applied to the virtual user using an active prosthesis.

Interestingly, in the condition where the gait policy was learned in parallel with the prosthesis policy, the additional observations of user intent provided no advantage to the joint performance of the human-prosthesis system (Figure 5.10).

In contrast, in the case where a pre-trained and frozen gait policy was used (hence no compensatory behaviour was adopted), conditioning the policy on user intent led to significant improvements in stability (Figure 5.11). In particular, the benefit was most apparent when the character transitioned between standing and walking. Without intent available, and with no compensatory movements, the assistive agent was unable to find an appropriate policy. This

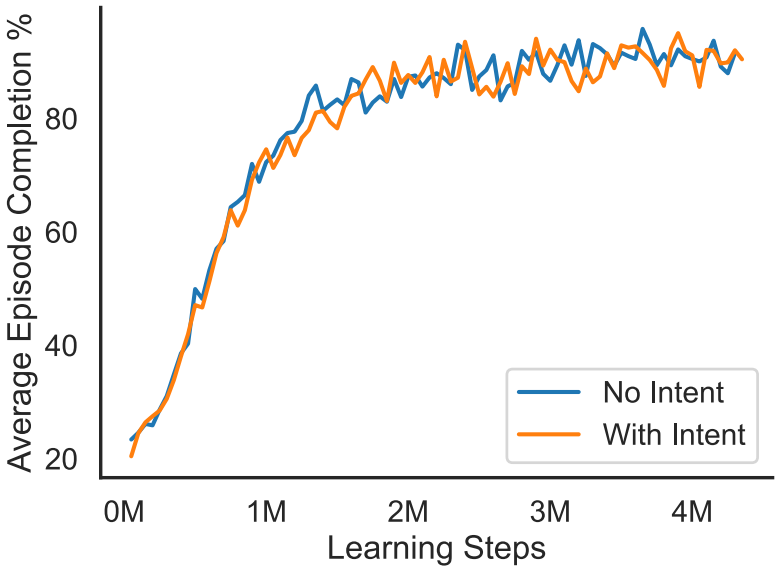


Figure 5.10: Collaborative performance evolution of learning the (pretrained) gait policy and prosthesis policy simultaneously (conditions 2 and 3). Providing the intent as additional observation yields no improvement in these conditions.

resulted in more frequent stumbling or tripping. When the intent was available, the prosthesis policy adjusted the control parameters immediately and continuously as soon as there was a change in intent, allowing the human-prosthesis team to transfer between locomotion modes (Figure 5.13). Without adaptation from the user, the mean time before the first trip event was 34 s and 12 s for the intent-driven and the non-intent-driven prosthesis respectively. If learning compensatory movements was allowed, the prosthesis user could walk for more than 3 minutes without tripping.

These two prosthesis conditions (4 and 5) were compared with the non-prosthesis-user (condition 1) through their average walking speeds, produced ankle torques and pose tracking error during steady-state straight walking (Table 5.1). Both prosthesis conditions led to slower walking when their gait policy was prevented from learning compensatory motion. Furthermore, the intent-driven prosthesis exhibited torque profiles closer in magnitude to those of the original gait policy (Figure 5.12). This potentially arises from the non-intent-driven prosthesis policy being unsure whether it should prepare for a transition to standing or keep walking.

Pose tracking error was evaluated using by quantifying the average total Euclidean distance of the joint positions in Cartesian space and the position of the corresponding joint in the reference

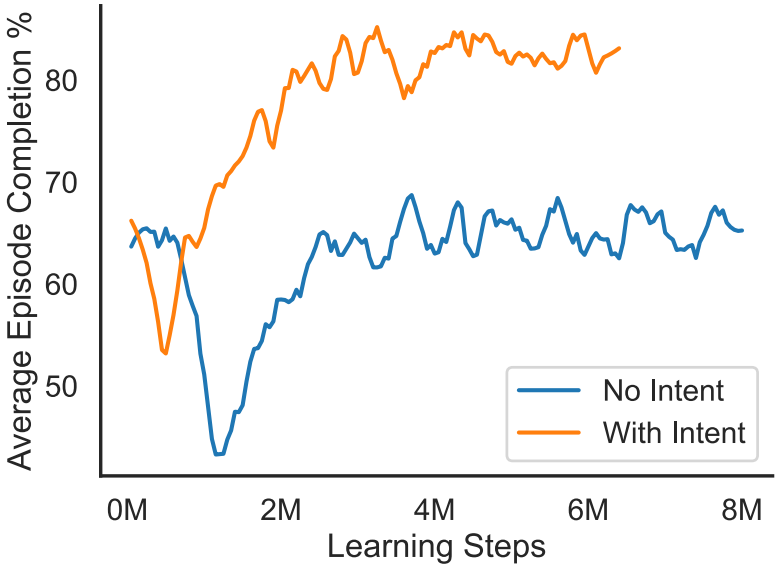


Figure 5.11: Collaborative performance of user-prosthesis system, when the gait policy was pretrained without a prosthesis and could no longer learn. Only the prosthesis policy was adapted (Conditions 4 and 5). Note that the gait policy is deterministic here, as it no longer needs to explore by injecting noise in its actions. This yields the higher performance at the start compared to Figure 5.10.

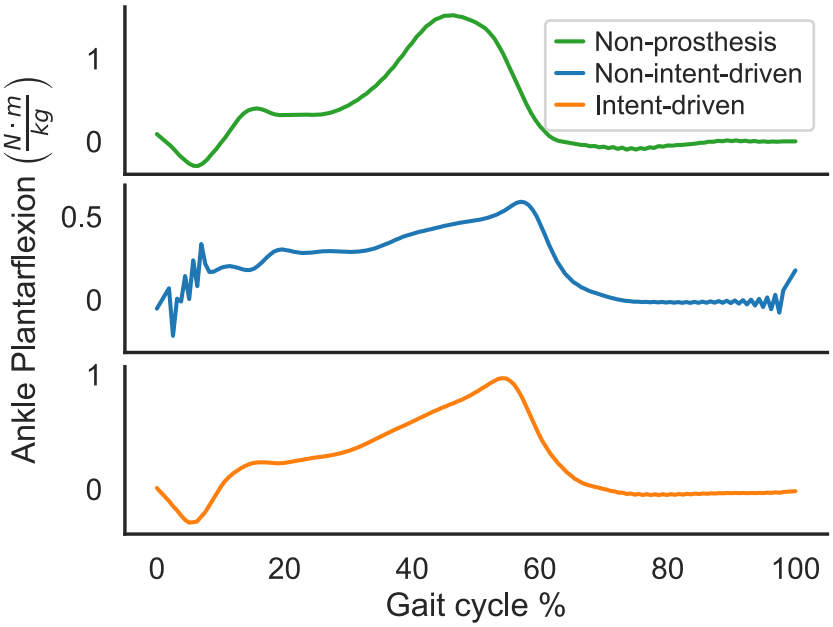


Figure 5.12: Plantar flexion torque output of the ankle of the non-prosthesis-user (condition 1), the prosthesis with no intent input (condition 4), and intent-driven prosthesis (condition 5).

animation (Equation 5.7) (Peng et al. 2021). The positions were relative to the location of the corresponding root joints (x^{root} and $\hat{x}^{\text{root}}$).

$$e_t = \frac{1}{N_{\text{joints}}} \sum_{j \in \text{joints}} \|(x_t^j - x_t^{\text{root}}) - (\hat{x}_t^j - \hat{x}_t^{\text{root}})\|_2 \tag{5.7}$$

Table 5.1: Performance of non-prosthesis (condition 1), non-intent-driven prosthesis (condition 4) and intent driven prosthesis (condition 5) during straight walking.

Measure	No Prosthesis	No intent	Intent
Walking speed $[\frac{\text{m}}{\text{s}}]$	1.30	0.98	1.05
Peak ankle torque $[\frac{\text{N m}}{\text{kg}}]$	1.52	0.59	0.97
Pose error [m]	0.026	0.045	0.043

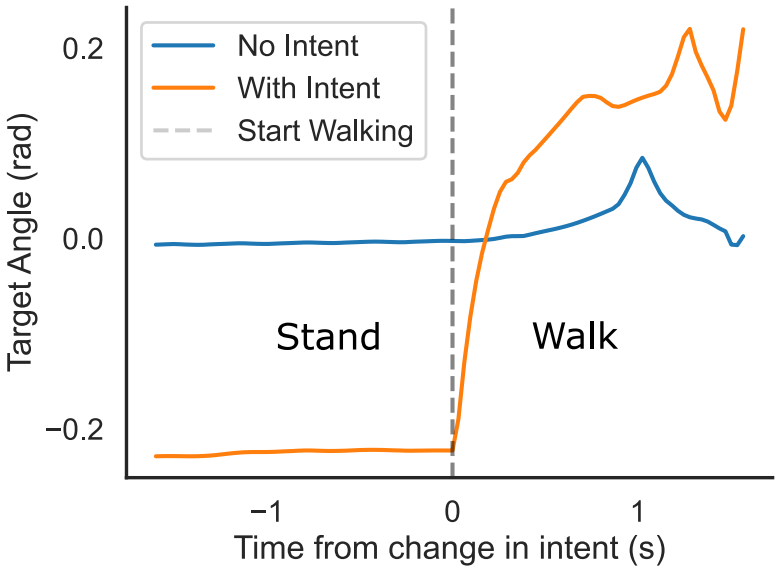


Figure 5.13: Trajectory of the equilibrium ankle angle in the PD control parameters, set by the intent-driven (condition 5) and non-intent-driven (condition 4) systems when there's a step change between the virtual user's intent (going from standing to walking).

5.4 Discussion

We have presented a human-prosthesis system model that provides key insights for its real-life equivalent. A high-level, abstract but continuous locomotive intent representation was shown to be helpful in non-cyclic gait scenarios, usually handled by classification processes (Tucker et al. 2015). By observation, most of the stability gained from relying on intent in conditions 4 and 5 was present when the agent was coming to a stop, or starting to walk from standing. The intent-driven policy was able to change its behaviour when the virtual user's goal changed, whereas the policy relying only on intrinsic sensing could react to movement already underway, leading to stumbling and trips (Figure 5.13). Sharing information through a limited channel is a known way of stabilising and improving performance in multi-agent learning scenarios (Zhang et al. 2021), and intent estimation is an experimentally plausible approach to represent this.

However, it is important to note that adaptive controllers, such as those learned through DRL, have challenging learning dynamics when collaborating with a second non-static agent (conditions 2 and 3). Since the human has significantly more agency over managing the gait, exploration in assistive policies can hinder the gait policy's ability to exploit its competence,

leading to lower the overall performance. This is apparent, for example, in the initial dip in the reward shown in Figure 5.11, where the gait policy was frozen. In a co-adaptation setting, this can encourage the prosthesis policy to act "lazily" (Sunehag et al. 2017), and let the human to learn safer, compensatory strategies. We believe this leads to an under-utilisation of the information available to the device policy, which explains the lack of difference when providing intent to the device (Figure 5.10). Indeed, the peak plantar flexion torque from the prosthesis in condition 3 is 0.5 $\frac{\mathrm{N\,m}}{\mathrm{kg}}$, half as much as in condition 5 where no compensation from the user is available. This type of situation may arise as well when applying reinforcement learning methods with real users who cannot risk falling over, who will also adapt their movements, potentially leading sub-optimal prosthesis policies. Through robust simulation environments, exploration can be enforced that discourages "lazy" policies (e.g., by disabling or slowing virtual user adaptation). Control parameters identified this way may provide starting points during fine tuning with real life users, which may avoid control strategies that do not fully take part in the shared autonomy of the movement they are supposed to assist. Although DRL was used to learn the device policy as well, the locomotion synthesizer could be combined with other, non-deep-learning based controllers (Bryan et al. 2021).

During the development of the environment, sources of instability, errors in the implementation and suitable parameter ranges were identified for the prosthesis controller. Solving these issues first in the context of the synthesised locomotion partially mitigates the risk discovering these experimentally, and could contribute to using laboratory, equipment, and most importantly participant time more efficiently. In addition to the DRL device policy, reimplementing an FSM controller served as a proof-of-concept for testing rule-based policies in simulation, and practising or evaluating their tuning process.

The construction of the learning environments was considerably accelerated through the use of Unity, a development platform designed for efficient editing of virtual scenes. Its potential use case to model biomechanics and robotics has been recognised before, with the main criticism being towards the built-in physics engine's prioritisation of performance over accuracy (C. K. Liu et al. 2021). Thanks to the recently released Unity plugin of the MuJoCo physics engine, it is possible to benefit from streamlined design and visualisation tools of Unity and still simulate

biomechanically validated interactions.

The presented actuation scheme, managed by the DRL policy originated from the computer graphics field (Bergamin et al. 2019). Despite this, there are rough parallels with current neuromechanical theories of human motor control. The controller consists of an open loop control signal assembled from motion templates, which was then modulated based on signals of abstracted proprioception and high-level goal observations. This is reminiscent of the concept of Central Pattern Generator based movement (Duysens et al. 1998). The motion-tracking and torque actuation-based method allowed for fast training of policies on limited hardware, while still producing plausible dynamics. For the purposes of simulating prosthesis use, torque actuation may be sufficient to provide the necessary context for the device. However, in partially over-actuated systems, like exoskeletons, realistic models of joint mechanics and actuation are more important, which is likely to involve musculoskeletal actuation (Lim et al. 2017). Lastly, an important consideration is the simulation of the gait pathology. Prosthesis use can be characterised by various patterns, such as gait asymmetry or specific compensations such as circumduction. These could be closely replicated by tracking patient motion instead of non-prosthesis user data. However, this could be counter-productive, as then the virtual user would actively resist assistance from the prosthesis that would mitigate these gait patterns, just to better imitate its reference. It is more appropriate to motivate the agent to match non-compensatory movements, and constrain it mechanically (e.g., by limiting joint power, adding extra weight) or "physiologically" (e.g., by penalising joint load to simulate pain, or introducing fatigue mechanisms) in a way that the expected gait patterns emerge on their own. The scope of this study did not include a detailed comparison with experimental prosthesis use; however, this is a key aspect to be investigated in future work.

MuJoCo in particular is an attractive physics engine to use for future work due its efficient computation, which extends to muscle modelling (H. Wang et al. 2022; Caggiano et al. 2022). Furthermore, due to its flexible collision constraint configuration (Todorov et al. 2012), surfaces of different materials and compliance can be easily simulated. This can be advantageous when trialling cushioned heels, or different types of challenging terrains.

Other key additions to this work would include other environments mimicking key activities of daily living, such as climbing stairs, tackling rough terrain or navigating crowds. The model of the prosthesis could also be extended. Key properties of the real socket-limb interface, such as pistoning, could be modelled to better capture the dynamics of gait during prosthesis use. This could reveal stress/strain relationships, which the optimisation process could take into account. The impact of aspects such as the number of DoFs, their range of motion, stiffness and power could also be further investigated to inform the design of new devices. The observations used as inputs to the policy are related to biological signals associated with human balancing (e.g., otholiths sensing linear acceleration of the head (Forbes et al. 2018)). Further biologically inspired input signals (e.g., the Golgi-tendon-like observations used in (J. Park et al. 2022)) could not only improve artificial locomotion policies, but have implications on their role in human motion learning. Similarly, more nuanced multi-agent learning schemes should be explored to robustly model short- and long-term coadaptation. Lastly, augmenting the reference motion synthesizer with a more diverse motion capture data set could have positive effects on performance. Motion-matching is prone to reuse only segments of its database, therefore not all reference gait cycles will be equally represented, which may bias the policy. More complex reference motion synthesizer, such as neural state machines (Starke et al. 2019) could induce a more diverse training set. Alternatively, other DRL controllers that do not rely on synchronised reference motion in the first place could be applied (Peng et al. 2021).

There are various benefits to using DRL methods for finding the gait policy. While there is a computational overhead associated with training a DRL policy, once trained they are cheap to evaluate. A single policy may be trained to be robust for a range of virtual users and devices (J. Park et al. 2022; Lee et al. 2021; Feng et al. 2022). Locomotion agents with differences in weight, height and other parameters could be generated to evaluate devices and their controllers on a diverse virtual user population. Furthermore, DRL strategies can mimic various walking styles with a given humanoid model (Peng et al. 2018). The function approximators commonly used in these methods can also establish connections between different representations of intent, sensory observations and the control policy. This is possible even if they are not directly accounted for in the cost or reward function, unlike trajectory optimisation methods.

The locomotion synthesis framework created and used in this study to train the gait policy was refactored, documented and released as an open-source project (Hodossy et al. 2023b). Additional features, such as the prosthesis environments updated for the latest package version along with more diverse walking environments (e.g., rough terrain and stairs) are planned additions for the future.

5.5 Conclusion

A gait policy learning environment was built through the combination of an accurate physics engine, a kinematic motion synthesizer and an accessible DRL framework. Autonomy over the simulated human's movement was shared between the gait policy and a second controller that operated a model of a unilateral transtibial prosthesis, forming a representative virtual test platform for wearable robotic devices. Controllers were trained and evaluated in a non-steady-state locomotion scenario involving walking, standing and turning. A continuous high-level intent representation was shown to be a useful control input, provided that compensatory gait patterns from the locomotion agent do not prevent the assistive device to capitalise on the additional information.

Human locomotion is capable of tackling various situations such as rough terrain or navigating crowds. By decoding the motor intent, assistive devices could adapt to their user's movements and their diverse environments. Using simulated gait and assistive devices, this complex design space can be explored in a low-cost and accessible way, promoting the transfer of the next generation of intent driven controllers from the lab to real-life scenarios.

Chapter 6

Modelling Locomotion with the IIT SoftFoot

6.1 Introduction

In the previous chapter the effects of the amputation and introduction of a transtibial prosthesis were modelled through the following steps:

- Altering the passive properties of the joints below the knee.
- Removing agency of the humanoid control policy over the control of the affected DoFs.
- Assigning the control to a second policy, representing an assistive device.

This is not an uncommon strategy in modelling prosthesis use (De Vree et al. 2021; Kidziński et al. 2020), as it avoids a series of otherwise difficult questions regarding the design and suspension of the prosthesis (i.e., attaching it to the human body). In this chapter we examine some of these questions and use them as an opportunity to evaluate a recently presented prosthetic foot prototype developed at the Istituto Italiano di Tecnologia (IIT), the SoftFoot (Piazza et al. 2024; Pace et al. 2024). Through this we demonstrate the potential application of our versatile

locomotion agents as virtual testbeds for evaluating passive lower limb prosthetics. Due to their lower cost, robustness and ease of maintenance, innovation in passive devices has a shorter path to impact than active powered devices. This is true in particular for low- or middle-income countries (Harkins et al. 2013). Modelling the SoftFoot was an appropriate choice due to its intermediate technology readiness level; it has been developed for long enough for user testing with physical prototypes, but there are still plenty of questions regarding its design choices that can benefit from computational modelling. Additionally, the SoftFoot's piece-wise rigid composition lends itself well for recreation in the MuJoCo physics engine, which (like most other human-scale physics engines (Erez et al. 2015)) makes this assumption in the simulated kinematic chains.

The SoftFoot incorporates the design principles of the SoftHand, IIT's robotic prosthetic hand, and adapts them to the lower limb (Della Santina et al. 2018). Both feature compliant mechanisms, constructed from piece-wise rigid elements interfacing each other through rolling contacts, secured through a combination of tendons elements of varying elasticity. As end effectors, they both strive for robustness and reliability when interacting with the environment using through their built-in compliance. Moreover, they are heavily underactuated with cables routed through their rigid elements to distribute forces and torques. The SoftFoot adopts an anthropomorphic design, mimicking the arches of the human foot. Its sole consists of five parallel chains of articulated elements, replicating the functionality of the human longitudinal arch by conforming around obstacles in the sagittal plane. The coupled system of elastic tendons and joints absorbs impacts and provides a larger support surfaces formed of multiple contact points with the ground and obstacles (Piazza et al. 2024). This in turn provides a more stable stance on uneven terrain, allowing the ZMP of the body to rest within the support. The articulation at the metatarsophalangeal joint of the SoftFoot provides a natural progression of the stance thanks to a more proximal pivot point compared to a rigid foot (where it would be located at the toes) (Piazza et al. 2024). This lowers the amount of torque necessary to lift off the heel for the same overall foot size. The SoftFoot is compliant to the walking surface but acts as a rigid lever when propelling the body forward. In contrast, a sole consisting of cushioned components could only replicate this level of compliance at the cost of stability, as linear and rotational

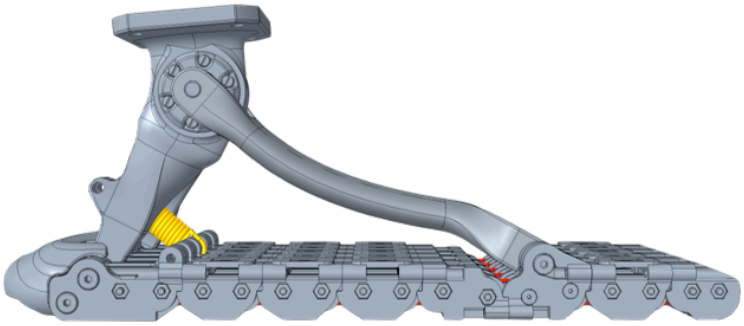


Figure 6.1: computer-aided design (CAD) model of the IIT SoftFoot (Piazza et al. 2024), rendered by Matteo Crotti.

stiffness are coupled in them.

Further anthropomorphic mechanisms in the SoftFoot include a high-stiffness cable that is routed below the sole joints, coupling and limiting their rotation. This is reminiscent of the plantar fascia in humans, and enables the stiffening-by-compression behaviour of the foot, distributing loads across the sole. The elastic components running above the joints play the role of the tendons and ligaments in maintaining the soles shape compliantly.

While there have been extensive efforts on computationally and analytically modelling the SoftFoot to select and justify its design choices (Piazza et al. 2024; Pace et al. 2024), these have all been performed outside the context of a subject using it dynamically. Instead, online validation was performed by IIT experimentally with a patient participant. In this chapter we performed the following extensions of our simulated prosthesis use model from Chapter 5:

- We replaced the "generic" prosthesis used for our virtual subject with a model of the SoftFoot.
- Attached the foot compliantly to the residual limb.
- The locomotion policy was adapted to the passive SoftFoot in unilateral and bilateral prosthesis use settings.
- Evaluated on tackling obstacles not seen during training.
- Our motion matching locomotion synthesizer was extended with a inverse-kinematics pipeline, procedurally adapting each step to conform to terrain involving steps and slopes.

- We contrasted the resulting motion with experimental data of a transtibial amputee using a real SoftFoot, to identify limitations and areas for improvement. This data was recorded as part of the Natural BionicS* project. The data acquisition was performed by IIT, collected and provided by Dr Anna Pace. Inverse kinematics and dynamics to extract joint angles and forces was performed using respective processing pipelines in OpenSim by Dr Hristo Dmitrov. Results with walking using an Ottobock Triton foot were also available with the same subjects. This commercial prosthetic foot features a carbon fibre structure with energy storage and return properties.

For this end, we translated and adapted the design files of the IIT SoftFoot to the MuJoCo physics engine. This work was done with jointly with Matteo Crotti from IIT.

6.2 Methods

6.2.1 SoftFoot in MuJoCo

The SoftFoot's design files were exported in the URDF format, an XML-based human readable markup language that describes the kinematic chains making up the prosthesis, their inertial properties and the parameters of the associated DoFs. As the features of URDF are a subset of MJCF, MuJoCo's own native model definition format, each elements and their attributes could be translated to MJCF to have full access of the physics engine's functionality. The associated 3-D mesh files of the foot geometry have been simplified and reduced in resolution (Figure 6.2).

Attempting to represent every aspect of system may not be appropriate when beginning to create a model (Hicks et al. 2015). Even without modelling the tendons, the kinematic complexity of the SoftFoot far outweighs that of the human body model we employ (see Section 5.2.2). Almost twice as many DoFs are present in the sole alone when compared to the rest of the humanoid (70 against 36). We considered a series of approaches to model each of the key mechanisms making up the SoftFoot (Table 6.1), including both tendon constrained (Figure

*DOI: 10.3030/810346

6.2) and joint constrained version (Figure 6.3). Additional summary on modelling decisions is available on Figure 6.4.

Table 6.1: Mechanism modelling complexity comparison. In digital editions of this thesis, most elements in this table are clickable, linking to relevant sections of the online MuJoCo documentation if available. Legend: *not yet implemented*, implemented, **selected**.

Mechanism	Increasing complexity →			
Attachment to user	Fixed joint	*3 DoF*	**6 DoF**	*Non-linear socket sys-id*
Plantar fascia cable	**Joint limits**	*Fixed tendon*	Piece-wise linear spatial tendon	*Wrapping tendon*
Plantar elasticity	**Joint stiffness**	*Elastic fixed tendon*	Individual spatial tendons	
Heel spring	**Joint stiffness**	*Elastic spatial tendon*		
Sole collisions	From mesh	**Collision primitives**	*Composite primitives*	
Closed-loop chain	**Equality constraint**			
Sole element rolling	**Joint equality constraints**	Rolling Contacts		

For the purposes of running the RL locomotion environments we selected the simplified version, as it qualitatively retains the compliant behaviour of the SoftFoot. We forgoed the tendon elements, and instead model their effect through joint properties (Figure 6.3). The elastic tendon stiffness was linearized around the joint. The rigid cable that plays the role of the plantar fascia was represented through joint limits preventing the collapse of the sole even in absence of the joint stiffness. There is ongoing work investigating the effects these simplifications have. As we build our understanding of these models, increased complexity in both the prosthesis and human model will be warranted for future studies.

The ability to robustly step on obstacles is one of the main proposed benefits of the SoftFoot. Therefore, we modelled obstacle types proposed by (Pace et al. 2023) in MuJoCo. Furthermore, we implemented tools to generate steps and slopes with a parameterised interface (Figure 6.5). Taking advantage of the flexibility of the virtual environment, we extended the functionality of

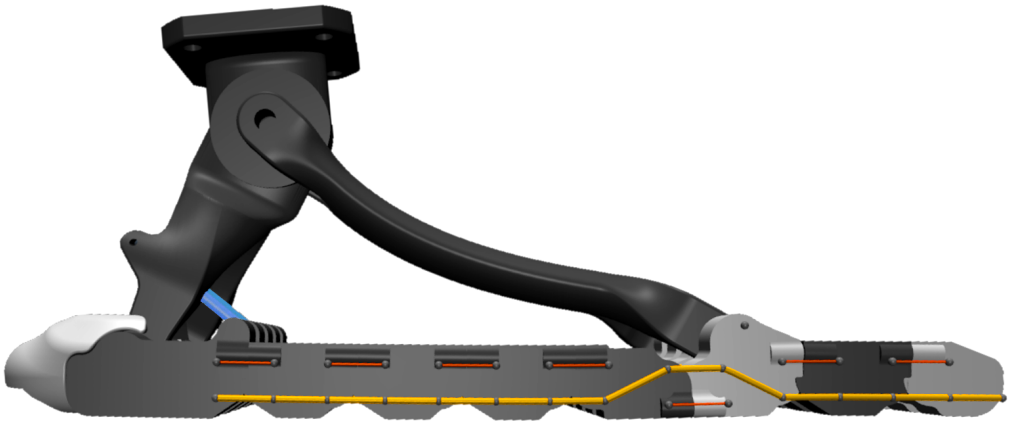


Figure 6.2: Tendon constrained model of the IIT SoftFoot in MuJoCo. The geometry mesh is simplified compared to Figure 6.1, leaving out extraneous elements such as screws and bolts. Elastic tendons shown in red, rigid tendon shown in yellow, and the sole-stabilizing spring shown in blue. The via-points of the piece-wise linear tendons shown with grey spheres between the segments.



Figure 6.3: Simplified model without tendons. The joint axes of the sole DoFs shown with blue vectors.

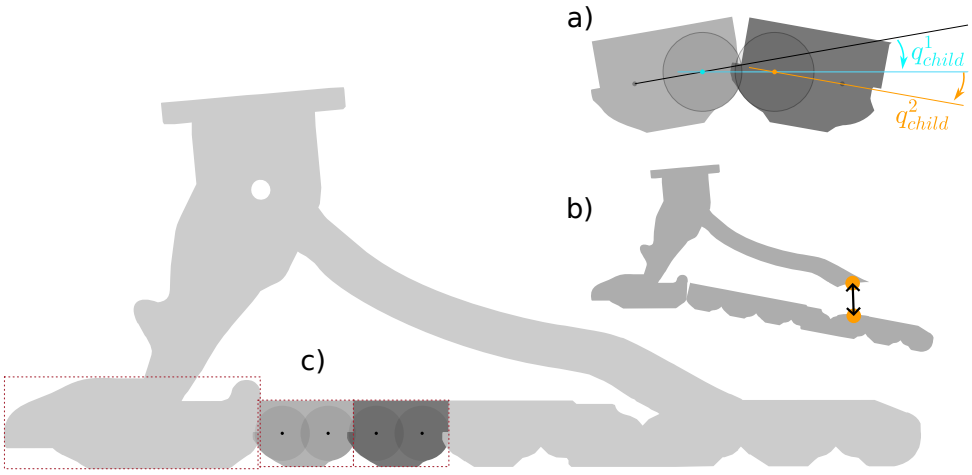


Figure 6.4: Implementation of the rolling motion of sole elements. *a)* Each sole element is modelled as having two parallel hinge degrees of freedom with a lateral axis. By constraining these two degrees of freedom to share an equal value (through MuJoCo's constraint enforcement tools (Google DeepMind 2024a)) no additional frictional contacts need to be simulated to model rolling without sliding between adjecent links. *b)* The same constraint system is used to enforce the connection between the metatarsal link and the foot arch in cartesian space, closing the kinematic loop. *c)* Collisions with the ground are simulated with mass-less colliders formed of the axis-aligned bounding box of each element. Inter-element collisions are disabled to speed up the collision detection.

the MuJoCo Unity plugin to provide support for the height field modelling features of MuJoCo to train on terrain of varying slopes (Figure 6.6). Similarly, environments involving steps have also been created (Figure 6.7).

6.2.2 Stance prediction and motion adaptation with footbase trajectories

Having information and control over where the locomotion agent is stepping can simplify evaluations with obstacles. We can either guide the agent to step on the obstacle, or place the obstacle over where the step is expected. By controlling the motion of the foot during swing, the feed-forward target kinematics we produce may be adapted for terrains of different slopes and height. We hypothesised that this can reduce the learning burden on the RL policy, leading to more capable locomotion agents with shorter training sessions, as it did for learning level ground walking. We followed the "footbase" method of Johansen (2009) and reimplemented parts of their semi-procedural animation adaptation pipeline, adjusted for our motion match-

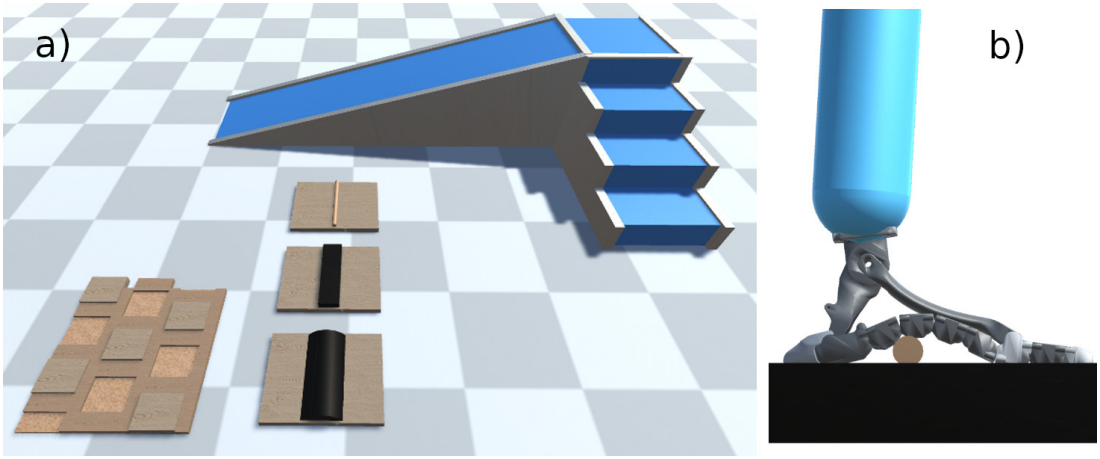


Figure 6.5: *a)* MuJoCo recreations of the uneven terrain (bottom left) and obstacle geometries in (Pace et al. 2023), rendered in Unity. Additionally, procedural stair and ramp generation is also demonstrated. *b)* MuJoCo model of the SoftFoot stepping on an obstacle while used by a locomotion agent.

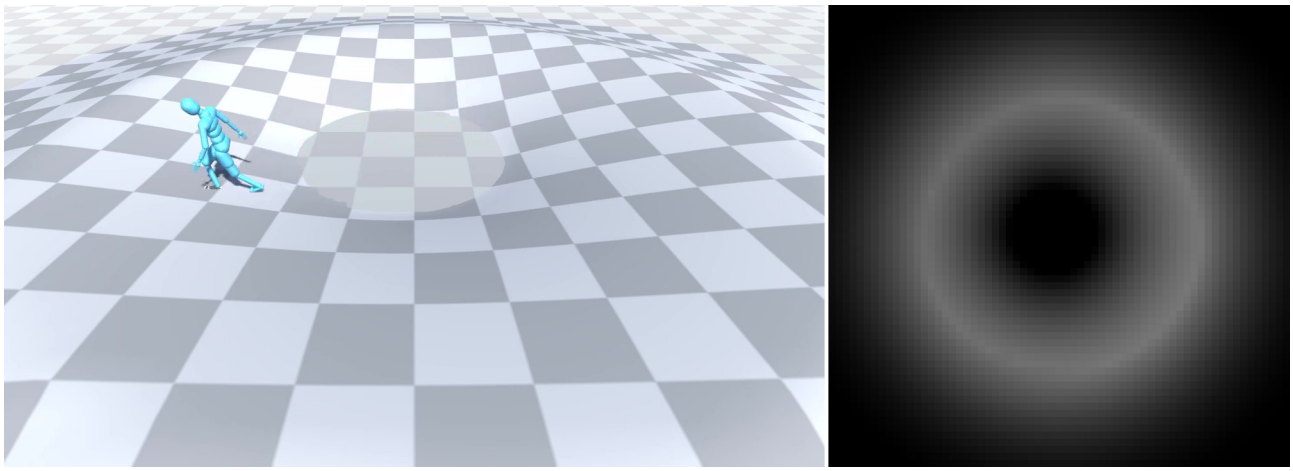


Figure 6.6: Height field stepping environment. By specifying the elevation in a grid, triangular prisms can be extruded to form a continuous smooth surface. As the locomotion task drives walking to arbitrary positions within the area of the valley, tackling slopes at various angles and directions is required from the agent. The height field data are shown on the left in a 65×65 grid, covering a $20m \times 20m$ area and a highest point corresponding to a 60 cm elevation.

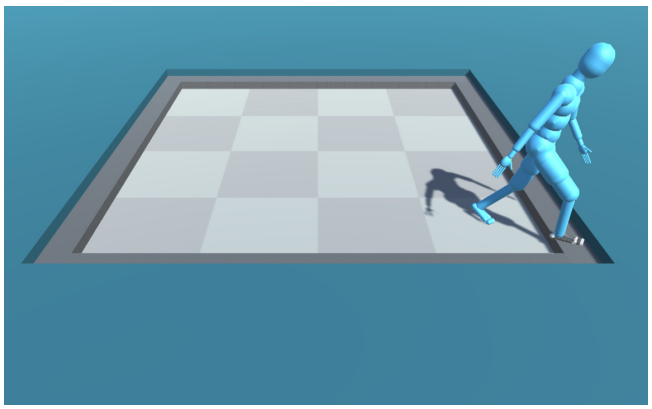


Figure 6.7: Stepping environments. Similarly as with the sloped environment, due to the random traversal, this environment involves both ascension and descension at various angles. The elevation of the blue platform is 10 cm.

ing system. This method parametrises the swing motion of the foot in terms of its position and orientation during the two stance phases surrounding it. The foot position and orientation during the middle of stance is referred to as a "footbase" by Johansen (2009). By shifting the footbase of the upcoming target stance foot placement, we can calculate the necessary offsets for the foot position at each time point to arrive at the new designated location, without intersecting the ground (Johansen 2009). With the new position of the foot determined, standard inverse kinematics approaches can be applied to calculate the necessary modifications to the knee and hip angles. We used Unity's built-in inverse kinematics framework for this purpose. The orientation of the foot was interpolated between the previous stance orientation, the mid-swing orientation from the source animation, and the orientation for the upcoming stance with a sinusoidal function (Johansen 2009). Similarly, an offset for the position of the pelvis was interpolated linearly based on the progress of its position along the vector pointing from the previous footbase and the upcoming one (Johansen 2009).

Through our use of a motion matching system, several simplifications can be made to the footbase method. Motion matching already supports adaptation to turns and different walking speeds, therefore offsets to the pelvis and foot positions were only applied in the vertical direction. To determine how to configure the upcoming footbase to lay normal to the upcoming elevated surface, we must know where it would have landed on level ground. Once again we can exploit motion matching, instead of trying to forecast based on kinematics. As a preprocessing step we identified the middle point of each stance phase in our motion capture datasets (using

another method described by Johansen (2009)) and stores their time points and locations as metadata for our animation. During the simulation, the upcoming target footbase location was available by searching through our metadata based on the time stamp of the currently playing animation. Then by casting a vertical ray towards the level ground footbase location, we found the elevated position at its interception with the ground surface. By repeating this two more times at the vertices of a equilateral triangle laying on the walking surface, we calculated the normal vector for the walking surface as well. As the animation continues playing, we eventually reach the target footbase. It then replaces the role of the preceding footbase, and a new target footbase is identified.

This pipeline was integrated into the motion synthesizer that provided the kinematic reference for the dynamically simulated agent, but the motion of the dynamic agent was not directly constrained by it. During a learning episode, the footbase location predictions were calculated from the current local frame of the dynamic agent. Due to perturbations or changing intent, the true footfall may happen at a different location than the prediction.

A limitation of our current approach is that walking speed and stride length would need to be dramatically reduced to create appropriate motion for continuous stair climbing, which our animation datasets lacked. For this reason only ascension of 2 sequential steps and sloped terrain were investigated.

6.2.3 Locomotion task

We remained in a non-steady state walking scenario for our locomotion task. Other behaviours such as running may have higher energy requirements, but in terms of learning stable policies, moving slower may in fact be more challenging (Kidziński et al. 2020), and walking represents a more common activity of daily living. We extended our environments to include elevated surfaces necessitating stepping up and down, and sloped terrains of various level of incline. As in Chapter 5, we trained our agents in several stages. We introduced more complex environments that specify transitions to and from standing, then stepping and finally tackling slopes, keeping conditions the same until the learning curves stabilise or at least 20 million learning steps have

passed. Curriculum learning strategies like this can enable learning more competent policies, than if the most challenging environment was used from the start (Kidziński et al. 2020).

The zero moment point or centre of pressure is an essential control concept for bipedal locomotion, with implications on the stability and acceleration of the human (Sardain et al. 2004). We explored the option of including the position of the centre of pressure, and the aggregate force vector scaled by the body weight in the observation vector of the locomotion agent, extending its length with 6 values.

We additionally include observations regarding the elevation and potential trip hazards. We provide the local position and normal-tangent orientation (see Section 4.2.4) of the current and past footbase bilaterally (36 observations).

We contrasted the SoftFoot with another virtual prosthesis with the same geometry and material properties, but without the compliant DoFs (a "rigidified" SoftFoot).The impact the compliant mechanism has on pose tracking error, gait symmetry, stance and overall stability was also compared to the rigid foot baseline.

6.2.4 Modelling the mechanical interface

How we model the mechanical interface between a wearable system and its user may be just as impactful as the design choices of the device itself, and deserves consideration (Hicks et al. 2015). In prosthetic interventions a bad socket fit with too much relative motion can completely undermine comfort, stability and usability (Eshraghi et al. 2012). Similarly, misalignment with the human and robotic kinematic chain can subject the user to pain and residual forces, and decrease the effectiveness of exoskeletons (Näf et al. 2018). A straightforward approach is using virtual spring forces to enforce attachment (K. He et al. 2023), alternatively Elastic Foundation Models could also be used for this purpose (J. I. Han et al. 2022). Pure linear translation may not fully explain the relative motion between assistive device and its user, as rotational motion often plays a role as well (Pew et al. 2020).

In this study we opted to model the prosthesis-tibia attachment via a weld constraint, which

non-rigidly couples the six DoFs of the foot's global transform to an attachment point on the virtual subject's tibia. This constraint is flexibly configurable, allowing for polynomial relationships between constraint violation and arising reaction forces, as well as the ratio between linear and rotational forces. The flexibility of this constraint may allow in the future to investigate system-identified models of the mechanical interface, however, for this study the default quadratic relationship was chosen. The length scale parameter was chosen to permit a relative motion of ≈ 3 mm displacement. Due to the available experimental dataset from Natural BionicS featuring an osseointegrated subject, we chose this very stiff relationship to roughly represent an osseointegrated suspension combined with a shock absorption element (Berge et al. 2004).

The attachment point was positioned to align the soles of the prosthetic foot with the original one, and rotated in the sagittal plane 5 degrees (in the dorsiflexion direction) to position the centre of pressure within the support of the feet when standing. In preliminary tests the agent was unable to learn stable standing without lifting the heel on the affected side without this rotation. As proper socket alignment is critical for effective prosthesis assisted locomotion, our choices definitely affect the resulting gait. Investigating how to best apply in simulation the expertise of prosthetists who perform this task manually needs to be investigated in the future. Lastly, it is worth mentioning that the issue of attachment is also relevant for sensor signal synthesis. For example, soft tissue artefacts due to compliant attachment influences IMU-based motion capture (Gu et al. 2023).

6.2.5 Stability and evaluations metrics

A common way to evaluate gait stability without perturbing the participant (which involves potential risk of falling) is through characterising the motion patterns through a the maximal Lyapunov exponent (also known as the local divergence exponent), calculated for a phase space of choice (Bruijn et al. 2012; Mehdizadeh 2018; Inagaki et al. 2023). Given starting states $\mathbf{s}_0$ and $\mathbf{s}_0 + \epsilon$ in the phase space (where ϵ is a small perturbation), a chaotic system will evolve these states to diverge from each other, measured by the Euclidean distance

$d(t) = \|\mathbf{s}(t|\mathbf{s}(0) = \mathbf{s}_0) - \mathbf{s}(t|\mathbf{s}(0) = \mathbf{s}_0 + \epsilon)\|_2$. This divergence can be modelled as an exponential process as given in Equation 6.1(Rosenstein et al. 1993):

$$d(t) = Ce^{\lambda t}, \tag{6.1}$$

where λ stands for the maxmial Lyapunov exponent and C is a constant representing the initial separation. This gives a measure of how chaotic a system is, by quantifying how quickly trajectories from similar states diverge. When applied to kinematic gait data, greater exponents are associated with lower walking stability (Mehdizadeh 2018). We apply Rosenstein's (1993) method to calculate the short-term (λ_S) and long-term (λ_L) estimates of the exponent. Following Bruijin et al. (2012), we extracted λ_S and λ_L from the divergence curves corresponding to 0-0.5 strides and 4-10 strides respectively. We defined the phase space with the two knee angle time series, with two additional lags of each angle with 0.05 seconds in-between them, forming a 6 dimensional phase space.

There's a wide range of other metrics applicable to quantify gait stability, each of them having many further variations (Inagaki et al. 2023; Vienne et al. 2017). While optimal gait during prosthesis use is expected to be asymmetric (Handford et al. 2018), increased differences between the gait parameters of the affected and intact side can indicate increased risk of joint problems (Nolan et al. 2003). For this reason, measures such as the absolute symmetry index (Kodesh et al. 2012; Dingwell et al. 1996) can be used to compare the suitability of different prosthetic interventions. We calculated this index for the percentage of each stride spent in the stance phase during straight walking following Equation 6.2 (Nolan et al. 2003):

$$\text{ASI} = \frac{I - P}{0.5(I + P)} \times 100, \tag{6.2}$$

where I and P correspond to the percentage of the stride spent in the stance phase on the intact and prosthesis side respectively.

The above recorded metrics are useful indicators of gait stability during analysis of motion

trials recorded with subjects. They indicate the robustness of the walking behaviour without needing to perform risky perturbation trials. We face no such restriction with virtual subjects, making them suitable to evaluate in fall-imminent situations induced by trips or impacts. We explored the stability when traversing the uneven terrain proposed by Pace (Pace et al. 2023) (see Figure 6.5) with and without external perturbations. To evaluate whether the compliance of the foot aids in balance recovery while navigating uneven terrain, we induced tripping events by applying a rotation impulse to the pelvis every 10 seconds. We selected a rotation axis uniformly from the unit sphere, and rotated the humanoid 20 degrees along it. If any part other than the feet of the agent contacted the ground, we treated the event as a fall and reset the agent to a walking state.

6.3 Results

6.3.1 Gait policy with SoftFoot and a Rigid foot

Successful locomotion policies could be learned with both the MuJoCo model of the SoftFoot, and its non-compliant counterpart on level ground, stepped and sloped environments. Plausible joint kinematics were possible to generate when introducing obstacles not observed during training, producing motion that resembled examples collected from a real amputee using the softfoot (Figure 6.12). This was in spite of not matching the subject's anthropometrics; the goal of this study was not to create a digital twin of the subject, but to create a new virtual participant who can take advantage of the compliant properties of the SoftFoot.

The hypothesis that the SoftFoot provides a shorter lever arm for lifting the heel is observable in lower peak knee extension moments when compared to the rigid foot (-26.5 ± 5.9 Nm vs. -37.7 ± 13.6 Nm). No such difference is apparent in knee flexion (23.6 ± 5.6 Nm vs. 23.1 ± 4.7 Nm). We also observe a greater foot clearance adopted with the rigid foot than the SoftFoot when descending from obstacles (Figure 6.8). This is potentially because there is more risk of tripping involved with the non-compliant toes. Similarly, we observe a greater increase in

the global pelvis height range when progressing over the stance phase on the obstacle with the rigid foot than with the SoftFoot (Figure 6.9). This pattern is also present in the experimental dataset. Changing the range of motion of the vertical CoM displacement within the gait cycle from the preferred one is associated with elevated energy expenditure (Gordon et al. 2009). Our locomotion agents with compliant prosthetic feet produced more consistent gait when traversing obstacles, which could suggest more efficient walking in prosthesis users.

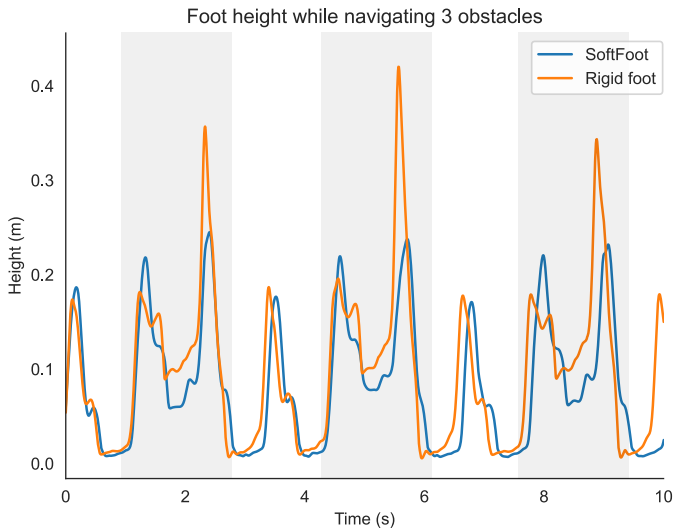


Figure 6.8: Height of the ankle when stepping on and off of obstacles in simulation. Shaded areas indicate period between the two swing phases preceding and following the obstacles. The agent with the rigid foot adopts greater foot clearance when unloading from an obstacle to reduce the risk of tripping.

Lyapunov exponents

Figure 6.10 shows the resulting estimates of the Lyapunov exponents, as defined in Equation 6.1 for 4 conditions: straight walking and level ground walking with turns for both unilateral and bilateral prosthesis use cases. We used 15000 samples from simulation for all conditions, and compared the results for the SoftFoot model and the rigid foot model. In addition to unilateral transtibial prosthesis use, we also included in our results conditions for bilateral prosthesis use, where a SoftFoot or rigid foot was used on both legs during training and evaluation of the gait policy.

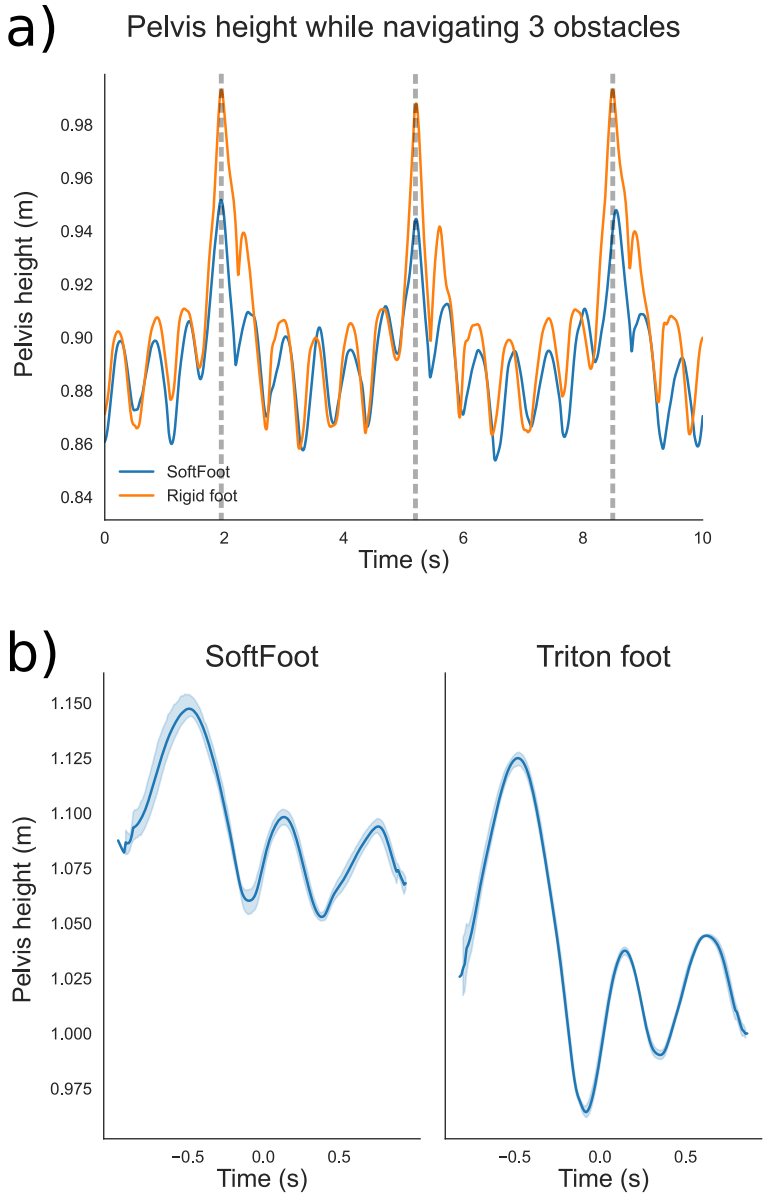


Figure 6.9: Height of the pelvis when tackling obstacles with an prosthesis. *a)* Results from simulation, dashed lines indicate the mid-stance when stepping on the chord-type obstacle. *b)* Corresponding experimental results from IIT for descending from a single obstacle. Graph starts with the contact with obstacle, repetitions aligned for $t = 0$ at the mid-swing of the prosthesis side, determined by the minimum knee angle. Two σ interval shown. The difference between off-obstacle mean height can be attributed to inter-session motion capture marker placement changes.

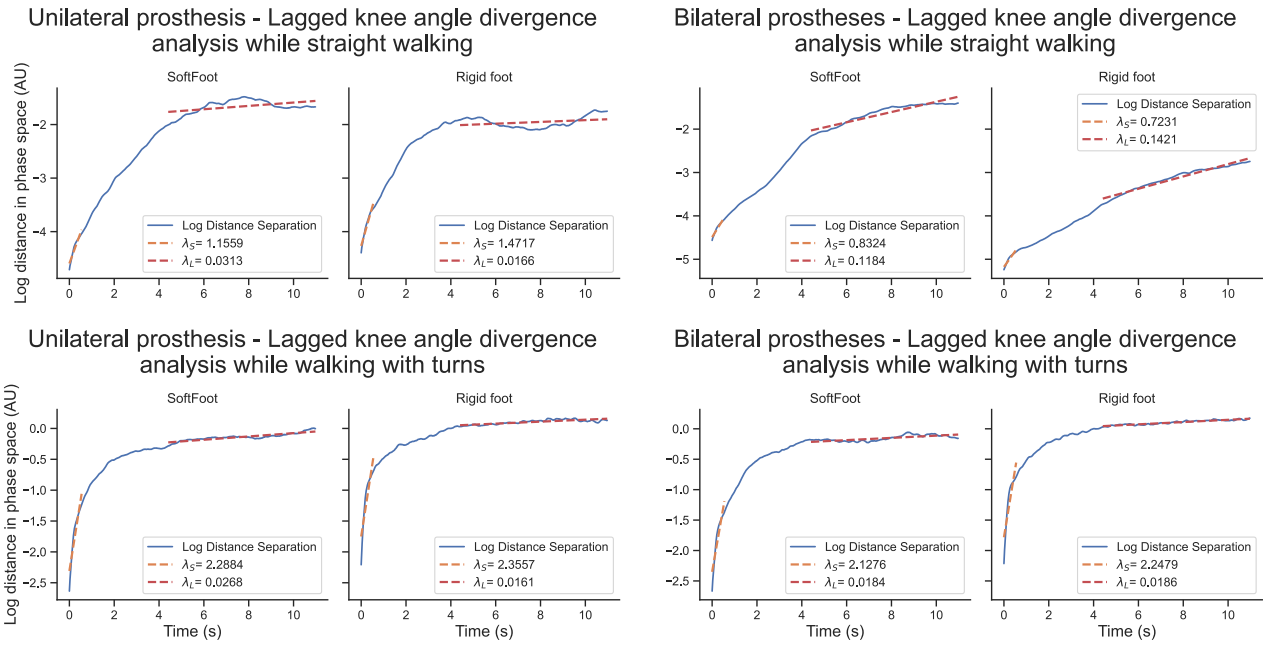


Figure 6.10: Maximal Lyapunov exponent approximations by fitting least-squares linear regression to short-term (λ_S) and long-term (λ_L) sections of the divergence curves, calculated for the SoftFoot and rigid foot model.

Absolute symmetry index

The mean absolute symmetry index (c.f. Equation 6.2) was lower in magnitude for the SoftFoot ($-0.30 \pm 2.01\%$) than for the rigid foot ($-7.05 \pm 4.38\%$). With a two-tailed t-test the rigid foot stance phase % distribution was significantly different than the intact side's ($p < 0.001$). The compliant foot did not exhibit a significant difference from its contralateral side ($p > 0.02$). We show the distributions bilaterally for the two foot conditions in the form of histograms in Figure 6.11.

Perturbation trials

Both rigid foot and SoftFoot policies can reliable perform level ground locomotion with turns while walking on the uneven terrain proposed by Pace (Pace et al. 2023) (see Figure 6.5) without falls, despite not including it during training. Rotation impulses were applied as perturbation to the humanoid as described in Section 6.2.5. After 30 such episodes we averaged the durations of walking bouts without fall events. The rigid foot agent walked without a fall for a mean of 31 seconds, the SoftFoot agent took 45 seconds before a fall.

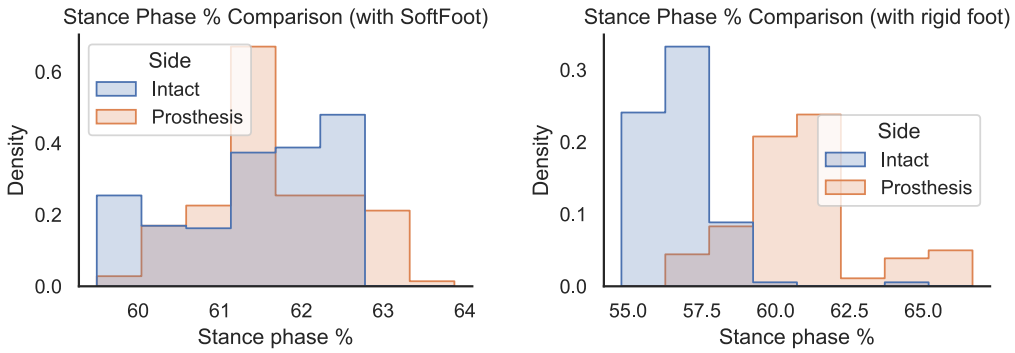


Figure 6.11: Histograms of the percentage of strides spent in stance phase for the prosthesis and contralateral sides for the two types of simulated prostheses investigated.

Lastly, we report the pose error as in Chapter 5's Equation 5.7 in Table 6.2.

Table 6.2: Pose error during straight walking and level ground walking with turns, calculated using Equation 5.7. Quantified for both SoftFoot and rigid foot useage, for unilateral (UL) and bilateral (BL) prosthesis use.

Condition	Pose error (m) with SoftFoot	Pose error (m) with rigid foot
UL straight	0.023 ± 0.007	0.026 ± 0.012
UL turning	0.044 ± 0.026	0.046 ± 0.024
BL straight	0.031 ± 0.017	0.041 ± 0.012
BL turning	0.040 ± 0.022	0.055 ± 0.024

6.3.2 Comparison with experimental data

On Figure 6.12, we visualise kinematics and dynamics when the virtual participant stepped on the chord-type obstacle introduced by Pace (2023) (also see Figure 6.13). In the top row we show the angle between the prosthesis' longitudinal axis along its sole and the ground. In both simulated and experimental SoftFoot, the early stance has a lower rate of change in this angle, as the foot conforms around the obstacle with multiple points of contact. The non-compliant devices instead roll on the surface of the obstacle's arc. In the case of simulated rigid device we observe a characteristic "wobble" in the early stance as the foot rocks back and forth once while the swing foot transitions forward. The dynamics of the knee is clearly different than the

inverse dynamics results from the experiment, only sharing coarse features. In particular, early stance knee extension is lower in our models.

6.4 Discussion

A key observation must be made regarding the simulated devices: While the behaviour using the SoftFoot resembles in multiple aspects the experimental data, the simulated rigid foot differs to a greater extent when compared against the experimental results using the non-compliant foot. The amputee participant used an Ottobock Triton device, which has a carbon fibre structure with energy storage and return mechanisms, tuned for efficient walking. Our simulated rigid foot lacks these properties. Comparing against the rigidified version of the SoftFoot emphasises the effect of its compliance, without introducing any other mechanisms or changes to the geometry. However, if we intend to predict the potential benefit of a device may have on the end user, we must implement models of existing alternatives available to them. For example, the energy storage and return capavilities of the Triton can potentially explain the reduced peak torques compared to the SoftFoot. Incorporating a way to replicate this function in the SoftFoot is therefore an interest. Future studies in this line could greatly benefit from an open source repository of validated P&O hardware which can be combined with simulated scenes in a modular way, similar to existing collections for robotic and musculoskeletal systems (Zakka et al. 2022; Delp et al. 2007). This effort may require collaboration from P&O manufacturers, to obtain permission to release potentially protected intellectual property, as well as to reduce the burden of reverse-engineering device properties.

The dynamics of the knee differ considerably between the simulated agent and the one estimated from experimental results. Our forward dynamics agents can adopt a range of different torque profiles that integrate to similar kinematics, and the inverse dynamic process is also sensitive joint centre definitions and choices of the humanoid model (Hicks et al. 2015) fitted for the subject. Adopting musculoskeletal humanoid models (perhaps by extending torque actuated policies like ours (Lee et al. 2019)) could help more realistic torque profiles emerge, as

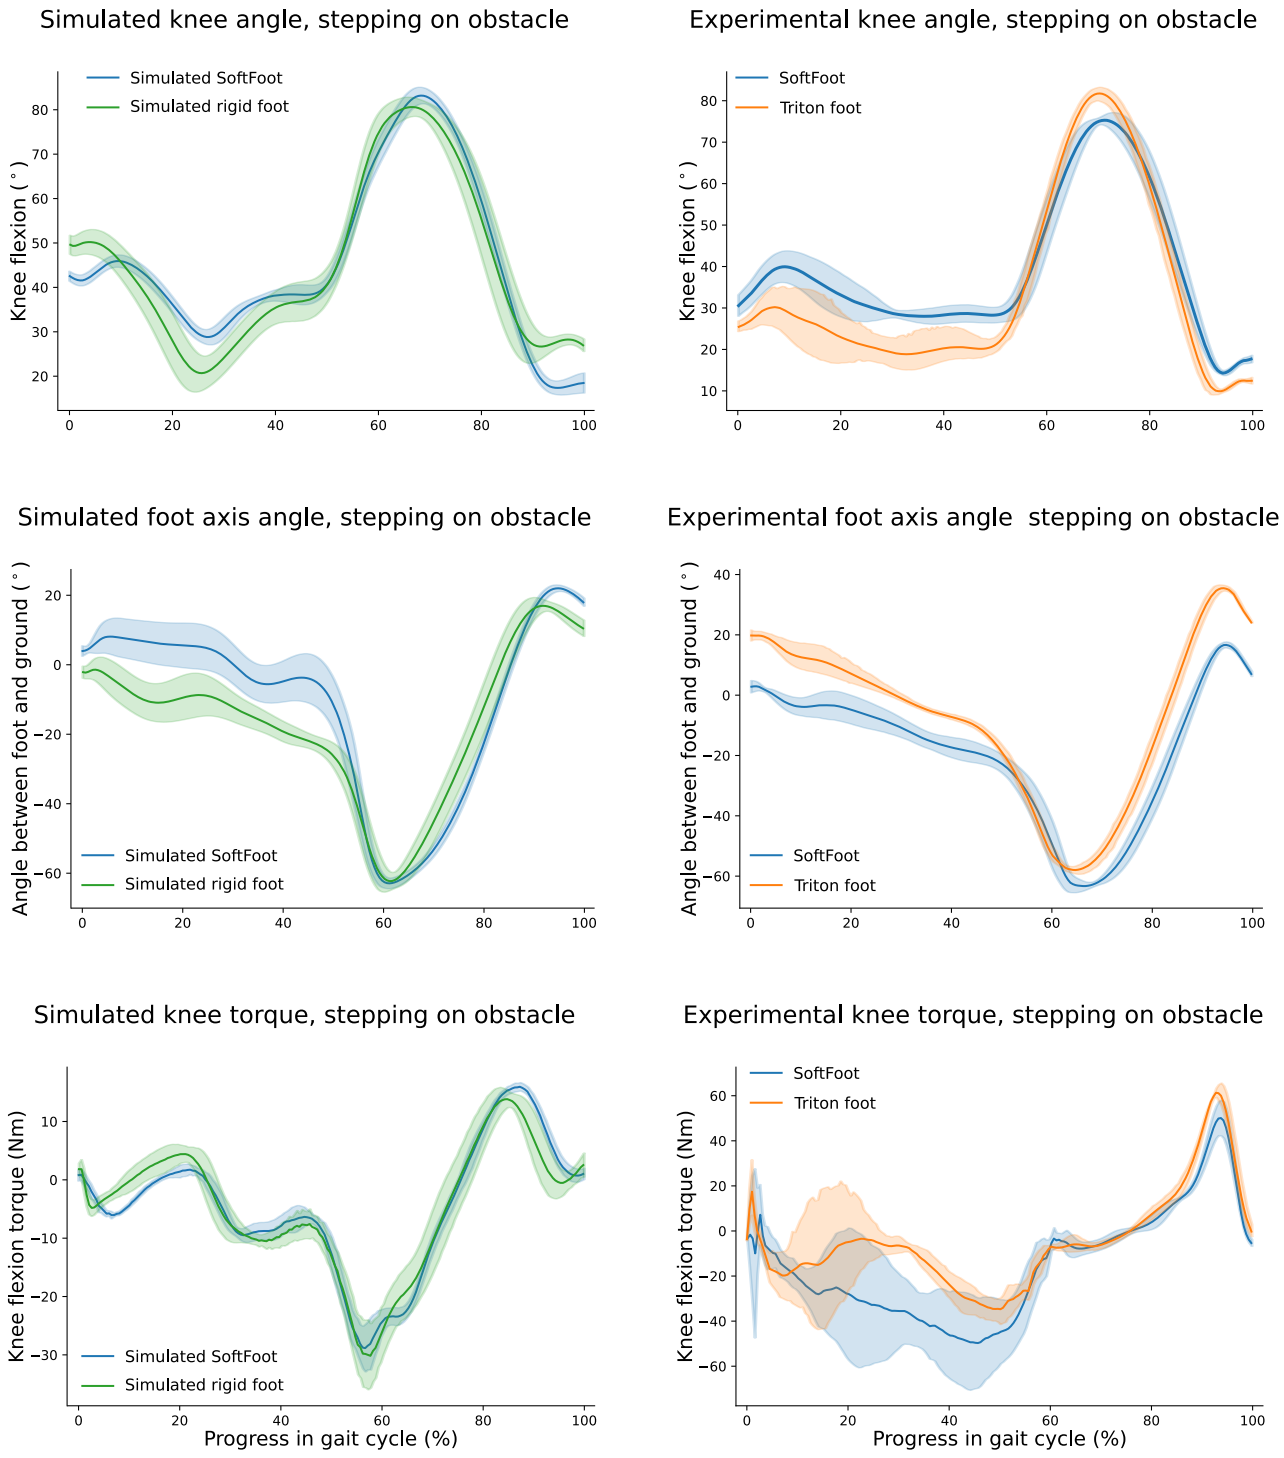


Figure 6.12: Foot and knee kinematics and dynamics compared between the simulated SoftFoot, simulated rigid foot, experimental SoftFoot and experimental Triton foot data when unloading from the chord-type obstacle. Experimental data provided by IIT and Pace (2023), determined using inverse kinematics and dynamics from motion-capture and load-cell data. All plots are aligned for $t = 0$ at the mid-swing of the prosthesis side, determined by the minimum knee angle. Two standard deviation interval shown. For an illustration of the motion see Figure 6.13.

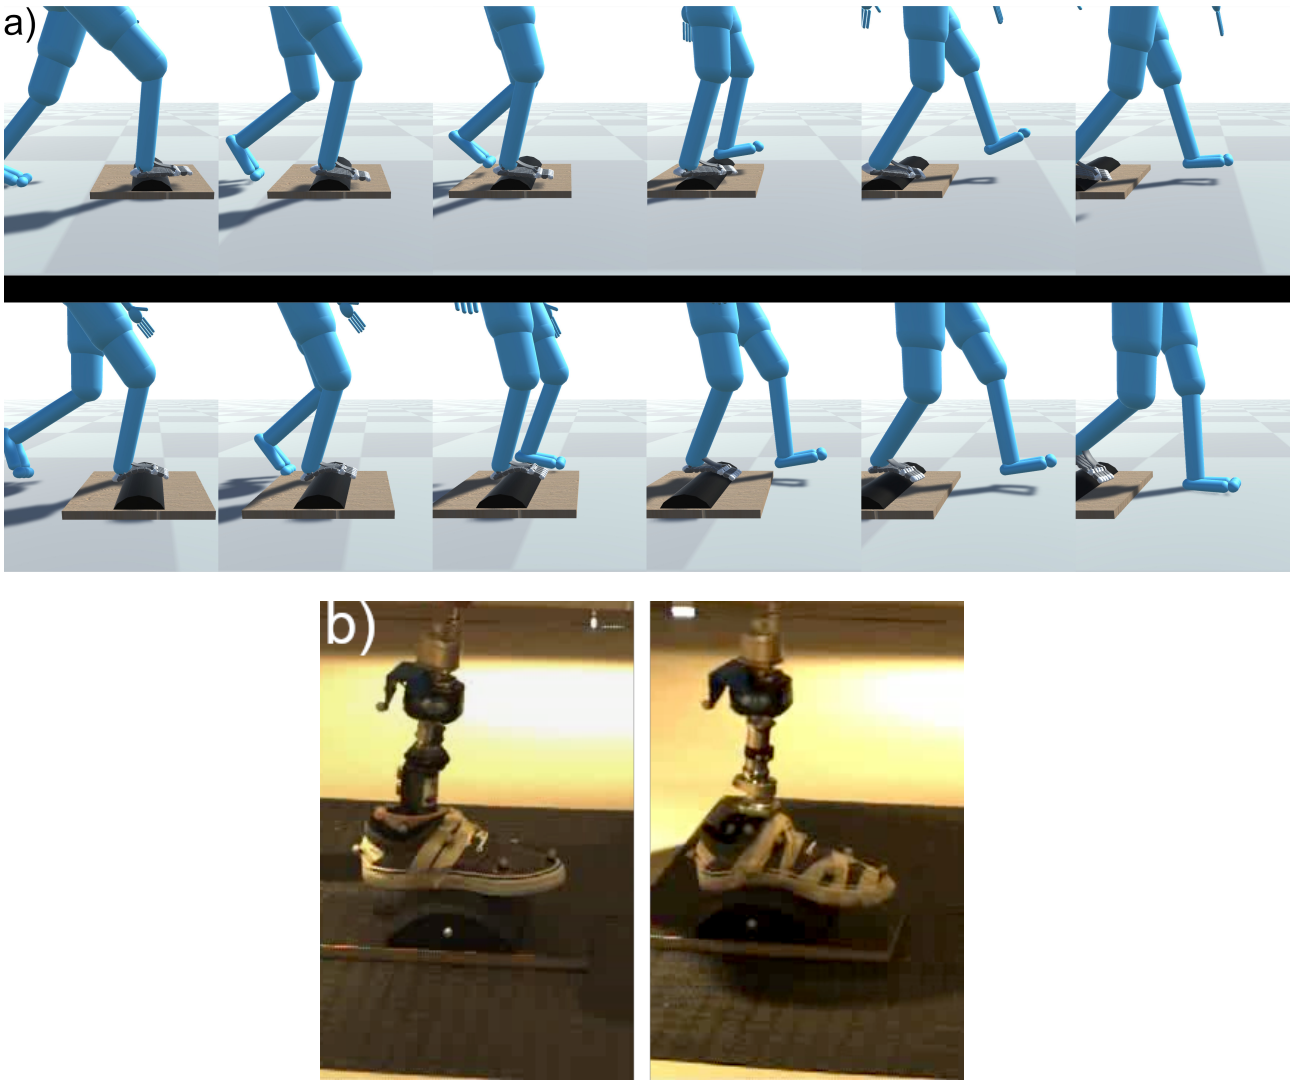


Figure 6.13: *a)* The simulated rigid foot and the SoftFoot stepping on the chord-type obstacle. *b)* Stepping on the same type of obstacle with the Triton foot and the SoftFoot during experiments.

the nonlinear, stateful and overactuated nature of muscle actuation can significantly influence resulting behaviours (Kidziński et al. 2020).

The variability of the knee angle kinematics in Figure 6.12 was lower than in the experimental data. In general, despite having access to many repetitions variations of gait cycles, motion matching has the possibility to converge to a single cycle of motion if the intent signal does not change. For an example, see Figure A.4. A potential way to mitigate this could be to penalise matching to frames recently used to incentivise using the full dataset, or to vary the input trajectory with a plausible low frequency signal (e.g. modulating the target walking speed or direction to a small extent)

The ability to generate testing equipment, such as the obstacles or slopes at no additional costs is a further benefit of virtual gait labs. Through this process, experimental setups can be prototyped as well. Alternatively, complex locomotion environments not easily feasible in real gait labs may be used for evaluation. Our sloped-locomotion agent was also tested on a previously unseen environment where the heightfield data was dynamically changed and moved with a pattern generated using Perlin noise. Even on constantly shifting surfaces the locomotion agent generated walking and suitable recovery movements to maintain balance, despite only seeing static height fields during training (Figure 6.14).

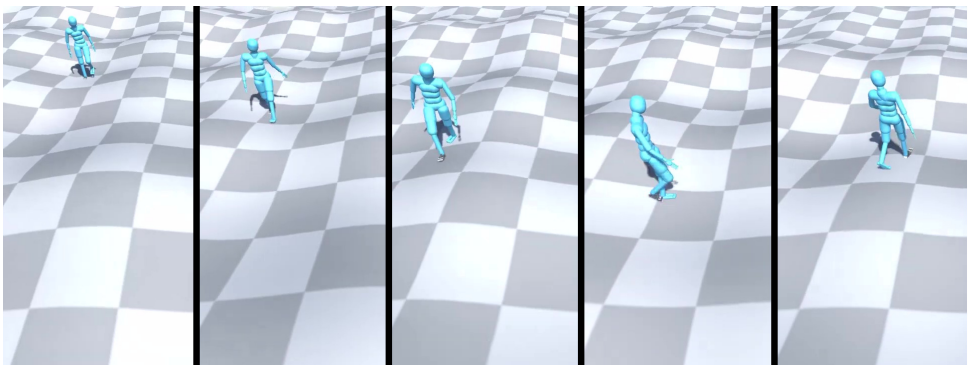


Figure 6.14: Evaluating the locomotion agent on a dynamically shifting height field environment.

We quantified key gait stability metrics such as the maximal Lyapunov exponent and gait symmetry in various walking conditions. We saw that λ_s, which is correlated with gait stability, appeared lower in walking tasks involving turns. A potential reason why λ_s was greater for

the rigid foot for straight walking could stem from the stereotypical nature of the generated straight walking patterns. Due to the lower variability compared to turning behaviours (see Figure A.4), the gait policy may be able to find a robust solution with the simpler foot. Once non-steady-state behaviours are involved with a greater potential for a loss of balance, the higher λ_s with the SoftFoot may indicate benefits it could provide thanks to the compliant toes and metatarsophalangeal joints. In the case of the symmetry index, the SoftFoot exhibited more comparable strides on either side than the rigid foot. However, we note that the sign of the symmetry index is unexpected in particular for the rigid foot, as more time was spent on the prosthesis side than on the leg under the control of the locomotion agent. This is in contrast to expected results from experimental studies, where it is suggested that optimal policies will spend more time on the intact side, where the device user has greater control over the stance (Nolan et al. 2003). This could once again indicate that the rigid foot is an inappropriate representation of commercial prosthetic feet, or could suggest that our heavily simplified, torque actuated motion tracking locomotion agent does not fully capture the constraints and goals of prosthesis users. It is possible that musculoskeletal actuation is needed to be implemented to better reproduce the extent of control and actuation the agent has over its body. Another potential improvement could be to limit the amount of proprioceptive information available on the prosthesis side, which may also induce the expected asymmetry. Lastly, altering the RL cost function to promote lower muscular effort solutions, or to penalise large loads on the residual limb could also be pursued for this purpose.

While these are useful measures, they fail to account for the responsive and adaptive nature of closed-loop control policies of humans, which may alter behaviour in case stability is compromised. Bruijn et al. (Bruijn et al. 2012) comment on this, claiming that "the probability of falling will thus not only depend on [the subject's] steady-state gait stability, but also on the quality of those recovery actions, which is not necessarily correlated with the stability of the gait pattern." This highlights how a virtual gait lab can complement experimental studies by testing situations and tasks not feasible in the real world.

6.4.1 Ablation studies

In our humanoid control we use the motion generated by motion matching as baseline, feedforward targets for PD actuation, which are then modulated by the RL policy to maintain stability. It is possible to instead rely on the control policy entirely to learn the absolute PD targets on their own (Peng et al. 2018). However, this is generally associated with slower learning (Peng et al. 2021). To test our assumptions of the benefits the feed-forward motion targets provide to the locomotion learning, we performed ablation studies in straight walking, level ground walking (with turns) and sloped walking settings. We compared the learning curves of non-prosthesis user humanoid control with and without the feedforward component (Figure 6.15). From these results we confirm that both early and asymptotic performance is positively affected with the inclusion of the feedforward control in our environments. In the case of straight walking, we also note a considerable reduction in the variance of asymptotic performance, where RL-only control could be prone to get stuck in local optima. On Figure Figure 6.15 we also see a phenomenon we observed multiple times during training, periods where the performance of the policy suddenly decreases, before recovering. We investigated what could cause this, without uncovering a definite cause. A potential source of it could be the scheduled adaptation of some of the hyperparameters, such as the learning rate or entropy regularisation in the ML-Agents package, which we used to manage the RL process (Juliani et al. 2018). These are supposed to be adapted to drive the network to converge to a near-deterministic solution, but as the learning and action dynamics change, it may also require additional learning from the policy to suit them.

We also included in the ablation study conditions for the ground reaction force (and ZMP) observations. Interestingly, this added observation only seems to provide marginal benefit, possibly the information is redundant as it may be possible for the agent to infer the ZMP from its existing proprioceptive observations.

We also tested whether adapting the feedforward motion using the footbase method provides a benefit. On Figure 6.16 we compare two agents using feedforward control, except one of them adapting the motion matched animation to the terrain with the procedural inverse kinematics

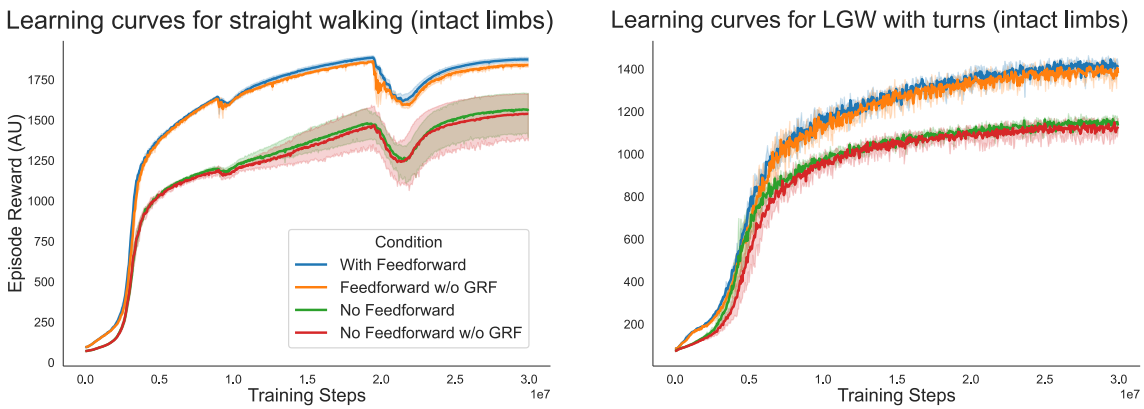


Figure 6.15: Comparing the effect the feedforward motion target and the ground reaction force (GRF) observations have on learning motion skills with RL. *Left:* Learning a continuous straight walking from a looping gait cycle animation. "Policy collapse" phenomenon can be observed later during the learning. 5 repetitions used for each condition, 1 standard deviation interval shown. *Right:* Training policies in a Level Ground Walking (LGW) setting involving turns with the same four conditions.

methods described in the previous section. Both agents start from the same baseline policy learned on level ground locomotion, placed in a new environment with the sloped terrain shown on Figure 6.6. The agent without the procedural motion synthesis had to learn appropriate corrections to the motion on its own. The agent with terrain adaptive feedforward motion had much higher performance both without training and after adaptation. We note that the kinematic motion synthesis serves both as the feedforward control and as the target for the reward calculation. As such, the best possible reward the agent can get for following the level-ground walking reference will be limited already even if the feedforward motion would not be used. Nonetheless, these results indicate that the additional processing introduced into the animation pipeline leads to better performance and more appropriate evaluation.

6.4.2 Future work

Our implementation of the procedural animation methods of Johansen (2009) benefited from many simplifications thanks to combining it with motion matching. However, we were limited by our motion capture dataset not including short stride length motion, due to which we could not generate plausible continuous stair ascension. This is a clear area for future improvement, addressed either through additional motion capture, or further procedural motion editing

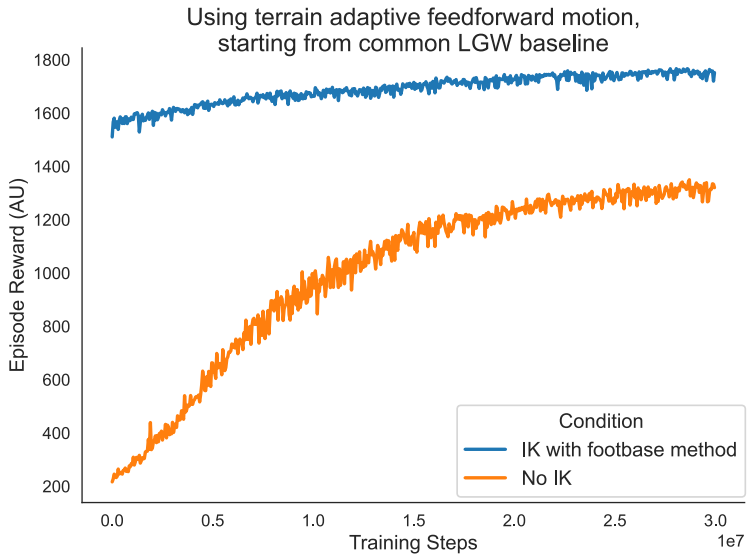


Figure 6.16: Including the inverse kinematics (IK) pipeline for adapting the level ground walking to the sloped environment shown in Figure 6.6 increases zero-shot and asymptotic performance.

techniques.

Further analysis of the effect of the mechanical interface model between user and device needs to be carried out. Not only can this lead to more accurate models of prosthesis use, it may have applications for embedded systems that could detect and warn users regarding a bad socket fit or alignment, developed in simulation. Additionally, considering how to include the effects of wearing a shoe over the prosthesis must be studied in the future. This should be done together with the creation of validated recreations of commercial devices.

Extending our learning environments to musculoskeletal implementations of the virtual humanoid is a key priority. This is particularly true if facets of prosthesis use such as metabolic cost, fatigue or muscular effort is to be investigated in the future. These are key concerns in P&O use, and accounting for them is the already the standard in simpler locomotion settings (e.g., straight stead-state walking) (Hicks et al. 2015; De Groote et al. 2021).

Our collaboration with IIT is ongoing, pursuing multiple goals:

- We are continuing our efforts to system-identify the SoftFoot to produce more accurate simulated models of the compliant mechanisms.
- We are investigating how to alter the humanoid model, and the RL cost function to create

locomotion policies that better reflect the real human-prosthesis system. For this reason, we are studying what further metrics and measures we can quantify from the simulation to highlight key limitations and discrepancies (as the symmetry index has done).

- We are exploring trialling the SoftFoot with an active bionic ankle inside our virtual environments, with different actuation and control strategies.

6.5 Conclusion

This chapter demonstrated how our dynamic locomotion agents can be used to explore the effects of design choices in the mechanical composition of a passive prosthetic foot. We extended our motion synthesis system to tackle varying slopes and anticipate obstacles. The gait resulting from stepping on an obstacle was compared between a rigid virtual foot and the compliant SoftFoot, and contrasted them against inverse kinematics and dynamics outcomes from an experimental study. Similarities in locomotion outcomes were observed, however, we have also highlighted key limitations and opportunities for future improvements in the models of the foot, the user and the interface between the two. The proof-of-concept virtual gait lab we have shown indicates that virtual environments could be used for prototyping lower limb P&O hardware prior to patient participation, which could lead to a better utilisation of the available resources for experiments in the future.

Chapter 7

Conclusion

Simulations and offline evaluations will not eliminate the need real users have for real P&O hardware. In the end, these devices are meant to fill biomechanical deficits impeding people from reaching functional and psychosocial goals during everyday life, outside of a virtual environment. However, there are countless ways to attempt to fill these needs. Models of P&O use help us make sense of our options and select an appropriate one, and play an instrumental role in realising novel P&O interventions.

Despite the diverse range of intent-driven locomotion tasks and prosthesis conditions investigated, our models are still coarse. They made a great number of assumptions and simplifications in both the neuromechanical model of the human, the structure and control of the prosthesis and the interface between them. This PhD could be viewed as a single iteration in the design cycle described by Hicks (2015). We defined research questions about modelling variable locomotion intent, prototyped new methods and tools to explore them, tested the robustness of our simulations and contrasted them against observations from the real-world systems they aimed to represent. The developed methods can be now leveraged to perform another round in this cycle with greater confidence, increasing the level of detail throughout the process.

7.1 Summary of thesis achievements

In Chapter 3 we studied the use of TCNs as simple but powerful learning-based tools for estimating motion intent during walking. Importantly, we identified and quantified failure cases that these otherwise promising methods are sensitive to, focusing on spatial electrode shifts. We proposed a straightforward but novel approach of randomly sampling bipolar sEMG channel combinations from a high-density grid during training of a deep learning model, which results in models that are able to generalise across a relatively large area.

Building on this, in Chapter 4 we adapted our intent estimator to predict the walking path trajectory during non-steady-state locomotion tasks involving various turns. This shift to mapping muscle activity and body posture to a higher-level representation of user goals meant more diverse tasks could be tackled. Furthermore, by modelling the human motion trajectory with a critically dampened system, the behaviour could be described in a low-dimensional parametric form.

Through this parametrisation, HL user intent could be synthesised in a virtual simulation of the locomotion task. In Chapter 5, we implemented a full body motion controller method which decompressed the intent representation to synchronised reference kinematics, which served as an imitation target for a humanoid simulated dynamically with a physics engine. Then a closed-loop control policy for the simulated participant was found using RL. By keeping this policy frozen, and introducing a second adaptive agent controlling a unilateral powered prosthesis, we demonstrated that only in our intent-driven prosthesis conditions could a reliable policy be learned during non-steady-state gait.

We leveraged the developed environments to serve as a virtual gait lab for modelling and evaluating a simulated version of the SoftFoot, IIT's compliant prosthesis design. Chapter 6 described the virtual SoftFoot's blueprint and our approach to modelling the mechanical interface between the residual limb and the device. We added varying slopes and steps to our roster of environments, and added a procedural inverse kinematics pipeline to adapt our reference motion to them. We compared gait control policies optimised for the compliant

SoftFoot and its rigid counterpart, and differentiated them in terms of their response to stepping on an obstacle.

Our methods demonstrate a virtual prototyping framework for both active and passive lower limb prosthetics for low-cost iterations and development of design ideas. The tools we developed for creating our learning environments were released and documented in the form of an open source package for Unity and MuJoCo (Hodossy et al. 2023b).

7.1.1 Understanding locomotion intent

A core question addressed by this thesis was how to model wearables that intuit the desired motion outcomes of their users. Lacking the biosignals that HMIs are usually driven by, we answered it by adopting an interpretable and low dimensional representation, which essentially describes *"where am I headed?"* and *"where would I rather end up?"*. The latter corresponds to information that goes beyond simple extrapolation of the current kinematics. As it depends on a free variable (the desired velocity), a wearable device can not deduce it without interpreting its user's intention through some channel like a myoelectric interface.

Through this, we gain the ability to differentiate and select walking behaviours from a large set of potential walking paths that involve turns, stops and walking at different speeds. However, it appears we are losing a lot of detail. Gait involves the coordination of dozens of muscles, thousands of motor units and unique kinematic and kinetic patterns for each of our DoFs. Can we really hope to seamlessly assist our walking without understanding all its complexity?

Part of the reason we can afford to be ignorant of the specifics is the high amounts of redundancy built into our movement. Different configurations of joint angles, muscle activations and neural signals can result in similar mechanical interactions with our environment (Burdet et al. 2013). This increases the robustness of our bodies' control, but also renders some of the underlying details redundant for wearable robotics. The cyclic and typical nature of our movement can also be exploited, which explains how a single channel sEMG could be used to predict continuous joint angles under controlled conditions in Chapter 3.

This however, does not mean that the aspects not captured by our intent representation are irrelevant. The estimated high-level intent may condition wearable device control like in Chapter 5, but assumptions must be built up about gait patterns not captured by HL intent. In our case this is done through RL, by collecting experience while sharing autonomy with the human agent, which slowly builds up the baseline behaviours that can be modulated by the intent signal.

7.2 Future work and potential applications

7.2.1 Robust bipolar sEMG interfaces

In addition to electrode position shift, there's a wide range of other sources of variability that myoelectric controllers need to be robust for. We briefly explored the effect of shifts in the walking speed (Figure 7.1). We observed non-uniform warping of the prediction's phase, indicating that the TCN still extracts useful information, but does not readily generalise to changes in the timing of muscle activations. Similarly as for electrode placement shifts, data augmentation methods could be pursued to mitigate this, such as dynamic time warping. Other factors such as fatigue and varying noise levels, and multi-subject models should be investigated. We applied the methods of Chapter 3 in a preliminary test to try this concept. For the Soleus, TCN knee angle estimator models trained on the recordings of Subjects 2-5 performed at $r^2 > 0.8$ on Subject 1's data. On the Tibialis this transfer was at $r^2 > 0.2$, while for the thigh muscles transferred models had negative r^2. This indicates potentially varying suitability of muscles for multi-subject generalisation, which should be established with a follow-up study.

The transfer of gait predictors across acquisition systems is a promising but highly challenging problem (Figures 3.11 and 3.13 illustrate this). However, to further develop the spatially robust TCN techniques of Chapter 3, this must be addressed. Progress towards this was made recently, exploring domain-adaptation techniques to improve transfer (Guez et al. 2023).

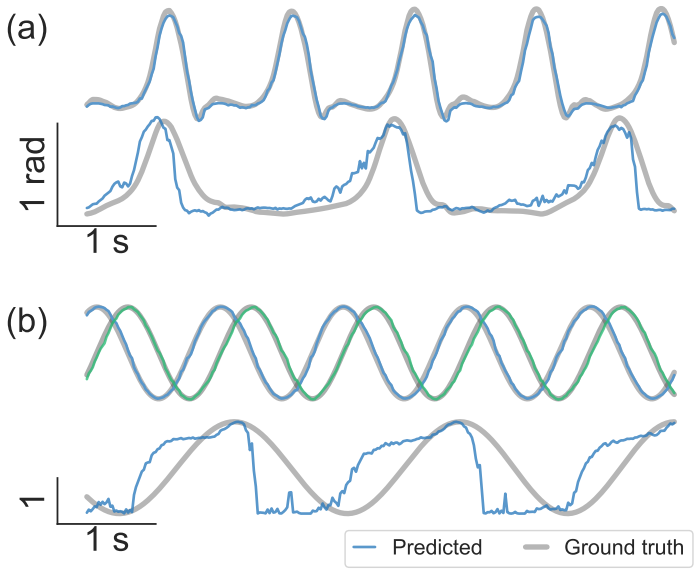


Figure 7.1: *a:* Training and predicting with a steady state knee angle estimator TCN for 0.8 m/s walking speed (top row, validation data), and predicting on a dataset of 0.3 m/s (bottom row). *b:* Estimating the two dimensional gait phase (as defined in Equation 3.2) instead of the knee angle, trained for 0.8 m/s (top row, validation data), and transferring to 0.3 m/s (bottom row, only first component shown for clarity). Subject 5's Soleus middle electrode data was used.

7.2.2 High-level intent estimation

Our work in Chapter 4 shows a promising connection between sEMG signals and HL locomotion goals, exploiting the coupled contributions of muscles preparing the body for the upcoming movement. The next steps in further studying this involve trialling estimators online, as well as addressing the usual concerns for biosignal based models for robustness already discussed in the previous subsection. The limits of how far in the future can we predict motion trajectories needs to be determined, as in online systems additional delays and processing times need to be accounted for (Smith et al. 2011). For this purpose, reducing the complexity of the network through pruning method and hyperparameter optimisation should also be carried out. Architectural improvements could be discovered by studying stateful predictors like recurrent networks, which could help stabilise our highly variable estimates.

7.2.3 Motion synthesis

Our implementation of motion matching followed closely on the information available from Bergamin (2019) and Clavet (2016). Other variants exist, such as learned motion matching (Holden et al. 2020), and could be used to supply more diverse motion types, sampled from a greater amount of motion capture data. For example, stair climbing can be included directly in the motion matching, instead of procedurally adding it using footbase trajectories (Section 6.2.2). Walking at a greater range of speeds, strafing, backwards steps and running are further behaviours of interest.

In the previous chapters we have only used the motion synthesizer to aid the motor control problem of the humanoid agent. The ability to generate full-body motion also has important applications in the control of the wearable device:

- *Full body pose extrapolation from sparse information (Reda et al. 2023):* Having limited measurements of the configuration of the user's body is a major hurdle in the controllers of wearable or assistive robotics (Martin et al. 2017; Maroger et al. 2021). A motion synthesizer like motion matching could fill in the gaps for the model of the humanoid.

- *Feedforward control for robotic P&O:* It is not uncommon in many settings for the device control to follow a predetermined or partially modulated trajectory, and expect the user to compensate with their remaining agency to maintain balance (K. He et al. 2023; Shushtari et al. 2022). The versatility and expressivity of motion synthesizer could help map changes in parameters to the detailed joint kinematics. We explored this idea in a simulated hip-knee-ankle prosthesis control scenario, where the target trajectories for the prosthesis were provided from the kinematic recordings of a second subject, matched dynamically to the main subject's movement.

7.2.4 Humanoid locomotion model

Removing the kinematic feedforward model

The actuation model we used relied on a baseline, feed-forward motion prediction generated through the motion matching method. Thanks to it, models could be trained much faster and to a higher capability (see Figures 6.15) and can be partially extended to novel tasks (such as adapting to stepping environments). However, this feed-forward baseline may prevent learning behaviours significantly diverging from the motion matching dataset, and requires additional computational work during inference to calculate. Furthermore, motion-tracking necessitates the careful definition of an imitation reward function. An alternate approach, adversarial imitation learning works by providing example motion features from a reference dataset, and the locomotion agent has freedom in generating behaviour as long as its resulting distribution of features is not too different from the reference (Peng et al. 2021; Merel et al. 2017). As such, no manual definition of the imitation reward function is needed. These methods were explored as part of this project, in a policy transfer setting from motion tracking to accelerate the learning process. Description of this experiment is presented in Appendix D. Additionally, adversarial imitation learning have been shown compatible with a high level control that learns its own representation space, instead of relying on manually defined ones (Tessler et al. 2023). Finding optimal intent representations are highly relevant for P&O device control, and the conditional latent models such as the ones described by Tessler (2023) could have important implications for this.

Musculoskeletal actuation

A torque actuated model was selected due to the drastically cheaper simulation cost and easier to learn action space. This reduction in model complexity was made to suit our research questions. Our humanoid controller's goal was to provide plausible kinematic and kinetic context to the prosthesis, where unlike with exoskeleton use, the human agent does not directly actuate any of the controlled DoFs. However, when appropriately configured, musculoskeletal human

structure can lead to the emergence of naturalistic gait patterns without necessitating any imitation learning or reliance on example motion capture (Song et al. 2015; J. Park et al. 2022). Although, often even with musculoskeletal models, imitation learning is necessary to enable meaningful learning of skill (De Vree et al. 2021; K. He et al. 2023; Lee et al. 2019). Removing imitation learning could, for example, empower the locomotion agent to pursue asymmetric gait patterns. Instead of trying to replicate walking with intact limbs (K. He et al. 2023; Lee et al. 2019), virtual subjects could be unrestrained in exploring strategies that reflect the neuromechanical differences between their sides. Despite gait symmetry being often pursued as indicator of healthy prosthesis use, asymmetry has been identified as being potentially optimal with prosthetic locomotion (Handford et al. 2018). Musculoskeletal humanoid models could also enable the more accurate investigation of energy expenditure, which is not appropriate to perform with torque actuated models (Kidziński et al. 2020). Therefore, replicating the methods presented in this thesis with musculoskeletal and more detailed neuromechanical modelling is of high priority. These models must be validated through comparisons with appropriate data from neuromechanical experiments, and tuned to explain to observed results (Hicks et al. 2015). Beyond experiments with living subjects, studies that employ cadaveric specimens could be used to uncover key mechanical properties of joints, tendons and muscles, otherwise unavailable to measure directly with non-invasive methods in vivo (Aubin et al. 2012).

Musculoskeletal modelling may also offer a way to reach another ambitious goal, end-to-end HMI simulation. Sensor units such IMUs or load cells can be readily characterised and represented in simulation. However, the synthesis of complex non-stationary biosignals like sEMG is a challenging research topic in its own right. There is significant need for it; in its absence many virtual prototype device's myoelectric interfaces (including the ones presented in this thesis) rely on abstract surrogate signals available in the simulation (Bustamante et al. 2021; Chappell et al. 2022; Gloumakov et al. 2022). The activation level of Hill-type muscle models provides a more physiological format to model the state of neural interface (Chadwick et al. 2009; Caggiano et al. 2022). This still requires the implementation and integration of an additional layer outside of simulation that extracts this measure. There are novel learning- and neurophysiological modelling-based methods being developed that propose to synthesise EMG

signals in silico (Ma et al. 2024b; Ma et al. 2024a; Jiang et al. 2008). Through systematic randomisation of neurophysiological parameters and electrode placements, a wide distribution of plausible EMG signals could be generated by musculoskeletal virtual subjects. Generalist HL intent estimators that are robust against these variations could be trained wholly in simulation, and their output could be used to control assistive robotics that influence the motion and the muscle activity that generates it in the first place. This is in contrast with current applications that generate EMG in an offline manner, processing but never influencing motion (Ma et al. 2024b).

Lastly, with musculoskeletal control policies it would be possible to explore our non-steady-state locomotion tasks for exoskeleton use, providing more versatile evaluations of these devices in simulation.

7.2.5 Architectural changes

The motor control policy used in Chapters 5 and 6 used a classic feedforward neural network. Many other architectures could be investigated for prosthesis user agents, such as:

- Encoder-decoder systems (K. He et al. 2023) that perform dimension reduction and enforce learning control synergies. Alternatively, subsystems that map the desired net joint torque from policies like ours to a musculoskeletal models, distributing the muscle activations appropriately could be further investigated (J. Park et al. 2022). Without suitable constraints, RL strategies can often lead to erratic and non-phyisiological muscle activity (De Vree et al. 2021; Caggiano et al. 2022).

- Long Short-Term Memory, TCN and causal transformer architectures, that can perform not only long term learning to static conditions, but to recognise and adapt on the shorter timescales to changing environments (Radosavovic et al. 2024; Akkaya et al. 2019). For example, if the agent equips a weighted backpack or starts walking on surfaces with less friction, agents with memory capabilities can compare their experience to a learnt internal model, and extrapolate and act accordingly adapted to the new situation. This way not

every detail that may change in the environment needs to be parametrised and provided to the agent to condition its behaviour, but instead the agent can infer it from its existing set of observations.

- Multi-subject / multi-device models. Data-driven methods are incentivised to exploit any possible advantage they can find in an environment, therefore any policy trained with a single set of anthropometrics, device and environmental configuration will be specific to that set. In a different context, we tackled the same issue in Chapter 3, where we show that by randomly sampling electrode placements, a single model can robustly work for many conditions. Domain randomisation can also be applied for motor control tasks, randomising anthropometric, kinetic or inertial parameters of the environment (Radosavovic et al. 2024; Akkaya et al. 2019; J. Park et al. 2022; Lee et al. 2021), or warping the shape and timing of the target motion (Lee et al. 2021). Through this, a single motor policy could be reused for many environments, removing the costly requirement to retrain for new conditions. Even more importantly, if a single assistive device's controller can be used for a wide range of virtual users, the potential for zero-shot transfer to a real user's device is greater. The domain randomisation method was recently shown feasible for a hip exoskeleton controller's transfer from simulation to real life (Luo et al. 2024). This indicates that this approach is perhaps even more promising than creating digital twins of the real user to train within the simulation. We have already developed tools to procedurally scale the virtual humanoid, which is necessary for these future steps (Figure 7.2).

As DRL is far from the only method to create physically simulated, dynamically stable locomotion agents, developments from alternative approaches such optimal control methods should be also considered and further reviewed (Nguyen et al. 2019). The implementation of concepts from these fields could benefit progress in RL agents, and likewise the other way around.

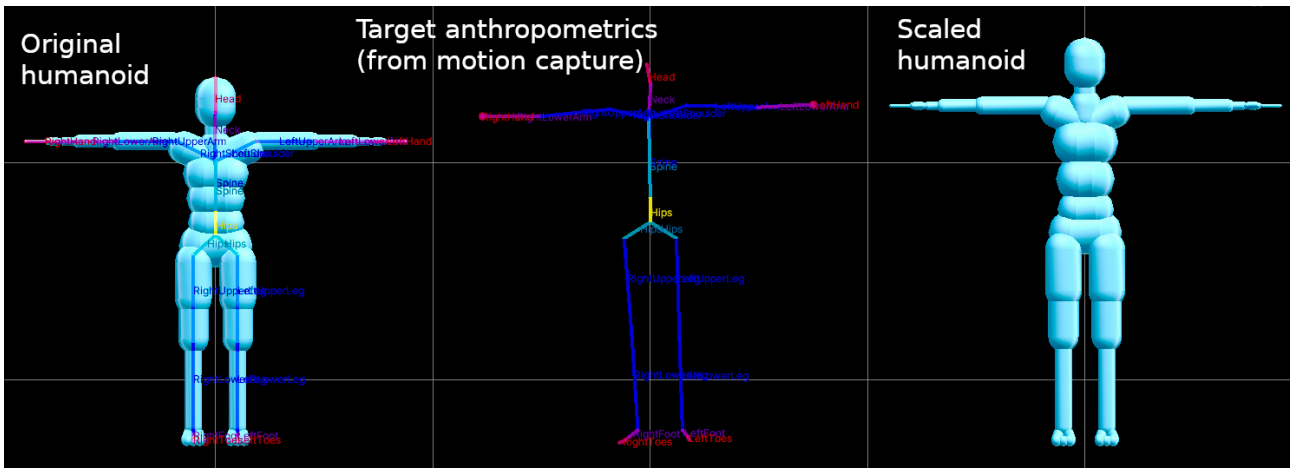


Figure 7.2: Humanoid model before and after adjusted to match new anthropometrics using scaling tools released in the Modular Agents package (Hodossy et al. 2023b).

Simultaneous multiagent learning

Another aspect related to the training of the locomotion policy is how the multiagent nature of coadaptation is handled. In Chapter 5 we show that radically different learning dynamics can arise if multiple agents are learning and interacting in parallel (Zhang et al. 2021). In simulation this can be resolved by freezing one of the policies in place (we adopted this approach in this thesis) (J. I. Han et al. 2022). Alternatively, the entire actuation problem can be combined under a single global policy, which is only a plausible neuromechanics model if perfect neural integration of the bionic limb is assumed (De Vree et al. 2021). In a real life setting if a controller is trained from scratch, or finetuned with a user, this issue cannot be avoided. It is possible to partially resolve this issue by performing the device adaptation on a much slower scale; allowing enough time for the human to "converge to an optimum" before adjusting device parameters (J. Zhang et al. 2017; Diaz et al. 2022; Bryan et al. 2021). However, this leads to lengthy calibration sessions that need to be performed in laboratory environments. Addressing the simultaneous multiagent learning problem could reduce the necessary adaptation time. Game theoretical methods may be the solution to resolve two of the issues in multiagent learning: the credit assignment problem, i.e. how to distribute the rewards between policies from achieving collaborative goals (D. Han et al. 2021), and quantifying and distributing the expected effort required to perform the necessary actions (Li et al. 2019). Once these novel learning strategies are demonstrated in simulations, they should be trialled in real life in human-in-the-loop settings



Figure 7.3: Locomotion inside a virtual home from an outside perspective, and from the perspective of a camera attached to a human approaching a stairwell.

with P&O devices or emulation hardware (Diaz et al. 2022).

7.2.6 Other applications

As it is used for augmenting datasets of computer vision systems integrated into semi-autonomous upper limb device control policies (Zhong et al. 2022), synthetic vision data also have use cases in environment recognition systems lower limb devices. This can help collect the large amounts of diverse data necessary to train these models (Kurbis et al. 2022). As a proof of concept, we attached a virtual camera to our locomotion agent navigating a domestic environment populated by realistic assets available for the Unity engine (Figure 7.3). In simulation, pixel perfect labels are available for the synthetic images, identifying features like stairs or obstacles. It may be possible to leverage these synthetic datasets to train systems that can select the appropriate locomotion mode for assistive devices (e.g. recognising incoming stair ascent), or to detect when a fall is imminent. Motion synthesisers like ours are also suitable testbeds to virtually prototype and determine details of novel human sensing devices, such as indoor radar-based gait analysis for at-home activity monitoring and fall identification (Hadjipanayi et al. 2024).

In Chapter 5 we have examined different conditions of the prosthesis receiving more or less information from its user (via intent estimators), but not the other way around. Bidirectional HMIs in neuroprostheses aim to re-establish the afferent communication between the bionic limb and its user (Petrini et al. 2019). In our simulations, interoceptive and exteroceptive

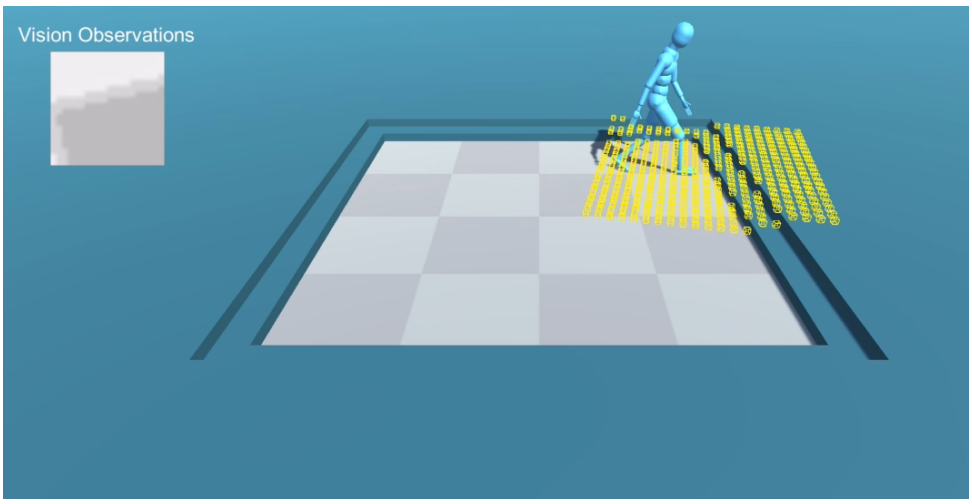


Figure 7.4: Spatial awareness through height field vision. By intercepting rays with the surrounding geometry, with origins arranged in a grid above the the humanoid, we construct a height map based on the distance each ray travels before hitting the ground. The resulting grey-scale image can be processed by a convolutional neural network or other encoding technique to create rudimentary "vision" for the agent. Yellow markers around agent show the intercept points for the rays, and the resulting image is shown in the top left corner.

sensing was used to direct the humanoid control policy. Through variations in the definition of the agent's observation space, it could be explored what details are most relevant to capture from the environment to provide to the user through sensory feedback. We performed a proof-of-concept experiment investigating whether RL agents learn compensatory gait patterns similar to prosthesis users in the presence of a sensory deficit, and whether these patterns disappear when providing simple exteroceptive sensing. Further details on this experiment are available in Appendix F. Additional research is also warranted in determining appropriate observation vectors for the locomotion agents. In our scenarios, only a narrow band of sensing was available to agents, appropriate for the simple locomotion environments they inhabited. If walking in larger spaces with complex obstacles, interactable objects or other humanoids is studied, egocentric vision (Merel et al. 2019), interaction graphs (Y. Zhang et al. 2023) or other parametrised representation of one's surroundings should be considered (another example is available on Figure 7.4). In a more general sense, theories about neuromechanical control, such as investigating the existence of human motor cost function or the structure of feedback mechanisms could be investigated through virtual environments like ours (Zhao et al. 2024).

7.2.7 Learning environments

Building learning environments with Unity empowered us to pursue ambitious research goals, while relying on the versatile tools for scene creation, behaviour scripting and animation control available for it. At the same time, communication of the RL framework with the Unity application forms a bottleneck in the learning process. We are currently investigating how to translate our environments to MJX, MuJoCo's hardware accelerated interface that can run massively parallel environments with orders of magnitude faster experience collection (Google DeepMind 2024b). Once sufficiently robust baseline policies can be trained via MJX for simple tasks, importing them to Unity for fine-tuning in complex environments could lead to faster iterations on our models of P&O use.

7.3 Next steps: Intent-driven powered prosthesis control with empirical uncertainty

In Chapter 5, the intent driven prosthesis received the target horizontal walking velocity as a surrogate signal of the virtual user's locomotion goals. This signal was also used to drive the motion synthesis. As a synthetic signal, it was consistent throughout the simulation and noise free. However, we have seen in Chapter 4 that this is not the case for a real intent estimator driven by biosignals.

Our current research goal is to evaluate the performance of our experimentally determined intent estimators, the output of which come with a level of uncertainty. We propose to model this uncertainty and superimpose it over the noiseless signal, allowing us to indirectly evaluate the performance of an integrated system in silico. This approach could be taken to model residuals in the trajectory space or in the parameter space (Method 2 vs. Method 3 in Chapter 4). We will start with the parametrised intent representation due to its lower dimensionality. We illustrate this proposal on Figure 7.5.

As an illustration of the properties of the residual of our high-level estimators, we show in Figure

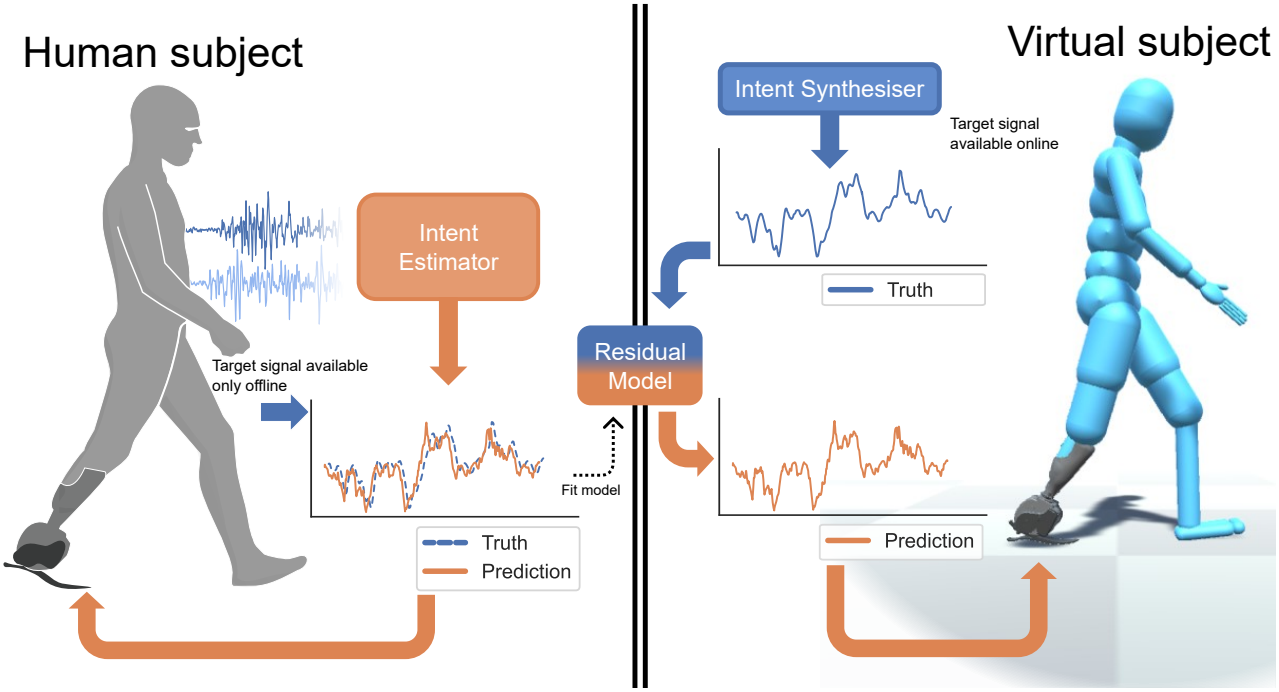


Figure 7.5: Proposal to evaluate the performance of an intent estimator trained from experimental data in simulation. The left side of the figure illustrates capturing biosignals and training the intent estimator to predict motion goals extracted offline (after the experiment). The intent estimator could then be used to provide input to the controller of an assistive device. The right side of the figure shows how a simulated locomotion is different; in virtual environments the ground truth motion goals are available even online (during the experiment), but there are no biosignals to use in the intent estimator. Instead, we can fit a model on the empirical performance of the intent estimator to predict the level of error given the true target. This residual model could then be used in the simulated environment to provide inputs to the virtual prosthesis at a realistic level of accuracy.

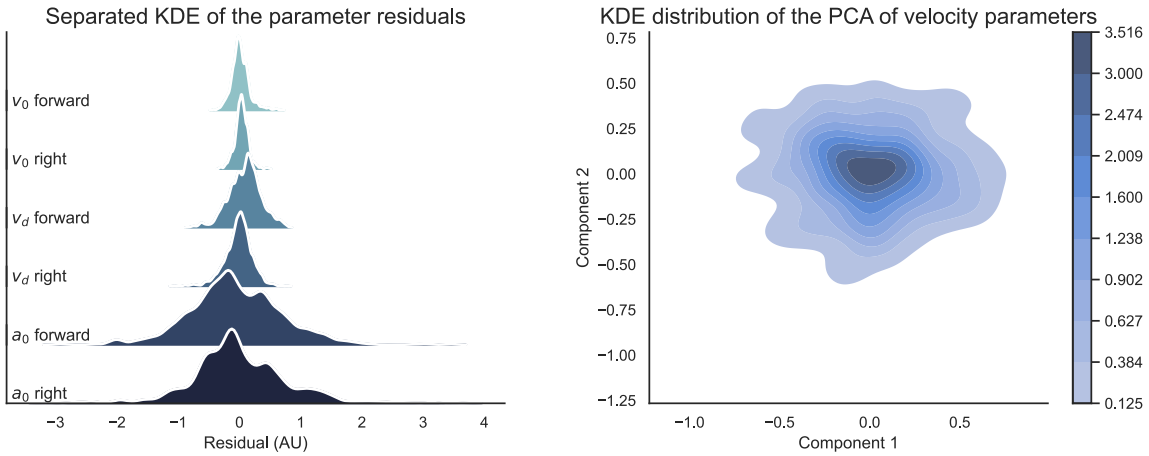


Figure 7.6: Subject 4's parameter space residual distributions for one of the EMG + Pose models trained with Method 3, approximated using kernel density estimation (KDE). In the left subfigure, we show the marginal distributions for each separate element. In the right subfigure we reduce the dimensionality to two via principal component analysis (PCA) and show their joint distributions.

7.6 the predicted distributions of the residuals in the parameters for one of the models trained with Method 3. We see that the uncertainty is not fully independent between the parameters, and that it may also be conditioned on the current true walking state (Figure 7.7).

In addition to the statistical dependence between the residual's dimensions (i.e., the residual is non-whitened), and the dependence between the residual vector and the true parameters, another factor to consider is the autocorrelation of the residual across time. This is partially due to the sliding window-based processing preformed by the TCN, which couples the variability across time. For example, a large point interference caused by a mechanical impact on the electrodes will negatively influence the prediction accuracy until that artifact exits the input window. This may suggest that a number of past residual samples equal to the window size should be used, however partial autocorrelation function analysis suggests that much shorter context sizes may be sufficient to capture the autoregressive properties of the residuals (Figure 7.8).

There are multiple approaches may be considered appropriate to generate synthetic residuals to satisfy more or less of these properties, including:

- Auto-regressive moving-average models. As the residual is strongly coupled across its

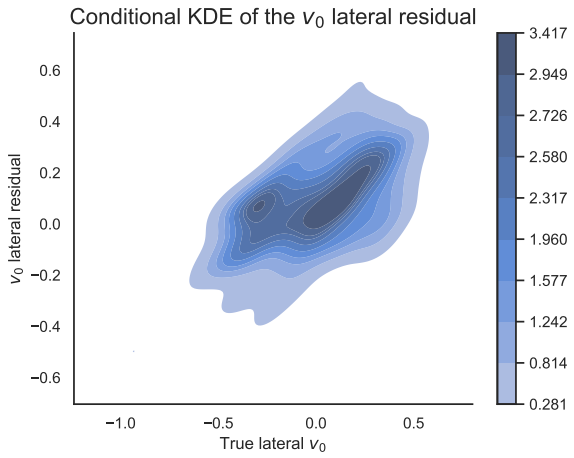


Figure 7.7: Subject 4's conditional KDE relating the residual of the predicted lateral v_0 to the true lateral v_0 for a model trained with Method 3 and pose-only inputs. The starting lateral velocity is consistently underestimated.

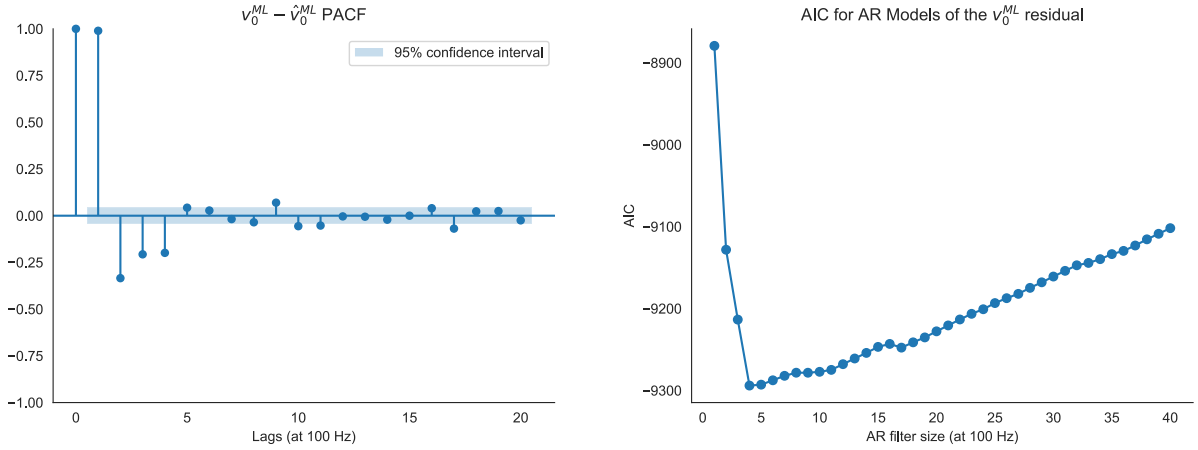


Figure 7.8: Yule-Walker partial autocorrelation function (PACF) of the mediolateral (ML) component of the starting velocity parameter's (v_0) residual, calculated for one of the models of Subject 1 from Chapter 4 on a validation "Angular Turns" trial. Method 3 used with the multimodal datasets. Additionally we provided the Akaike information criterion (AIC) measure for AR models of different size, suggesting a window size of ~ 40 ms to be appropriate to model the autoregressive properties of the residual. Repeating this analysis for the v_d paramter's residual in the mediolateral direction suggests a ~ 90 ms window .

multiple dimensions, vector autoregression may also be considered instead of dimension-wise separate models. Either approach relies on an assumption of stationarity, therefore they may require additional care before applying to diverse behavioural data like ours.

- Gaussian process regression modelling may be able to capture the residual's heteroskedasticity (e.g. Figure 7.7, another example from previous chapters is shown on Figure 3.12). Similarly, both temporal and feature space covariances of the residual can be determined with this approach. Multiple response variables (i.e., the outputs of the regressor) are commonly modelled independently from each other (B. Wang et al. 2015), although we observe in Figure 7.6 that this assumption may not be appropriate in our case, warranting additional care when employing this approach.

- Bayesian inference could be incorporated directly into the estimator networks, providing us directly with the estimated distribution from the output layer (Lampinen et al. 2001). By examining the learned posterior distributions of the parameters, an inverse mapping from the true signals could be established.

We propose to investigate these methods starting with a Gaussian process regressor that maps the ground truth signal to the residual of one of the intent estimators trained in Chapter 4. We could then use this model of the state-dependent covariance to inject noise with some properties resembling our intent estimators trained with experimental data. With this we can then evaluate if the level of accuracy achieved would be sufficient to assist prosthetic control policies, by relying on the distorted signal in the prosthesis use models of Chapter 5. We aim to characterise the sensitivity of the prosthesis controller to uncertainty in the intent estimation by scaling the covariance of the modelled noise. We will combine this investigation with using our models of the SoftFoot from Chapter 6 to trial different actuation and control schemes with the intent signal for a intent-driven, powered compliant prosthetic ankle-foot system.

The field of P&O is inherently tied to modelling. Intentionally or not, creating these devices is about attempting to capture what role our limbs play in our lives, and identifying what biomechanical deficits arise when they are affected or missing. By including the imperfections of the HMI in our models, we are continuing the work towards the goals of this PhD: understanding

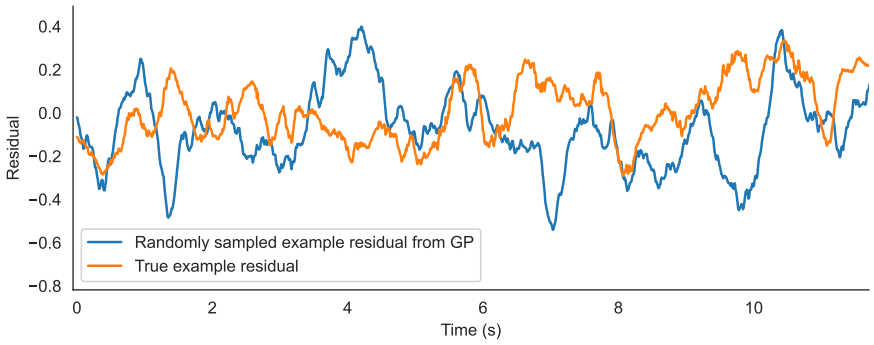


Figure 7.9: Time varying Gaussian process with length-scales and covariance scales fitted on example residuals through optimising the likelihood. A Matérn kernel with $\nu = 1$ was used. We sampled a random trajectory from the process, which should be different than the experimental signals, but resemble them in signal characteristics.

how wearable devices and the people who use them can overcome the shared challenge of restoring movement.

Bibliography

Abul-Haj, Cary and Hogan, Neville (1987). "An emulator system for developing improved elbow-prosthesis designs". In: *IEEE transactions on biomedical engineering* 9, pp. 724–737.

Adobe Systems Incorporated (2019). *Adobe's Mixamo*. `https://www.mixamo.com`. Online; accessed 18-Jun-2021.

Aftabi, Hamidreza, Nasiri, Rezvan, and Ahmadabadi, Majid Nili (2021). "Simulation-based biomechanical assessment of unpowered exoskeletons for running". In: *Scientific Reports* 11.1, p. 11846.

Akkaya, Ilge, Andrychowicz, Marcin, Chociej, Maciek, Litwin, Mateusz, McGrew, Bob, Petron, Arthur, Paino, Alex, Plappert, Matthias, Powell, Glenn, Ribas, Raphael, et al. (2019). "Solving rubik's cube with a robot hand". In: *arXiv preprint arXiv:1910.07113*.

Alimusaj, Merkur, Fradet, Laetitia, Braatz, Frank, Gerner, Hans J, and Wolf, Sebastian I (2009). "Kinematics and kinetics with an adaptive ankle foot system during stair ambulation of transtibial amputees". In: *Gait & posture* 30.3, pp. 356–363.

Ameri, Ali, Akhaee, Mohammad Ali, Scheme, Erik, and Englehart, Kevin (2020). "A Deep Transfer Learning Approach to Reducing the Effect of Electrode Shift in EMG Pattern Recognition-Based Control". In: *IEEE transactions on neural systems and rehabilitation engineering* 28.2, pp. 370–379.

Anam, Khairul and Al-Jumaily, Adel Ali (2012). "Active exoskeleton control systems: State of the art". In: *Procedia Engineering* 41, pp. 988–994.

Anders, Christoph, Wagner, Heiko, Puta, Christian, Grassme, Roland, Petrovitch, Alexander, and Scholle, Hans-Christoph (2007). "Trunk muscle activation patterns during walking at different speeds". In: *Journal of Electromyography and Kinesiology* 17.2, pp. 245–252.

Arazpour, Mokhtar, Bani, Monireh Ahmadi, and Hutchins, Stephen W (2013). "Reciprocal gait orthoses and powered gait orthoses for walking by spinal cord injury patients". In: *Prosthetics and orthotics international* 37.1, pp. 14–21.

Armiger, Robert S, Tenore, Francesco V, Bishop, William E, Beaty, James D, Bridges, Michael M, Burck, James M, Vogelstein, R Jacob, and Harshbarger, Stuart D (2011). "A real-time virtual integration environment for neuroprosthetics and rehabilitation". In: *Johns Hopkins APL Technical Digest* 30.3, pp. 198–206.

Asif, Muhammad, Tiwana, Mohsin Islam, Khan, Umar Shahbaz, Qureshi, Waqar Shahid, Iqbal, Javaid, Rashid, Nasir, and Naseer, Noman (2021). "Advancements, trends and future prospects of lower limb prosthesis". In: *IEEE Access* 9, pp. 85956–85977.

Aubin, Patrick M., Whittaker, Eric, and Ledoux, William R. (2012). "A Robotic Cadaveric Gait Simulator With Fuzzy Logic Vertical Ground Reaction Force Control". In: *IEEE Transactions on Robotics* 28.1, pp. 246–255. DOI: 10.1109/TRO.2011.2164958.

Azocar, Alejandro F, Mooney, Luke M, Duval, Jean-François, Simon, Ann M, Hargrove, Levi J, and Rouse, Elliott J (2020). "Design and clinical implementation of an open-source bionic leg". In: *Nature biomedical engineering* 4.10, pp. 941–953.

Bai, Shaojie, Kolter, J Zico, and Koltun, Vladlen (2018). "An empirical evaluation of generic convolutional and recurrent networks for sequence modeling". In: *arXiv preprint arXiv:1803.01271*.

Barrett, Samuel, Taylor, Matthew E, and Stone, Peter (2010). "Transfer learning for reinforcement learning on a physical robot". In: *Ninth International Conference on Autonomous Agents and Multiagent Systems-Adaptive Learning Agents Workshop (AAMAS-ALA)*. Vol. 1.

Baud, Romain, Manzoori, Ali Reza, Ijspeert, Auke, and Bouri, Mohamed (2021). "Review of control strategies for lower-limb exoskeletons to assist gait". In: *Journal of NeuroEngineering and Rehabilitation* 18.1, pp. 1–34.

Beil, Jonas, Perner, Gernot, and Asfour, Tamim (2015). "Design and control of the lower limb exoskeleton KIT-EXO-1". In: *2015 IEEE International Conference on Rehabilitation Robotics (ICORR)*. IEEE, pp. 119–124.

Bergamin, Kevin, Clavet, Simon, Holden, Daniel, and Forbes, James Richard (2019). "DReCon: data-driven responsive control of physics-based characters". In: *ACM Transactions on Graphics (TOG)* 38.6, pp. 1–11.

Berge, Jocelyn S, Klute, Glenn K, and Czerniecki, Joseph M (2004). "Mechanical properties of shock-absorbing pylons used in transtibial prostheses". In: *J. Biomech. Eng.* 126.1, pp. 120–122.

Betthauser, Joseph L, Krall, John T, Bannowsky, Shain G, Lévay, György, Kaliki, Rahul R, Fifer, Matthew S, and Thakor, Nitish V (2019). "Stable responsive EMG sequence prediction and adaptive reinforcement with temporal convolutional networks". In: *IEEE Transactions on Biomedical Engineering* 67.6, pp. 1707–1717.

Biddiss, Elaine and Chau, Tom (2007). "Upper-limb prosthetics: critical factors in device abandonment". In: *American journal of physical medicine & rehabilitation* 86.12, pp. 977–987.

Bogue, Robert (2018). "Exoskeletons–a review of industrial applications". In: *Industrial Robot: An International Journal* 45.5, pp. 585–590.

Bollo, David (2016). "Inertialization: High-Performance Animation Transitions in 'Gears of War'". In: *Proc. of GDC 2018*.

— (2017). "High performance animation in Gears of War 4". In: *ACM SIGGRAPH 2017 Talks*. Association for Computing Machinery, pp. 1–2.

Booth, Joe and Ivanov, Vladimir (2020). "Realistic Physics Based Character Controller". In: *arXiv preprint arXiv:2006.07508*.

Boser, Quinn A, Valevicius, Aïda M, Lavoie, Ewen B, Chapman, Craig S, Pilarski, Patrick M, Hebert, Jacqueline S, and Vette, Albert H (2018). "Cluster-based upper body marker models for three-dimensional kinematic analysis: Comparison with an anatomical model and reliability analysis". In: *Journal of biomechanics* 72, pp. 228–234.

Branch, Mary Ann, Coleman, Thomas F, and Li, Yuying (1999). "A subspace, interior, and conjugate gradient method for large-scale bound-constrained minimization problems". In: *SIAM Journal on Scientific Computing* 21.1, pp. 1–23.

Bruijn, Sjoerd M, Bregman, Daan JJ, Meijer, Onno G, Beek, Peter J, and Dieën, Jaap H van (2012). "Maximum Lyapunov exponents as predictors of global gait stability: a modelling approach". In: *Medical engineering & physics* 34.4, pp. 428–436.

Bryan, Gwendolyn M et al. (2021). "Optimized hip-knee-ankle exoskeleton assistance reduces the metabolic cost of walking with worn loads". In: *Journal of neuroengineering and rehabilitation* 18.1, pp. 1–13.

Bunderson, Nathan E. (2014). "Real-time control of an interactive impulsive virtual prosthesis." eng. In: *IEEE transactions on neural systems and rehabilitation engineering : a publication of the IEEE Engineering in Medicine and Biology Society* 22.2, pp. 363–370. DOI: 10.1109/TNSRE.2013.2274599.

Burden, Adrian (2010). "How should we normalize electromyograms obtained from healthy participants? What we have learned from over 25 years of research". In: *Journal of electromyography and kinesiology* 20.6, pp. 1023–1035.

Burdet, Etienne, Franklin, David W, and Milner, Theodore E (2013). *Human robotics: neuromechanics and motor control*. MIT press.

Bustamante, Samuel, Peters, Jan, Schölkopf, Bernhard, Grosse-Wentrup, Moritz, and Jayaram, Vinay (2021). "ArmSym: A Virtual Human–Robot Interaction Laboratory for Assistive Robotics". In: *IEEE Transactions on Human-Machine Systems* 51.6, pp. 568–577. DOI: 10.1109/THMS.2021.3106865.

Buttner, Michael (2015). "Motion Matching-The Road to Next-Gen Animation". In: *Nucl. ai Conference*.

Caggiano, Vittorio, Wang, Huawei, Durandau, Guillaume, Sartori, Massimo, and Kumar, Vikash (2022). "MyoSuite–A contact-rich simulation suite for musculoskeletal motor control". In: *arXiv preprint arXiv:2205.13600*.

Camargo, Jonathan, Flanagan, Will, Csomay-Shanklin, Noel, Kanwar, Bharat, and Young, Aaron (2021). "A machine learning strategy for locomotion classification and parameter estimation using fusion of wearable sensors". In: *IEEE Transactions on Biomedical Engineering* 68.5, pp. 1569–1578.

Campanini, I., Merlo, A., Degola, P., Merletti, R., Vezzosi, G., and Farina, D. (2006). "Effect of electrode location on EMG signal envelope in leg muscles during gait". In: *Journal of electromyography and kinesiology* 17.4, pp. 515–526.

Caputo, Joshua M and Collins, Steven H (2014). "A universal ankle–foot prosthesis emulator for human locomotion experiments". In: *Journal of biomechanical engineering* 136.3.

Caulcrick, Christopher, Huo, Weiguang, Hoult, Will, and Vaidyanathan, Ravi (2021). "Model Predictive Control for Human-Centred Lower Limb Robotic Assistance". In: *IEEE Transactions on Medical Robotics and Bionics* 3.4, pp. 980–991. DOI: https://doi.org/10.1109/TMRB.2021.3105141.

Ceccato, Jean-Charles, De Sèze, Mathieu, Azevedo, Christine, and Cazalets, Jean-René (2009). "Comparison of trunk activity during gait initiation and walking in humans". In: *PLoS One* 4.12, e8193.

Chadwick, Edward K., Blana, Dimitra, Bogert, Antonie J. van den, and Kirsch, Robert F. (2009). "A Real-Time, 3-D Musculoskeletal Model for Dynamic Simulation of Arm Movements". In: *IEEE Transactions on Biomedical Engineering* 56.4, pp. 941–948. DOI: 10.1109/TBME.2008.2005946.

Chang, Eunhee, Kim, Hyun Taek, and Yoo, Byounghyun (2020). "Virtual reality sickness: a review of causes and measurements". In: *International Journal of Human–Computer Interaction* 36.17, pp. 1658–1682.

Chao, E.Y., Laughman, R.K., Schneider, E., and Stauffer, R.N. (1983). "Normative data of knee joint motion and ground reaction forces in adult level walking". In: *Journal of biomechanics* 16.3, pp. 219–233.

Chappell, Digby, Son, Honn Wee, Clark, Angus B., Yang, Zeyu, Bello, Fernando, Kormushev, Petar, and Rojas, Nicolas (2022). "Virtual Reality Pre-Prosthetic Hand Training With Physics Simulation and Robotic Force Interaction". In: *IEEE Robotics and Automation Letters* 7.2, pp. 4550–4557. DOI: 10.1109/LRA.2022.3151569.

Chen, Baojun, Feng, Yanggang, and Wang, Qining (2016). "Combining vibrotactile feedback with volitional myoelectric control for robotic transtibial prostheses". In: *Frontiers in neurorobotics* 10, p. 8.

Chen, Baojun, Wang, Qining, and Wang, Long (2014). "Adaptive slope walking with a robotic transtibial prosthesis based on volitional EMG control". In: *IEEE/ASME Transactions on mechatronics* 20.5, pp. 2146–2157.

Chollet, François et al. (2015). *Keras*. https://keras.io.

Chu, J.-U., Moon, I., and Mun, M.-S. (2006). "A Real-Time EMG Pattern Recognition System Based on Linear-Nonlinear Feature Projection for a Multifunction Myoelectric Hand". In: *IEEE Transactions on Biomedical Engineering* 53.11, pp. 2232–2239. DOI: 10.1109/TBME.2006.883695.

Cimolato, Andrea, Driessen, Josephus JM, Mattos, Leonardo S, De Momi, Elena, Laffranchi, Matteo, and De Michieli, Lorenzo (2022). "EMG-driven control in lower limb prostheses: A topic-based systematic review". In: *Journal of NeuroEngineering and Rehabilitation* 19.1, p. 43.

Clavet, Simon (2016). "Motion matching and the road to next-gen animation". In: *Proc. of GDC*. Vol. 2016.

Clement, RGE, Bugler, Kate E, and Oliver, Christopher W (2011). "Bionic prosthetic hands: A review of present technology and future aspirations". In: *The surgeon* 9.6, pp. 336–340.

CMU, CMU (2003). "Graphics Lab Motion Capture Database". In: *URL http://mocap. cs. cmu. edu*.

Coapt (2024). *Complete Control System GEN2, Handbook (V 6.1)*. URL: https://coaptengineering.com/wp-content/uploads/Coapt-Gen2-Handbook.pdf.

Cohen, Joel E (2003). "Human population: the next half century". In: *Science* 302.5648, pp. 1172–1175.

Colombo, Giorgio, Facoetti, Giancarlo, Regazzoni, Daniele, and Rizzi, Caterina (2013). "A full virtual approach to design and test lower limb prosthesis: This paper reports a software platform for design and validation of lower limb prosthesis in a completely virtual environment, potentially replacing current manual process". In: *Virtual and Physical Prototyping* 8.2, pp. 97–111.

Community, Blender Online (2018). *Blender - a 3D modelling and rendering package*. Blender Foundation. Stichting Blender Foundation, Amsterdam. URL: http://www.blender.org.

Conti, Edoardo, Madhavan, Vashisht, Petroski Such, Felipe, Lehman, Joel, Stanley, Kenneth, and Clune, Jeff (2018). "Improving exploration in evolution strategies for deep reinforcement learning via a population of novelty-seeking agents". In: *Advances in neural information processing systems* 31.

Cybathlon ETH, Zürich (2020). *CYBATHLON 2020 Global Edition - Results*. https://cybathlon.ethz.ch/en/event/results. Online; accessed 22-Dec-2020.

De Groote, Friedl and Falisse, Antoine (2021). "Perspective on musculoskeletal modelling and predictive simulations of human movement to assess the neuromechanics of gait". In: *Proceedings of the Royal Society B* 288.1946, p. 20202432.

De Luca, Carlo J (1997). "The use of surface electromyography in biomechanics". In: *Journal of applied biomechanics* 13.2, pp. 135–163.

De Vree, Leanne and Carloni, Raffaella (2021). "Deep reinforcement learning for physics-based musculoskeletal simulations of healthy subjects and transfemoral prostheses' users during normal walking". In: *IEEE Transactions on Neural Systems and Rehabilitation Engineering* 29, pp. 607–618.

Della Santina, Cosimo, Piazza, Cristina, Grioli, Giorgio, Catalano, Manuel G, and Bicchi, Antonio (2018). "Toward dexterous manipulation with augmented adaptive synergies: The pisa/iit softhand 2". In: *IEEE Transactions on Robotics* 34.5, pp. 1141–1156.

Delp, Scott L, Anderson, Frank C, Arnold, Allison S, Loan, Peter, Habib, Ayman, John, Chand T, Guendelman, Eran, and Thelen, Darryl G (2007). "OpenSim: open-source software to create and analyze dynamic simulations of movement". In: *IEEE transactions on biomedical engineering* 54.11, pp. 1940–1950.

Delsys Incorporated, USA (2021). *Delsys Trigno® Sensors*. https://delsys.com/trigno-avanti/. Online; accessed 16-07-2021.

Dembia, Christopher L., Bianco, Nicholas A., Falisse, Antoine, Hicks, Jennifer L., and Delp, Scott L. (2020). "OpenSim Moco: Musculoskeletal optimal control". In: *PLOS Computational Biology* 16.12. DOI: 10.1371/journal.pcbi.1008493.

Deutsch, Judith, Lewis, Jeffrey A, and Burdea, Grigore (2007a). "Technical and patient performance using a virtual reality-integrated telerehabilitation system: preliminary finding". In: *IEEE Transactions on Neural Systems and Rehabilitation Engineering* 15.1, pp. 30–35.

Deutsch, Judith and Mirelman, A. (2007b). "Virtual Reality-Based Approaches to Enable Walking for People Poststroke." In: *Top. stroke rehabil.* 14.6, pp. 45–53. DOI: 10.1310/tsr1406-45.

Diaz, Maria Alejandra, Voss, Matthias, Dillen, Arnau, Tassignon, Bruno, Flynn, Louis, Geeroms, Joost, Meeusen, Romain, Verstraten, Tom, Babič, Jan, Beckerle, Philipp, et al. (2022). "Human-in-the-loop optimization of wearable robotic devices to improve human–robot interaction: A systematic review". In: *IEEE Transactions on Cybernetics* 53.12, pp. 7483–7496.

Dietterich, Thomas G (1998). "Approximate statistical tests for comparing supervised classification learning algorithms". In: *Neural computation* 10.7, pp. 1895–1923.

Dietz, Volker (2002). "Proprioception and locomotor disorders". In: *Nature Reviews Neuroscience* 3.10, pp. 781–790.

Dingwell, JB, Davis, BL, and Frazder, DM (1996). "Use of an instrumented treadmill for real-time gait symmetry evaluation and feedback in normal and trans-tibial amputee subjects". In: *Prosthetics and orthotics international* 20.2, pp. 101–110.

Ditunno, PL, Patrick, M, Stineman, M, and Ditunno, JF (2008). "Who wants to walk? Preferences for recovery after SCI: a longitudinal and cross-sectional study". In: *Spinal cord* 46.7, pp. 500–506.

Duysens, Jacques and Van de Crommert, Henry WAA (1998). "Neural control of locomotion; Part 1: The central pattern generator from cats to humans". In: *Gait & posture* 7.2, pp. 131–141.

Eames, MHA, Cosgrove, A, and Baker, R (1999). "Comparing methods of estimating the total body centre of mass in three-dimensions in normal and pathological gaits". In: *Human movement science* 18.5, pp. 637–646.

Erez, Tom, Tassa, Yuval, and Todorov, Emanuel (2015). "Simulation tools for model-based robotics: Comparison of bullet, havok, mujoco, ode and physx". In: *2015 IEEE international conference on robotics and automation (ICRA)*. IEEE, pp. 4397–4404.

Eshraghi, Arezoo, Osman, Noor Azuan Abu, Gholizadeh, Hossein, Karimi, Mohammad, and Ali, Sadeeq (2012). "Pistoning assessment in lower limb prosthetic sockets". In: *Prosthetics and Orthotics International* 36.1, pp. 15–24.

Falisse, Antoine, Serrancolí, Gil, Dembia, Christopher L, Gillis, Joris, Jonkers, Ilse, and De Groote, Friedl (2019). "Rapid predictive simulations with complex musculoskeletal models suggest that diverse healthy and pathological human gaits can emerge from similar control strategies". In: *Journal of The Royal Society Interface* 16.157, p. 20190402.

Fang, Bin, Zhou, Quan, Sun, Fuchun, Shan, Jianhua, Wang, Ming, Xiang, Cheng, and Zhang, Qin (2020). "Gait Neural Network for Human-Exoskeleton Interaction". In: *Frontiers in Neurorobotics* 14.

Farina, Dario, Jensen, Winnie, and Akay, Metin (2013a). "Introduction to neural engineering for motor rehabilitation". In: John Wiley & Sons. Chap. 14.

— (2013b). "Introduction to neural engineering for motor rehabilitation". In: John Wiley & Sons. Chap. 19.

Farris, Dominic James, Hicks, Jennifer L., Delp, Scott L, and Sawicki, Gregory S (2014). "Musculoskeletal modelling deconstructs the paradoxical effects of elastic ankle exoskeletons on plantar-flexor mechanics and energetics during hopping". In: *Journal of Experimental Biology* 217.22, pp. 4018–4028.

Featherstone, Roy (2010). "A beginner's guide to 6-d vectors (part 1)". In: *IEEE robotics & automation magazine* 17.3, pp. 83–94.

Feng, Gilbert, Zhang, Hongbo, Li, Zhongyu, Peng, Xue Bin, Basireddy, Bhuvan, Yue, Linzhu, Song, Zhitao, Yang, Lizhi, Liu, Yunhui, Sreenath, Koushil, et al. (2022). "GenLoco: Generalized Locomotion Controllers for Quadrupedal Robots". In: *arXiv preprint arXiv:2209.05309*.

Fey, Nicholas P, Klute, Glenn K, and Neptune, Richard R (2011). "The influence of energy storage and return foot stiffness on walking mechanics and muscle activity in below-knee amputees". In: *Clinical Biomechanics* 26.10, pp. 1025–1032.

Flash, Tamar and Hogan, Neville (1985). "The coordination of arm movements: an experimentally confirmed mathematical model". In: *Journal of neuroscience* 5.7, pp. 1688–1703.

Fliegel, O and Feuer, SG (1966). "Historical development of lower-extremity prostheses". In: *Archives of Physical Medicine and Rehabilitation* 47.5, pp. 275–285.

Forbes, Patrick A, Chen, Anthony, and Blouin, Jean-Sébastien (2018). "Sensorimotor control of standing balance". In: *Handbook of clinical neurology* 159, pp. 61–83.

Fougner, Anders, Stavdahl, Øyvind, Kyberd, Peter J, Losier, Yves G, and Parker, Philip A (2012). "Control of upper limb prostheses: Terminology and proportional myoelectric control—A review". In: *IEEE Transactions on neural systems and rehabilitation engineering* 20.5, pp. 663–677.

Fussell, Levi, Bergamin, Kevin, and Holden, Daniel (2021). "Supertrack: Motion tracking for physically simulated characters using supervised learning". In: *ACM Transactions on Graphics (TOG)* 40.6, pp. 1–13.

Galan, Ferran and Baker, Stuart N. (2015). "Deafferented controllers: a fundamental failure mechanism in cortical neuroprosthetic systems". In: *Frontiers in Behavioral Neuroscience* 9. DOI: 10.3389/fnbeh.2015.00186.

Gehlhar, Rachel, Tucker, Maegan, Young, Aaron J, and Ames, Aaron D (2023). "A review of current state-of-the-art control methods for lower-limb powered prostheses". In: *Annual Reviews in Control*.

Geijtenbeek, Thomas (2019). "Scone: Open source software for predictive simulation of biological motion". In: *Journal of Open Source Software* 4.38, p. 1421.

— (2021). *The Hyfydy Simulation Software*. https://hyfydy.com. URL: https://hyfydy.com.

Geijtenbeek, Thomas and Pronost, Nicolas (2012). "Interactive character animation using simulated physics: A state-of-the-art review". In: *Computer graphics forum*. Vol. 31. Wiley Online Library, pp. 2492–2515.

Ghannadi, Borna, Sharif Razavian, Reza, and McPhee, John (2018). "Configuration-dependent optimal impedance control of an upper extremity stroke rehabilitation manipulandum". In: *Frontiers in Robotics and AI* 5, p. 124.

Gillies, Marco (2009). "Learning finite-state machine controllers from motion capture data". In: *IEEE transactions on computational intelligence and AI in games* 1.1, pp. 63–72.

Gloumakov, Yuri, Bimbo, Joao, and Dollar, Aaron M. (2022). "Trajectory Control–An Effective Strategy for Controlling Multi-DOF Upper Limb Prosthetic Devices". In: *IEEE Transactions on Neural Systems and Rehabilitation Engineering* 30, pp. 420–430. DOI: 10.1109/TNSRE.2022.3151055.

Goniewicz, Krzysztof, Burkle, Frederick M, Dzhus, Marta, and Khorram-Manesh, Amir (2023). "Ukraine's healthcare crisis: sustainable strategies for navigating conflict and rebuilding for a resilient future". In: *Sustainability* 15.15, p. 11602.

Google DeepMind (2024a). *Computation - MuJoCo Documentation*. https://mujoco.readthedocs.io/en/latest/computation/index.html#constraint-model. (Visited on 10/14/2024).

— (2024b). *MuJoCo XLA (MJX) - MuJoCo Documentation*. URL: https://mujoco.readthedocs.io/en/stable/mjx.html#mujoco-xla-mjx (visited on 10/07/2024).

Gordon, Keith E, Ferris, Daniel P, and Kuo, Arthur D (2009). "Metabolic and mechanical energy costs of reducing vertical center of mass movement during gait". In: *Archives of physical medicine and rehabilitation* 90.1, pp. 136–144.

Graimann, Bernhard and Dietl, Hans (2013). "Introduction to Upper Limb Prosthetics". In: *Introduction to Neural Engineering for Motor Rehabilitation*. John Wiley & Sons, Ltd. Chap. 14, pp. 267–290. DOI: 10.1002/9781118628522.ch14.

Gregg, Robert D, Dhaher, Yasin Y, Degani, Amir, and Lynch, Kevin M (2012). "On the mechanics of functional asymmetry in bipedal walking". In: *IEEE Transactions on Biomedical Engineering* 59.5, pp. 1310–1318.

Gregg, Robert D., Lenzi, Tommaso, Hargrove, Levi J., and Sensinger, Jonathon W. (2014). "Virtual Constraint Control of a Powered Prosthetic Leg: From Simulation to Experiments With Transfemoral Amputees". In: *IEEE Transactions on Robotics* 30.6, pp. 1455–1471. DOI: 10.1109/TRO.2014.2361937.

Grood, Edward S and Suntay, Wilfredo J (1983). "A joint coordinate system for the clinical description of three-dimensional motions: application to the knee". In: *Journal of biomechanical engineering* 105.2, pp. 136–144.

Gu, Chenyu, Lin, Weicong, He, Xinyi, Zhang, Lei, and Zhang, Mingming (2023). "IMU-based motion capture system for rehabilitation applications: A systematic review". In: *Biomimetic Intelligence and Robotics* 3.2, p. 100097.

Guez, Annika, Hodossy, Balint K, Farina, Dario, and Vaidyanathan, Ravi (2023). "Transferring Gait Predictors Across EMG Acquisition Systems with Domain Adaptation". In: *2023 International Conference on Rehabilitation Robotics (ICORR)*. IEEE, pp. 1–6.

Gusman, Jacob, Mastinu, Enzo, and Ortiz-Catalan, Max (2017). "Evaluation of Computer-Based Target Achievement Tests for Myoelectric Control". In: *IEEE Journal of Translational Engineering in Health and Medicine* 5, pp. 1–10. DOI: 10.1109/JTEHM.2017.2776925.

Ha, David, Dai, Andrew, and Le, Quoc V (2016). "Hypernetworks". In: *arXiv preprint arXiv:1609.09106*.

Hadjipanayi, Charalambos, Yin, Maowen, Bannon, Alan, Rapeaux, Adrien, Banger, Matthew, Haar, Shlomi, Lande, Tor Sverre, McGregor, Alison H, and Constandinou, Timothy G (2024). "Remote Gait Analysis Using Ultra-Wideband Radar Technology Based on Joint Range-Doppler-Time Representation". In: *IEEE Transactions on Biomedical Engineering*.

Han, Dongge, Lu, Chris Xiaoxuan, Michalak, Tomasz, and Wooldridge, Michael (2021). "Multiagent model-based credit assignment for continuous control". In: *arXiv preprint arXiv:2112.13937*.

Han, Jong In, Lee, Jeong-Hoon, Choi, Ho Seon, Kim, Jung-Hoon, and Choi, Jongeun (2022). "Policy Design for an Ankle-Foot Orthosis Using Simulated Physical Human–Robot Interaction via Deep Reinforcement Learning". In: *IEEE Transactions on Neural Systems and Rehabilitation Engineering* 30, pp. 2186–2197.

Handford, Matthew L. and Srinivasan, Manoj (2018). "Energy-Optimal Human Walking With Feedback-Controlled Robotic Prostheses: A Computational Study". In: *IEEE Transactions on Neural Systems and Rehabilitation Engineering* 26.9, pp. 1773–1782. DOI: 10.1109/TNSRE.2018.2858204.

Hara, Masayuki, Shokur, Solaiman, Yamamoto, Akio, Higuchi, Toshiro, Gassert, Roger, and Bleuler, Hannes (2010). "Virtual environment to evaluate multimodal feedback strategies for augmented navigation of the visually impaired." eng. In: *Annual International Conference of the IEEE Engineering in Medicine and Biology Society. IEEE Engineering in Medicine and Biology Society. Annual International Conference* 2010, pp. 975–978. DOI: 10.1109/IEMBS.2010.5627611.

Hargrove, Levi J., Miller, Laura, Turner, Kristi, and Kuiken, Todd (2018). "Control within a virtual environment is correlated to functional outcomes when using a physical prosthesis." eng. In: *Journal of neuroengineering and rehabilitation* 15.Suppl 1, p. 60. DOI: 10.1186/s12984-018-0402-y.

Hargrove, Levi J., Simon, Ann M, Lipschutz, Robert D, Finucane, Suzanne B, and Kuiken, Todd A (2011). "Real-time myoelectric control of knee and ankle motions for transfemoral amputees". In: *JAMA* 305.15, pp. 1542–1544.

Hargrove, Levi J., Simon, Ann M., Lipschutz, Robert, Finucane, Suzanne B., and Kuiken, Todd A. (2013). "Non-weight-bearing neural control of a powered transfemoral prosthesis". In: *Journal of NeuroEngineering and Rehabilitation* 10.1, p. 62. DOI: 10.1186/1743-0003-10-62.

Harkins, Colette S, McGarry, Anthony, and Buis, Arjan (2013). "Provision of prosthetic and orthotic services in low-income countries: A review of the literature". In: *Prosthetics and Orthotics International* 37.5, pp. 353–361.

Harvey, Félix G, Yurick, Mike, Nowrouzezahrai, Derek, and Pal, Christopher (2020). "Robust motion in-betweening". In: *ACM Transactions on Graphics (TOG)* 39.4, pp. 60–1.

Hase, Kazunori and Stein, RB (1999). "Turning strategies during human walking". In: *Journal of neurophysiology* 81.6, pp. 2914–2922.

Hauschild, Markus, Davoodi, Rahman, and Loeb, Gerald E. (2007). "A Virtual Reality Environment for Designing and Fitting Neural Prosthetic Limbs". In: *IEEE Transactions on Neural Systems and Rehabilitation Engineering* 15.1, pp. 9–15. DOI: 10.1109/TNSRE.2007.891369.

He, Jiayuan, Sheng, Xinjun, Zhu, Xiangyang, and Jiang, Ning (2019). "A novel framework based on position verification for robust myoelectric control against sensor shift". In: *IEEE Sensors Journal* 19.21, pp. 9859–9868.

— (2020). "Position Identification for Robust Myoelectric Control Against Electrode Shift". In: *IEEE transactions on neural systems and rehabilitation engineering* 28.12, pp. 3121–3128.

He, Kaibo, Zuo, Chenhui, Shao, Jing, and Sui, Yanan (2023). *Self Model for Embodied Intelligence: Modeling Full-Body Human Musculoskeletal System and Locomotion Control with Hierarchical Low-Dimensional Representation.* arXiv:2312.05473 [cs]. DOI: 10.48550/arXiv.2312.05473.

Hermens, Hermie J, Freriks, Bart, Merletti, Roberto, Stegeman, Dick, Blok, Joleen, Rau, Günter, Disselhorst-Klug, Cathy, and Hägg, Göran (1999). "European recommendations for surface electromyography". In: *Roessingh research and development* 8.2, pp. 13–54.

Hicks, Jennifer L., Uchida, Thomas K., Seth, Ajay, Rajagopal, Apoorva, and Delp, Scott L. (2015). "Is My Model Good Enough? Best Practices for Verification and Validation of Musculoskeletal Models and Simulations of Movement". In: *Journal of Biomechanical Engineering* 137.020905. DOI: 10.1115/1.4029304.

Hidler, Joseph M. and Wall, Anji E. (2005). "Alterations in muscle activation patterns during robotic-assisted walking". In: *Clinical Biomechanics* 20.2, pp. 184–193.

Higgins, Irina, Matthey, Loic, Glorot, Xavier, Pal, Arka, Uria, Benigno, Blundell, Charles, Mohamed, Shakir, and Lerchner, Alexander (2016). "Early visual concept learning with unsupervised deep learning". In: *arXiv preprint arXiv:1606.05579.*

Hill, Deborah, Holloway, Catherine Sarah, Ramirez, Dafne Zuleima Morgado, Smitham, Peter, and Pappas, Yannis (2017). "What are user perspectives of exoskeleton technology? A literature review". In: *International journal of technology assessment in health care* 33.2, pp. 160–167.

Hodossy, Balint K (2024). *Mj-Unity Tutorials: Introductory set of tutorials for using the Unity plugin of MuJoCo with the ML-Agents framework.* URL: https://github.com/Balint-H/mj-unity-tutorial (visited on 09/10/2024).

Hodossy, Balint K and Farina, Dario (2023a). "Shared Autonomy Locomotion Synthesis With a Virtual Powered Prosthetic Ankle". In: *IEEE Transactions on Neural Systems and Rehabilitation Engineering* 31, pp. 4738–4748.

Hodossy, Balint K and Llobera, Joan (2023b). *Modular Agents: Extensions to the ML-Agents toolkit, focusing on humanoid control in Unity*. original-date: 2023-08-16T11:02:16Z. URL: https://github.com/Balint-H/modular-agents (visited on 10/11/2023).

Hodossy, Balint K, Llobera, Joan, and Mancuso, Matteo (2024). *Balint-H/drecon-unity: Implementation of the DReCon motion synthesis system with the MuJoCo, ML-Agents and Modular-Agents packages in Unity*. URL: https://github.com/Balint-H/drecon-unity (visited on 10/29/2024).

Höfer, Sebastian, Bekris, Kostas, Handa, Ankur, Gamboa, Juan Camilo, Mozifian, Melissa, Golemo, Florian, Atkeson, Chris, Fox, Dieter, Goldberg, Ken, Leonard, John, et al. (2021). "Sim2Real in robotics and automation: Applications and challenges". In: *IEEE transactions on automation science and engineering* 18.2, pp. 398–400.

Hogan, Neville (1984). "Impedance control: An approach to manipulation". In: *1984 American control conference*. IEEE, pp. 304–313.

Holden, Daniel, Kanoun, Oussama, Perepichka, Maksym, and Popa, Tiberiu (2020). "Learned motion matching". In: *ACM Transactions on Graphics (TOG)* 39.4, pp. 53–1.

Holden, Daniel, Komura, Taku, and Saito, Jun (2017). "Phase-functioned neural networks for character control". In: *ACM transactions on graphics* 36.4, pp. 1–13.

Hollerbach, John M (1982). "Computers, brains and the control of movement". In: *Trends in Neurosciences* 5, pp. 189–192.

Huang, He, Zhang, Fan, Hargrove, Levi J., Dou, Zhi, Rogers, Daniel R, and Englehart, Kevin B (2011). "Continuous locomotion-mode identification for prosthetic legs based on neuromuscular–mechanical fusion". In: *IEEE Transactions on Biomedical Engineering* 58.10, pp. 2867–2875.

Huang, He, Zhou, Ping, Li, Guanglin, and Kuiken, Todd A. (2008). "An Analysis of EMG Electrode Configuration for Targeted Muscle Reinnervation Based Neural Machine Interface". In: *IEEE transactions on neural systems and rehabilitation engineering* 16.1, pp. 37–45.

Inagaki, Takashi, Akiyama, Yasuhiro, Okamoto, Shogo, Mayumi, Takuya, and Yamada, Yoji (2023). "Relationship between gait stability indices and gait parameters comprising joint angles based on walking data from 288 people". In: *Nagoya Journal of Medical Science* 85.2, p. 211.

Ivanenko, Yuri P., Cappellini, Germana, Molinari, Marco, and Lacquaniti, Francesco (2013). "Motor Control Modules of Human Movement in Health and Disease". In: *Introduction to Neural Engineering for Motor Rehabilitation*. John Wiley & Sons, Ltd. Chap. 3, pp. 39–60. DOI: https://doi.org/10.1002/9781118628522.ch3.

Jang, Chul Ho, Yang, Hee Seung, Yang, Hea Eun, Lee, Seon Yeong, Kwon, Ji Won, Yun, Bong Duck, Choi, Jae Yung, Kim, Seon Nyeo, and Jeong, Hae Won (2011). "A survey on activities of daily living and occupations of upper extremity amputees". In: *Annals of rehabilitation medicine* 35.6, pp. 907–921.

Jegou, Herve, Douze, Matthijs, and Schmid, Cordelia (2010). "Product quantization for nearest neighbor search". In: *IEEE transactions on pattern analysis and machine intelligence* 33.1, pp. 117–128.

Jiang, Ning, Englehart, Kevin B, and Parker, Philip A (2008). "Extracting simultaneous and proportional neural control information for multiple-DOF prostheses from the surface electromyographic signal". In: *IEEE transactions on Biomedical Engineering* 56.4, pp. 1070–1080.

Jimenez-Fabian, René and Verlinden, Olivier (2012). "Review of control algorithms for robotic ankle systems in lower-limb orthoses, prostheses, and exoskeletons". In: *Medical engineering & physics* 34.4, pp. 397–408.

Johansen, Rune Skovbo (2009). "Automated semi-procedural animation for character locomotion". In: *Aarhus Universitet, Institut for Informations Medievidenskab*.

John Hopkins Applied Physics Laboratory (2017). *Revolutionizing Prosthetics, Virtual Integration Environment*. `https://www.jhuapl.edu/Prosthetics/ResearchVIE`. Online; accessed 12-Jan-2021.

Juliani, Arthur, Berges, Vincent-Pierre, Teng, Ervin, Cohen, Andrew, Harper, Jonathan, Elion, Chris, Goy, Chris, Gao, Yuan, Henry, Hunter, Mattar, Marwan, et al. (2018). "Unity: A general platform for intelligent agents". In: *arXiv preprint arXiv:1809.02627*.

Kang, Inseung, Molinaro, Dean D, Duggal, Srijan, Chen, Yanrong, Kunapuli, Pratik, and Young, Aaron J (2021). "Real-Time Gait Phase Estimation for Robotic Hip Exoskeleton Control during Multimodal Locomotion". In: *IEEE Robotics and Automation Letters* 6.2, pp. 3491–3497.

Kannape, Oliver A and Herr, Hugh M (2014). "Volitional control of ankle plantar flexion in a powered transtibial prosthesis during stair-ambulation". In: *2014 36th Annual International Conference of the IEEE Engineering in Medicine and Biology Society*. IEEE, pp. 1662–1665.

Kapelner, Tamas, Negro, Francesco, Aszmann, Oskar C, and Farina, Dario (2017). "Decoding motor unit activity from forearm muscles: perspectives for myoelectric control". In: *IEEE Transactions on Neural Systems and Rehabilitation Engineering* 26.1, pp. 244–251.

Karrenbach, Maxim, Preechayasomboon, Pornthep, Sauer, Peter, Boe, David, and Rombokas, Eric (2022). "Deep learning and session-specific rapid recalibration for dynamic hand gesture recognition from EMG". In: *Frontiers in Bioengineering and Biotechnology* 10. DOI: `10.3389/fbioe.2022.1034672`.

Karulkar, Roopak M and Wensing, Patrick M (2021). "Using footsteps to estimate changes in the desired gait speed of an exoskeleton user". In: *IEEE Robotics and Automation Letters* 6.4, pp. 6781–6788.

Kaya, Mahmut and Bilge, Hasan Şakir (2019). "Deep metric learning: A survey". In: *Symmetry* 11.9, p. 1066.

Kharb, Ashutosh, Saini, Vipin, Jain, YK, and Dhiman, Surender (2011). "A review of gait cycle and its parameters". In: *IJCEM International Journal of Computational Engineering & Management* 13, pp. 78–83.

Kidziński, Lukasz, Ong, Carmichael, Mohanty, Sharada Prasanna, Hicks, Jennifer L., Carroll, Sean, Zhou, Bo, Zeng, Hongsheng, Wang, Fan, Lian, Rongzhong, Tian, Hao, et al. (2020). "Artificial intelligence for prosthetics: Challenge solutions". In: *The NeurIPS'18 Competition*. Springer, pp. 69–128.

Kilgore, Kevin L, Scherer, Marcia, Bobblitt, Rod, Dettloff, Jocelyn, Dombrowski, David M, Godbold, Nicholas, Jatich, James W, Morris, Robert, Penko, Jennifer S, Schremp, Eric S, et al. (2001). "Neuroprosthesis consumers' forum: consumer priorities for research directions". In: *Journal of rehabilitation research and development* 38.6, pp. 655–660.

Kim, Keun-Tae, Guan, Cuntai, and Lee, Seong-Whan (2019). "A subject-transfer framework based on single-trial EMG analysis using convolutional neural networks". In: *IEEE Transactions on Neural Systems and Rehabilitation Engineering* 28.1, pp. 94–103.

Kluger, David T., Joyner, Janell S., Wendelken, Suzanne M., Davis, Tyler S., George, Jacob A., Page, David M., Hutchinson, Douglas T., Benz, Heather L., and Clark, Gregory A. (2019). "Virtual Reality Provides an Effective Platform for Functional Evaluations of Closed-Loop Neuromyoelectric Control". In: *IEEE Transactions on Neural Systems and Rehabilitation Engineering* 27.5, pp. 876–886. DOI: 10.1109/TNSRE.2019.2908817.

Knaepen, Kristel, Beyl, Pieter, Duerinck, Saartje, Hagman, Friso, Lefeber, Dirk, and Meeusen, Romain (2014). "Human-Robot Interaction: Kinematics and Muscle Activity Inside a Powered Compliant Knee Exoskeleton". In: *IEEE transactions on neural systems and rehabilitation engineering* 22.6, pp. 1128–1137.

Kodesh, Einat, Kafri, Michal, Dar, Gali, and Dickstein, Ruth (2012). "Walking speed, unilateral leg loading, and step symmetry in young adults". In: *Gait & posture* 35.1, pp. 66–69.

Koenig, Nathan and Howard, Andrew (2004). "Design and use paradigms for gazebo, an open-source multi-robot simulator". In: *2004 IEEE/RSJ international conference on intelligent robots and systems (IROS)(IEEE Cat. No. 04CH37566)*. Vol. 3. Ieee, pp. 2149–2154.

Kovar, Lucas, Gleicher, Michael, and Pighin, Frédéric (2008). "Motion graphs". In: *ACM SIGGRAPH 2008 classes*, pp. 1–10.

Krausz, Nili E and Hargrove, Levi J. (2019). "A survey of teleceptive sensing for wearable assistive robotic devices". In: *Sensors* 19.23, p. 5238.

Kuiken, Todd A, Li, Guanglin, Lock, Blair A, Lipschutz, Robert D, Miller, Laura A, Stubblefield, Kathy A, and Englehart, Kevin B (2009). "Targeted muscle reinnervation for real-time myoelectric control of multifunction artificial arms". In: *Jama* 301.6, pp. 619–628.

Kuiken, Todd A, Lowery, MM, and Stoykov, NS (2003). "The effect of subcutaneous fat on myoelectric signal amplitude and cross-talk". In: *Prosthetics and orthotics international* 27.1, pp. 48–54.

Kumar, Vikash (2015). *Modular Prosthetic Limb + SHAP test suites | Legacy MuJoCo Forum*. URL: https://roboti.us/forum/index.php?resources/modular-prosthetic-limb-shap-test-suites.19/ (visited on 09/26/2024).

Kurbis, Andrew Garrett, Laschowski, Brokoslaw, and Mihailidis, Alex (2022). "Stair recognition for robotic exoskeleton control using computer vision and deep learning". In: *2022 International Conference on Rehabilitation Robotics (ICORR)*. IEEE, pp. 1–6.

Kyeong, Seulki and Kim, Jung (2020). "Enhanced torque estimation method from multi-channel surface electromyography compensating electrode location variation". In: *2020 42nd Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC)*. IEEE, pp. 3965–3968.

Lampinen, Jouko and Vehtari, Aki (2001). "Bayesian approach for neural networks—review and case studies". In: *Neural networks* 14.3, pp. 257–274.

Lapatki, Bernd G., Oostenveld, Robert, Dijk, Johannes P. Van, Jonas, Irmtrud E., Zwarts, Machiel J., and Stegeman, Dick F. (2010). "Optimal placement of bipolar surface EMG electrodes in the face based on single motor unit analysis". In: *Psychophysiology* 47.2, pp. 299–314.

Lea, Colin, Flynn, Michael D, Vidal, Rene, Reiter, Austin, and Hager, Gregory D (2017). "Temporal convolutional networks for action segmentation and detection". In: *proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 156–165.

Lea, Colin, Vidal, Rene, Reiter, Austin, and Hager, Gregory D (2016). "Temporal convolutional networks: A unified approach to action segmentation". In: *European Conference on Computer Vision*. Springer, pp. 47–54.

Leboeuf, Fabien, Baker, R, Barré, Arnaud, Reay, Julie, Jones, Richard, and Sangeux, Morgan (2019). "The conventional gait model, an open-source implementation that reproduces the past but prepares for the future". In: *Gait & posture* 69, pp. 235–241.

Lee, Seunghwan, Park, Moonseok, Lee, Kyoungmin, and Lee, Jehee (2019). "Scalable muscle-actuated human simulation and control". In: *ACM Transactions On Graphics (TOG)* 38.4, pp. 1–13.

Lee, Seyoung, Lee, Sunmin, Lee, Yongwoo, and Lee, Jehee (2021). "Learning a family of motor skills from a single motion clip". In: *ACM Transactions on Graphics (TOG)* 40.4, pp. 1–13.

Li, Yanan, Carboni, Gerolamo, Gonzalez, Franck, Campolo, Domenico, and Burdet, Etienne (2019). "Differential game theory for versatile physical human–robot interaction". In: *Nature Machine Intelligence* 1.1, pp. 36–43.

Lim, Bokman, Hyoung, Seungyong, Lee, Jusuk, Seo, Keehong, Jang, Junwon, and Shim, Youngbo (2017). "Simulating gait assistance of a hip exoskeleton: Case studies for ankle pathologies". In: *2017 IEEE international conference on robotics and automation (ICRA)*. IEEE, pp. 1022–1027.

Lindén, Henrik, Petersen, Peter C, Vestergaard, Mikkel, and Berg, Rune W (2022). "Movement is governed by rotational neural dynamics in spinal motor networks". In: *Nature* 610.7932, pp. 526–531.

Liu, C Karen and Negrut, Dan (2021). "The role of physics-based simulators in robotics". In: *Annual Review of Control, Robotics, and Autonomous Systems* 4.1.

Liu, Jiaqing, Wang, Can, He, Bailin, Li, Pengbo, and Wu, Xinyu (2022). "Metric Learning for Robust Gait Phase Recognition for a Lower Limb Exoskeleton Robot Based on sEMG". In: *IEEE Transactions on Medical Robotics and Bionics*.

Lock, Blair, Englehart, Kevin, and Hudgins, Bernard (2005). "Real-time myoelectric control in a virtual environment to relate usability vs. accuracy". In: *MyoElectric Controls Symposium*. Citeseer, pp. 122–127.

Lopez-Delis, A, Delisle-Rodriguez, D, Villa-Parra, Ana Cecilia, and Bastos-Filho, T (2015). "Knee motion pattern classification from trunk muscle based on sEMG signals". In: *2015 37th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*. IEEE, pp. 2604–2607.

Luo, Shuzhen, Jiang, Menghan, Zhang, Sainan, Zhu, Junxi, Yu, Shuangyue, Dominguez Silva, Israel, Wang, Tian, Rouse, Elliott, Zhou, Bolei, Yuk, Hyunwoo, et al. (2024). "Experiment-free exoskeleton assistance via learning in simulation". In: *Nature* 630.8016, pp. 353–359.

Luo, Ying-Sheng, Soeseno, Jonathan Hans, Chen, Trista Pei-Chun, and Chen, Wei-Chao (2020). "CARL: Controllable Agent with Reinforcement Learning for Quadruped Locomotion". In: *ACM Transactions on Graphics (Proceedings of SIGGRAPH 2020)* 39.4.

Ma, Shihan, Clarke, Alexander Kenneth, Maksymenko, Kostiantyn, Deslauriers-Gauthier, Samuel, Sheng, Xinjun, Zhu, Xiangyang, and Farina, Dario (2024a). "Conditional generative models for simulation of EMG during naturalistic movements". In: *IEEE Transactions on Neural Networks and Learning Systems*.

Ma, Shihan, Mendez Guerra, Irene, Caillet, Arnault Hubert, Zhao, Jiamin, Clarke, Alexander Kenneth, Maksymenko, Kostiantyn, Deslauriers-Gauthier, Samuel, Sheng, Xinjun, Zhu, Xiangyang, and Farina, Dario (2024b). "NeuroMotion: Open-source platform with neuromechanical and deep network modules to generate surface EMG signals during voluntary movement". In: *PLOS Computational Biology* 20.7, e1012257.

Maroger, Isabelle, Ramuzat, Noëlie, Stasse, Olivier, and Watier, Bruno (2021). "Human Trajectory Prediction Model and Its Coupling With a Walking Pattern Generator of a Humanoid Robot". In: *IEEE Robotics and Automation Letters* 6.4, pp. 6361–6369. DOI: 10.1109/LRA.2021.3092750.

Martin, Anne E. and Gregg, Robert D. (2017). "Stable, Robust Hybrid Zero Dynamics Control of Powered Lower-Limb Prostheses". In: *IEEE Transactions on Automatic Control* 62.8, pp. 3930–3942. DOI: 10.1109/TAC.2017.2648040.

McDonald, Cody L, Westcott-McCoy, Sarah, Weaver, Marcia R, Haagsma, Juanita, and Kartin, Deborah (2021). "Global prevalence of traumatic non-fatal limb amputation". In: *Prosthetics and orthotics international* 45.2, pp. 105–114.

Mehdizadeh, Sina (2018). "The largest Lyapunov exponent of gait in young and elderly individuals: A systematic review". In: *Gait & posture* 60, pp. 241–250.

Mello, Roger GT, Oliveira, Liliam F, and Nadal, Jurandir (2007). "Digital Butterworth filter for subtracting noise from low magnitude surface electromyogram". In: *Computer methods and programs in biomedicine* 87.1, pp. 28–35.

Meredith, Maddock, Maddock, Steve, et al. (2001). "Motion capture file formats explained". In: *Department of Computer Science, University of Sheffield* 211, pp. 241–244.

Merel, Josh, Botvinick, Matthew, and Wayne, Greg (2019). "Hierarchical motor control in mammals and machines". en. In: *Nature Communications* 10.1, p. 5489. DOI: 10.1038/s41467-019-13239-6.

Merel, Josh, Hasenclever, Leonard, Galashov, Alexandre, Ahuja, Arun, Pham, Vu, Wayne, Greg, Teh, Yee Whye, and Heess, Nicolas (2018). "Neural probabilistic motor primitives for humanoid control". In: *arXiv preprint arXiv:1811.11711*.

Merel, Josh, Tassa, Yuval, TB, Dhruva, Srinivasan, Sriram, Lemmon, Jay, Wang, Ziyu, Wayne, Greg, and Heess, Nicolas (2017). "Learning human behaviors from motion capture by adversarial imitation". In: *arXiv preprint arXiv:1707.02201*.

Merletti, R. and Muceli, S. (2019). "Tutorial. Surface EMG detection in space and time: Best practices". In: *Journal of electromyography and kinesiology* 49, p. 102363.

Millán, José del R, Rupp, Rüdiger, Müller-Putz, Gernot, Murray-Smith, Roderick, Giugliemma, Claudio, Tangermann, Michael, Vidaurre, Carmen, Cincotti, Febo, Kubler, Andrea, Leeb, Robert, et al. (2010). "Combining brain–computer interfaces and assistive technologies: state-of-the-art and challenges". In: *Frontiers in neuroscience* 4, p. 161.

Mitchell, Natasha (2018). *ISRCTN - ISRCTN15043643: STEPFORWARD: Patient acceptability of a novel prosthetic device*. en. DOI: 10.1186/ISRCTN15043643. URL: https://www.isrctn.com/ISRCTN15043643 (visited on 09/26/2024).

Moisio, Kirsten C, Sumner, Dale R, Shott, Susan, and Hurwitz, Debra E (2003). "Normalization of joint moments during gait: a comparison of two techniques". In: *Journal of biomechanics* 36.4, pp. 599–603.

Moissenet, Florent and Schreiber, Céline (2019). *A multimodal dataset of human gait at different walking speeds established on injury-free adult participants*. https://figshare.com/articles/dataset/A_multimodal_dataset_of_human_gait_at_different_walking_speeds/7734767/8. Online; accessed 09-Jan-2021. DOI: 10.6084/m9.figshare.7734767.v8.

Mombaur, Katja, Truong, Anh, and Laumond, Jean-Paul (2010). "From human to humanoid locomotion—an inverse optimal control approach". In: *Autonomous robots* 28, pp. 369–383.

Mourot, Lucas, Hoyet, Ludovic, Le Clerc, François, Schnitzler, François, and Hellier, Pierre (2022). "A Survey on Deep Learning for Skeleton-Based Human Animation". In: *Computer Graphics Forum*. Vol. 41. Wiley Online Library, pp. 122–157.

Müller, Meinard, Röder, Tido, Clausen, Michael, Eberhardt, Bernhard, Krüger, Björn, and Weber, Andreas (2007). "Mocap database hdm05". In: *Institut für Informatik II, Universität Bonn* 2.7.

Murley, George S., Menz, Hylton B., Landorf, Karl B., and Bird, Adam R. (2009). "Reliability of lower limb electromyography during overground walking: A comparison of maximal- and sub-maximal normalisation techniques". In: *Journal of biomechanics* 43.4, pp. 749–756.

Murtagh, Elaine M, Mair, Jacqueline L, Aguiar, Elroy, Tudor-Locke, Catrine, and Murphy, Marie H (2021). "Outdoor walking speeds of apparently healthy adults: A systematic review and meta-analysis". In: *Sports Medicine* 51.1, pp. 125–141.

Näf, Matthias B, Junius, Karen, Rossini, Marco, Rodriguez-Guerrero, Carlos, Vanderborght, Bram, and Lefeber, Dirk (2018). "Misalignment compensation for full human-exoskeleton kinematic compatibility: State of the art and evaluation". In: *Applied Mechanics Reviews* 70.5, p. 050802.

Nakamura, Go, Shibanoki, Taro, Kurita, Yuichi, Honda, Yuichiro, Masuda, Akito, Mizobe, Futoshi, Chin, Takaaki, and Tsuji, Toshio (2017). "A virtual myoelectric prosthesis training system capable of providing instructions on hand operations". In: *International Journal of Advanced Robotic Systems* 14.5.

National Health Service (2019). *NHS Obesity Guide*. `https://www.nhs.uk/conditions/obesity/`. Online; accessed 10-Jan-2021.

Nawfel, Jena L., Englehart, Kevin B., and Scheme, Erik J. (2022). "The Influence of Training With Visual Biofeedback on the Predictability of Myoelectric Control Usability". In: *IEEE Transactions on Neural Systems and Rehabilitation Engineering* 30, pp. 878–892. DOI: `10.1109/TNSRE.2022.3162421`.

Nguyen, Vinh Q, Johnson, Russell T, Sup, Frank C, and Umberger, Brian R (2019). "Bilevel optimization for cost function determination in dynamic simulation of human gait". In: *IEEE Transactions on Neural Systems and Rehabilitation Engineering* 27.7, pp. 1426–1435.

Nolan, Lee, Wit, Andrzej, Dudziñski, Krzysztof, Lees, Adrian, Lake, Mark, and Wychowañski, Michał (2003). "Adjustments in gait symmetry with walking speed in trans-femoral and trans-tibial amputees". In: *Gait & posture* 17.2, pp. 142–151.

Ofli, Ferda, Chaudhry, Rizwan, Kurillo, Gregorij, Vidal, René, and Bajcsy, Ruzena (2013). "Berkeley mhad: A comprehensive multimodal human action database". In: *2013 IEEE Workshop on Applications of Computer Vision (WACV)*. IEEE, pp. 53–60.

Orendurff, Michael S, Segal, Ava D, Berge, Jocelyn S, Flick, Kevin C, Spanier, David, and Klute, Glenn K (2006). "The kinematics and kinetics of turning: limb asymmetries associated with walking a circular path". In: *Gait & posture* 23.1, pp. 106–111.

Oskoei, Mohammadreza Asghari and Hu, Huosheng (2007a). "Myoelectric control systems — A survey". In: *Biomedical Signal Processing and Control* 2.4, pp. 275–294.

— (2007b). "Myoelectric control systems-A survey". In: *Biomedical signal processing and control* 2.4, pp. 275–294.

OT Bioelettronica (2022). *OTBioLab user manual v3.4*. `https://www.otbioelettronica.it/files/47/Software/15/OTBioLab-User-Manual-ENG.pdf`. [Online; accessed 19-Jan-2022].

Pace, Anna, Grioli, Giorgio, Ghezzi, Alice, Bicchi, Antonio, and Catalano, Manuel Giuseppe (2024). "Investigating the performance of soft robotic adaptive feet with longitudinal and transverse arches". In: *Frontiers in Robotics and AI* 11, p. 1375515.

Pace, Anna, Proksch, Lukas, Grioli, Giorgio, Aszmann, Oskar C, Bicchi, Antonio, and Catalano, Manuel G (2023). "An Experimental Setup to Test Obstacle-Dealing Capabilities of Prosthetic Feet". In: *2023 International Conference on Rehabilitation Robotics (ICORR)*. IEEE, pp. 1–6.

Park, Jungnam, Min, Sehee, Chang, Phil Sik, Lee, Jaedong, Park, Moon Seok, and Lee, Jehee (2022). "Generative GaitNet". In: *ACM SIGGRAPH 2022 Conference Proceedings*, pp. 1–9.

Park, Soohwan, Ryu, Hoseok, Lee, Seyoung, Lee, Sunmin, and Lee, Jehee (2019). "Learning predict-and-simulate policies from unorganized human motion data". In: *ACM Transactions on Graphics (TOG)* 38.6, pp. 1–11.

Pavllo, Dario, Feichtenhofer, Christoph, Auli, Michael, and Grangier, David (2020). "Modeling human motion with quaternion-based neural networks". In: *International Journal of Computer Vision* 128, pp. 855–872.

Pawłowski, Maciej, Wróblewska, Anna, and Sysko-Romańczuk, Sylwia (2023). "Effective techniques for multimodal data fusion: A comparative analysis". In: *Sensors* 23.5, p. 2381.

Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., and Duchesnay, E. (2011). "Scikit-learn: Machine Learning in Python". In: *Journal of Machine Learning Research* 12, pp. 2825–2830.

Peeraer, Louis, Aeyels, B, and Van der Perre, Georges (1990). "Development of EMG-based mode and intent recognition algorithms for a computer-controlled above-knee prosthesis". In: *Journal of biomedical engineering* 12.3, pp. 178–182.

Peng, Xue Bin, Abbeel, Pieter, Levine, Sergey, and Van de Panne, Michiel (2018). "Deepmimic: Example-guided deep reinforcement learning of physics-based character skills". In: *ACM Transactions On Graphics (TOG)* 37.4, pp. 1–14.

Peng, Xue Bin, Coumans, Erwin, Zhang, Tingnan, Lee, Tsang-Wei, Tan, Jie, and Levine, Sergey (2020). "Learning agile robotic locomotion skills by imitating animals". In: *arXiv preprint arXiv:2004.00784*.

Peng, Xue Bin, Ma, Ze, Abbeel, Pieter, Levine, Sergey, and Kanazawa, Angjoo (2021). "Amp: Adversarial motion priors for stylized physics-based character control". In: *ACM Transactions on Graphics (TOG)* 40.4, pp. 1–20.

Perry, Briana N, Armiger, Robert S, Yu, Kristin E, Alattar, Ali A, Moran, Courtney W, Wolde, Mikias, McFarland, Kayla, Pasquina, Paul F, and Tsao, Jack W (2018). "Virtual integration environment as an advanced prosthetic limb training platform". In: *Frontiers in neurology* 9, p. 785.

Peters, Maria, Crowe, John, Piéri, Jean-Francois, Quartero, Hendrik, Hayes-Gill, Barrie, James, David, Stinstra, Jeroen, and Shakespeare, Simon (2001). "Monitoring the fetal heart noninvasively: a review of methods". In: *Journal of perinatal medicine* 29.5, pp. 408–416.

Petrini, Francesco Maria, Valle, Giacomo, Bumbasirevic, Marko, Barberi, Federica, Bortolotti, Dario, Cvancara, Paul, Hiairrassary, Arthur, Mijovic, Pavle, Sverrisson, Atli Örn, Pedrocchi,

Alessandra, et al. (2019). "Enhancing functional abilities and cognitive integration of the lower limb prosthesis". In: *Science translational medicine* 11.512, eaav8939.

Pew, Corey A, Roelker, Sarah A, Klute, Glenn K, and Neptune, Richard R (2020). "Analysis of the relative motion between the socket and residual limb in transtibial amputees while wearing a transverse rotation adapter". In: *Journal of Applied Biomechanics* 37.1, pp. 21–29.

Pfeifer, Serge, Riener, Robert, and Vallery, Heike (2014). "Knee stiffness estimation in physiological gait". In: *2014 36th Annual International Conference of the IEEE Engineering in Medicine and Biology Society*. IEEE, pp. 1607–1610.

Pham, Hai-Trieu, Kim, Jung-Ja, Nguyen, Tan Loc, and Won, Yonggwan (2015). "3D motion matching algorithm using signature feature descriptor". In: *Multimedia tools and applications* 74.3, pp. 1125–1136.

Phillips, Betsy and Zhao, Hongxin (1993). "Predictors of assistive technology abandonment". In: *Assistive technology* 5.1, pp. 36–45.

Piazza, Cristina, Della Santina, Cosimo, Grioli, Giorgio, Bicchi, Antonio, and Catalano, Manuel G (2024). "Analytical Model and Experimental Testing of the SoftFoot: An Adaptive Robot Foot for Walking Over Obstacles and Irregular Terrains". In: *IEEE Transactions on Robotics*.

Pradhan, Ashirbad, Malagon, Gemma, Lagacy, Rebecca, Chester, Victoria, and Kuruganti, Usha (2020). "Effect of age and sex on strength and spatial electromyography during knee extension". In: *Journal of physiological anthropology* 39, pp. 1–11.

Quintero, Hugo, Farris, Ryan, Hartigan, Clare, Clesson, Ismari, and Goldfarb, Michael (2011). "A powered lower limb orthosis for providing legged mobility in paraplegic individuals". In: *Topics in spinal cord injury rehabilitation* 17.1, pp. 25–33.

Radosavovic, Ilija, Xiao, Tete, Zhang, Bike, Darrell, Trevor, Malik, Jitendra, and Sreenath, Koushil (2024). "Real-world humanoid locomotion with reinforcement learning". In: *Science Robotics* 9.89, eadi9579. DOI: 10.1126/scirobotics.adi9579.

Rainoldi, A., Melchiorri, G., and Caruso, I. (2004). "A method for positioning electrodes during surface EMG recordings in lower limb muscles". In: *Journal of neuroscience methods* 134.1, pp. 37–43.

Ramachandran, Prajit, Zoph, Barret, and Le, Quoc V (2017). "Searching for activation functions". In: *arXiv preprint arXiv:1710.05941*.

Reda, Daniele, Won, Jungdam, Ye, Yuting, Panne, Michiel van de, and Winkler, Alexander (2023). "Physics-based Motion Retargeting from Sparse Inputs". In: *Proceedings of the ACM on Computer Graphics and Interactive Techniques* 6.3, pp. 1–19.

Riener, Robert (2016). "The Cybathlon promotes the development of assistive technology for people with physical disabilities". In: *Journal of neuroengineering and rehabilitation* 13.1, p. 49.

Rodriguez-Tapia, Bernabe, Soto, Israel, Martínez, Daniela M, and Arballo, Norma Candolfi (2020). "Myoelectric interfaces and related applications: Current state of EMG signal processing–A systematic review". In: *IEEE Access* 8, pp. 7792–7805.

Rohmer, E., Singh, S. P. N., and Freese, M. (2013). "CoppeliaSim (formerly V-REP): a Versatile and Scalable Robot Simulation Framework". In: *Proc. of The International Conference on Intelligent Robots and Systems (IROS)*. www.coppeliarobotics.com.

Rosenstein, Michael T, Collins, James J, and De Luca, Carlo J (1993). "A practical method for calculating largest Lyapunov exponents from small data sets". In: *Physica D: Nonlinear Phenomena* 65.1-2, pp. 117–134.

Rudenko, Andrey, Palmieri, Luigi, Herman, Michael, Kitani, Kris M, Gavrila, Dariu M, and Arras, Kai O (2020). "Human motion trajectory prediction: A survey". In: *The International Journal of Robotics Research* 39.8, pp. 895–935.

Rusu, Andrei A, Vecerik, Mel, Rothörl, Thomas, Heess, Nicolas, Pascanu, Razvan, and Hadsell, Raia (2016). "Sim-to-real robot learning from pixels with progressive nets". In: *arXiv preprint arXiv:1610.04286*.

Sacco, Isabel C. N., Gomes, Aline A., Otuzi, Mitie E., Pripas, Denise, and Onodera, Andrea N. (2009). "A method for better positioning bipolar electrodes for lower limb EMG recordings during dynamic contractions". In: *Journal of neuroscience methods* 180.1, pp. 133–137.

Sardain, Philippe and Bessonnet, Guy (2004). "Forces acting on a biped robot. Center of pressure-zero moment point". In: *IEEE Transactions on Systems, Man, and Cybernetics-Part A: Systems and Humans* 34.5, pp. 630–637.

Schaik, Loeke van, Geertzen, Jan HB, Dijkstra, Pieter U, and Dekker, Rienk (2019). "Metabolic costs of activities of daily living in persons with a lower limb amputation: A systematic review and meta-analysis". In: *PloS one* 14.3, e0213256.

Scheme, Erik J. and Englehart, Kevin B. (2013). "Validation of a Selective Ensemble-Based Classification Scheme for Myoelectric Control Using a Three-Dimensional Fitts' Law Test". In: *IEEE Transactions on Neural Systems and Rehabilitation Engineering* 21.4, pp. 616–623. DOI: 10.1109/TNSRE.2012.2226189.

Schulman, John, Wolski, Filip, Dhariwal, Prafulla, Radford, Alec, and Klimov, Oleg (2017). "Proximal policy optimization algorithms". In: *arXiv preprint arXiv:1707.06347*.

Schweitzer, Wolf, Thali, Michael J, and Egger, David (2018). "Case-study of a user-driven prosthetic arm design: bionic hand versus customized body-powered technology in a highly demanding work environment". In: *Journal of neuroengineering and rehabilitation* 15.1, p. 1.

Seymour, Ron (2002). *Prosthetics and orthotics: lower limb and spinal*. Lippincott Williams & Wilkins.

Shah, Syed Ghulam Sarwar, Robinson, Ian, and AlShawi, Sarmad (2009). "Developing medical device technologies from users' perspectives: a theoretical framework for involving users in the development process". In: *International journal of technology assessment in health care* 25.4, pp. 514–521.

Shamaei, Kamran, Sawicki, Gregory S, and Dollar, Aaron M (2013). "Estimation of quasi-stiffness of the human hip in the stance phase of walking". In: *PloS one* 8.12, e81841.

Sharifi Renani, Mohsen, Eustace, Abigail M, Myers, Casey A, and Clary, Chadd W (2021). "The use of synthetic imu signals in the training of deep learning models significantly improves the accuracy of joint kinematic predictions". In: *Sensors* 21.17, p. 5876.

Sheng, Bo, Tang, Lihua, Xie, Shengquan, Deng, Chao, and Zhang, Yanxin (2019). "Alterations in muscle activation patterns during robot-assisted bilateral training: A pilot study". In: *Proceedings of the Institution of Mechanical Engineers. Part H, Journal of engineering in medicine* 233.2, pp. 219–231.

Shenoy, Pradeep, Miller, Kai J, Crawford, Beau, and Rao, Rajesh PN (2008). "Online electromyographic control of a robotic prosthesis". In: *IEEE transactions on biomedical engineering* 55.3, pp. 1128–1135.

Sherman, Michael A, Seth, Ajay, and Delp, Scott L (2011). "Simbody: multibody dynamics for biomedical research". In: *Procedia Iutam* 2, pp. 241–261.

Al-Shuka, Hayder FN, Rahman, Mohammad H, Leonhardt, Steffen, Ciobanu, Ileana, and Berteanu, Mihai (2019). "Biomechanics, actuation, and multi-level control strategies of power-augmentation lower extremity exoskeletons: An overview". In: *International Journal of Dynamics and Control* 7, pp. 1462–1488.

Shultz, Amanda Huff, Mitchell, Jason E, Truex, Don, Lawson, Brian E, and Goldfarb, Michael (2013). "Preliminary evaluation of a walking controller for a powered ankle prosthesis". In: *2013 IEEE International Conference on Robotics and Automation*. IEEE, pp. 4838–4843.

Shushtari, Mohammad, Nasiri, Rezvan, and Arami, Arash (2022). "Online Reference Trajectory Adaptation: A Personalized Control Strategy for Lower Limb Exoskeletons". In: *IEEE Robotics and Automation Letters* 7.1, pp. 128–134. DOI: 10.1109/LRA.2021.3115572.

Simao, Miguel, Mendes, Nuno, Gibaru, Olivier, and Neto, Pedro (2019). "A review on electromyography decoding and pattern recognition for human-machine interaction". In: *Ieee Access* 7, pp. 39564–39582.

Simon, Ann M., Hargrove, Levi J., Lock, Blair A., and Kuiken, Todd A. (2011). "The Target Achievement Control Test: Evaluating real-time myoelectric pattern recognition control of a multifunctional upper-limb prosthesis". In: *Journal of rehabilitation research and development* 48.6, pp. 619–627.

Simon, Ann M., Ingraham, Kimberly A, Fey, Nicholas P, Finucane, Suzanne B, Lipschutz, Robert D, Young, Aaron J, and Hargrove, Levi J (2014). "Configuring a powered knee and ankle prosthesis for transfemoral amputees within five specific ambulation modes". In: *PloS one* 9.6, e99387.

Simpson, Lisa A, Eng, Janice J, Hsieh, Jane TC, Wolfe, and Spinal Cord Injury Rehabilitation Evidence (SCIRE) Research Team, Dalton L the (2012). "The health and life priorities of individuals with spinal cord injury: a systematic review". In: *Journal of neurotrauma* 29.8, pp. 1548–1555.

Sinke, Maaike, Chadwell, Alix, and Smit, Gerwin (2022). "State of the art of prosthesis simulators for the upper limb: a narrative review". In: *Annals of Physical and Rehabilitation Medicine* 65.6, p. 101635.

Smith, Lauren H., Hargrove, Levi J., Lock, Blair A., and Kuiken, Todd A. (2011). "Determining the Optimal Window Length for Pattern Recognition-Based Myoelectric Control: Balancing the Competing Effects of Classification Error and Controller Delay". In: *IEEE Transactions on Neural Systems and Rehabilitation Engineering* 19.2, pp. 186–192. DOI: 10.1109/TNSRE.2010.2100828.

Soares, Alcimar, Andrade, Adriano, Lamounier, Edgard, and Carrijo, Renato (2003). "The development of a virtual myoelectric prosthesis controlled by an EMG pattern recognition system based on neural networks". In: *Journal of Intelligent Information Systems* 21, pp. 127–141.

Song, Seungmoon and Geyer, Hartmut (2015). "A neural circuitry that emphasizes spinal feedback generates diverse behaviours of human locomotion". en. In: *The Journal of Physiology* 593.16, pp. 3493–3511. DOI: 10.1113/JP270228.

Song, Seungmoon, Kidziński, Łukasz, Peng, Xue Bin, Ong, Carmichael, Hicks, Jennifer L., Levine, Sergey, Atkeson, Christopher G, and Delp, Scott L (2021). "Deep reinforcement learning for modeling human locomotion control in neuromechanical simulation". In: *Journal of neuroengineering and rehabilitation* 18.1, pp. 1–17.

Srinivasan, SS, Carty, MJ, Calvaresi, PW, Clites, TR, Maimon, BE, Taylor, CR, Zorzos, AN, and Herr, H (2017). "On prosthetic control: A regenerative agonist-antagonist myoneural interface". In: *Science Robotics* 2.6.

Stango, Antonietta, Negro, Francesco, and Farina, Dario (2015). "Spatial Correlation of High Density EMG Signals Provides Features Robust to Electrode Number and Shift in Pattern Recognition for Myocontrol". In: *IEEE transactions on neural systems and rehabilitation engineering* 23.2, pp. 189–198.

Starke, Sebastian, Zhang, He, Komura, Taku, and Saito, Jun (2019). "Neural state machine for character-scene interactions." In: *ACM Trans. Graph.* 38.6, pp. 209–1.

Staudenmann, Didier, Roeleveld, Karin, Stegeman, Dick F., and van Dieën, Jaap H. (2010). "Methodological aspects of SEMG recordings for force estimation – A tutorial and review". In: *Journal of Electromyography and Kinesiology* 20.3, pp. 375–387. DOI: https://doi.org/10.1016/j.jelekin.2009.08.005.

Steinhardt, Cynthia R., Bettthauser, Joseph, Hunt, Christopher, and Thakor, Nitish (2018). "Registration of EMG Electrodes to Reduce Classification Errors due to Electrode Shift". In: IEEE, pp. 1–4.

Subbaswamy, Adarsh and Saria, Suchi (2020). "From development to deployment: dataset shift, causality, and shift-stable models in health AI". In: *Biostatistics* 21.2, pp. 345–352.

Sunehag, Peter, Lever, Guy, Gruslys, Audrunas, Czarnecki, Wojciech Marian, Zambaldi, Vinicius, Jaderberg, Max, Lanctot, Marc, Sonnerat, Nicolas, Leibo, Joel Z, Tuyls, Karl, et al. (2017). "Value-decomposition networks for cooperative multi-agent learning". In: *arXiv preprint arXiv:1706.05296*.

Sutton, Richard S and Barto, Andrew G (2018). *Reinforcement learning: An introduction*. MIT press.

Sylos-Labini, F., Scaleia, V. La, Avella, A., Pisotta, I., Tamburella, F., Scivoletto, G., Molinari, M., Wang, S., Wang, L., Asseldonk, E. Van, Kooij, H. Van der, Hoellinger, T., Cheron, G., Thorsteinsson, F., Ilzkovitz, M., Gancet, J., Hauffe, R., Zanov, F., Lacquaniti, F., and Ivanenko, Y. P. (2014). "EMG patterns during assisted walking in the exoskeleton". In: *Frontiers in human neuroscience* 8, p. 423.

Tan, Jie, Liu, Karen, and Turk, Greg (2011). "Stable proportional-derivative controllers". In: *IEEE Computer Graphics and Applications* 31.4, pp. 34–44.

Tassa, Yuval, Erez, Tom, and Todorov, Emanuel (2012). "Synthesis and stabilization of complex behaviors through online trajectory optimization". In: *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, pp. 4906–4913.

Tessler, Chen, Kasten, Yoni, Guo, Yunrong, Mannor, Shie, Chechik, Gal, and Peng, Xue Bin (2023). "Calm: Conditional adversarial latent models for directable virtual characters". In: *ACM SIGGRAPH 2023 Conference Proceedings*, pp. 1–9.

Tkach, DC, Lipschutz, Robert D, Finucane, Suzanne B, and Hargrove, Levi J (2013). "Myoelectric neural interface enables accurate control of a virtual multiple degree-of-freedom foot-ankle prosthesis". In: *2013 IEEE 13th International Conference on Rehabilitation Robotics (ICORR)*. IEEE, pp. 1–4.

Todorov, Emanuel, Erez, Tom, and Tassa, Yuval (2012). "Mujoco: A physics engine for model-based control". In: *2012 IEEE/RSJ international conference on intelligent robots and systems*. IEEE, pp. 5026–5033.

Tomovic, R and McGhee, Robert B (1966). "A finite state approach to the synthesis of bioengineering control systems". In: *IEEE Transactions on Human Factors in Electronics* HFE-7.2, pp. 65–69.

Torricelli, D and Pons, Jose L (2018). "Eurobench: Preparing robots for the real world". In: *International Symposium on Wearable Robotics*. Springer, pp. 375–378.

Trent, Lauren, Intintoli, Michelle, Prigge, Pat, Bollinger, Chris, Walters, Lisa Smurr, Conyers, Dan, Miguelez, John, and Ryan, Tiffany (2020). "A narrative review: current upper limb prosthetic options and design". In: *Disability and Rehabilitation: Assistive Technology*.

Tschiedel, Michael, Russold, Michael Friedrich, and Kaniusas, Eugenijus (2020). "Relying on more sense for enhancing lower limb prostheses control: a review". In: *Journal of Neuro-Engineering and Rehabilitation* 17.1, pp. 1–13.

Tucker, Michael R, Olivier, Jeremy, Pagel, Anna, Bleuler, Hannes, Bouri, Mohamed, Lambercy, Olivier, Millán, José del R, Riener, Robert, Vallery, Heike, and Gassert, Roger (2015). "Control strategies for active lower extremity prosthetics and orthotics: a review". In: *Journal of neuroengineering and rehabilitation* 12.1, pp. 1–30.

Tunyasuvunakool, Saran, Muldal, Alistair, Doron, Yotam, Liu, Siqi, Bohez, Steven, Merel, Josh, Erez, Tom, Lillicrap, Timothy, Heess, Nicolas, and Tassa, Yuval (2020). "dm_control: Software and tasks for continuous control". In: *Software Impacts* 6, p. 100022.

Ubisoft La Forge (2020). *Ubisoft La Forge Animation Dataset LAFAN1*. `https://github.com/ubisoft/ubisoft-laforge-animation-dataset`. Online; accessed 18-Jun-2021.

Unity Technologies (2021). *Unity User Manual*. https://docs.unity3d.com/Manual/. Online; accessed 12-Jan-2021.

— (2022). *Unity - Manual: State Machine Basics*. en. URL: https://docs.unity3d.com/2022.3/Documentation/Manual/StateMachineBasics.html (visited on 04/11/2024).

Urbanchek, Melanie G, Kung, Theodore A, Frost, Christopher M, Martin, David C, Larkin, Lisa M, Wollstein, Adi, and Cederna, Paul S (2016). "Development of a regenerative peripheral nerve interface for control of a neuroprosthetic limb". In: *BioMed research international* 2016.

Varol, Huseyin Atakan, Sup, Frank, and Goldfarb, Michael (2009). "Multiclass real-time intent recognition of a powered lower limb prosthesis". In: *IEEE Transactions on Biomedical Engineering* 57.3, pp. 542–551.

Vaughan, Christopher L, Davis, Brian L, and O'connor, Jeremy C (1992). *Dynamics of human gait*. Vol. 2. Human Kinetics.

Venugopal, G, Navaneethakrishna, M, and Ramakrishnan, S (2014). "Extraction and analysis of multiple time window features associated with muscle fatigue conditions using sEMG signals". In: *Expert Systems with Applications* 41.6, pp. 2652–2659.

Verweij, David, Esteves, Augusto, Khan, Vassilis-Javed, and Bakker, Saskia (2017). "WaveTrace: motion matching input using wrist-worn motion sensors". In: *Proceedings of the 2017 CHI Conference Extended Abstracts on Human Factors in Computing Systems*, pp. 2180–2186.

Vicon Motion Systems Ltd (2020). *Vicon Nexus Documentation*. https://docs.vicon.com/display/Nexus212/Nexus+Documentation. [Online; accessed 17-Jan-2022].

Vicon Motion Systems Ltd, UK (2021). *Vicon Vantage Motion Capture Camera*. https://www.vicon.com/hardware/cameras/vantage/. Online; accessed 16-Jul-2021.

Vienne, Alienor, Barrois, Remi P, Buffat, Stephane, Ricard, Damien, and Vidal, Pierre-Paul (2017). "Inertial sensors to assess gait quality in patients with neurological disorders: a systematic review of technical and analytical challenges". In: *Frontiers in psychology* 8, p. 817.

Wang, Bo and Chen, Tao (2015). "Gaussian process regression with multiple response variables". In: *Chemometrics and Intelligent Laboratory Systems* 142, pp. 159–165.

Wang, Chia-Yu E, Bobrow, James E, and Reinkensmeyer, David J (2005). "Dynamic motion planning for the design of robotic gait rehabilitation". In: *Journal of biomechanical engineering* 127.4, pp. 672–679.

Wang, Huawei, Caggiano, Vittorio, Durandau, Guillaume, Sartori, Massimo, and Kumar, Vikash (2022). "MyoSim: Fast and physiologically realistic MuJoCo models for musculoskeletal and exoskeletal studies". In: *2022 International Conference on Robotics and Automation (ICRA)*. IEEE, pp. 8104–8111.

Wang, Qian, Zhang, Xu, Chen, Xiang, Chen, Ruizhi, Chen, Wei, and Chen, Yuwei (2010). "A novel pedestrian dead reckoning algorithm using wearable EMG sensors to measure walking

strides". In: *2010 Ubiquitous Positioning Indoor Navigation and Location Based Service*. IEEE, pp. 1–8.

Wang, Ziyu, Merel, Josh S, Reed, Scott E, Freitas, Nando de, Wayne, Gregory, and Heess, Nicolas (2017). "Robust imitation of diverse behaviors". In: *Advances in Neural Information Processing Systems*, pp. 5320–5329.

Wentink, EC, Schut, VGH, Prinsen, EC, Rietman, Johan S, and Veltink, Peter H (2014). "Detection of the onset of gait initiation using kinematic sensors and EMG in transfemoral amputees". In: *Gait & posture* 39.1, pp. 391–396.

Windrich, Michael, Grimmer, Martin, Christ, Oliver, Rinderknecht, Stephan, and Beckerle, Philipp (2016). "Active lower limb prosthetics: a systematic review of design issues and solutions". In: *Biomedical engineering online* 15, pp. 5–19.

Winter, David A (1984). "Kinematic and kinetic patterns in human gait: variability and compensating effects". In: *Human movement science* 3.1-2, pp. 51–76.

Won, Jungdam, Gopinath, Deepak, and Hodgins, Jessica (2020). "A scalable approach to control diverse behaviors for physically simulated characters". In: *ACM Transactions on Graphics (TOG)* 39.4, pp. 33–1.

World Health Organization (2017). *WHO standards for prosthetics and orthotics*.

Wu, Amy R, Dzeladini, Florin, Brug, Tycho JH, Tamburella, Federica, Tagliamonte, Nevio L, Van Asseldonk, Edwin HF, Van Der Kooij, Herman, and Ijspeert, Auke J (2017). "An adaptive neuromuscular controller for assistive lower-limb exoskeletons: A preliminary study on subjects with spinal cord injury". In: *Frontiers in neurorobotics* 11, p. 30.

Wu, Le, Zhang, Xu, Wang, Kun, Chen, Xiang, and Chen, Xun (2020). "Improved High-Density Myoelectric Pattern Recognition Control Against Electrode Shift Using Data Augmentation and Dilated Convolutional Neural Network". In: *IEEE transactions on neural systems and rehabilitation engineering* 28.12, pp. 2637–2646.

Wu, Sai-Kit, Waycaster, Garrett, and Shen, Xiangrong (2011). "Electromyography-based control of active above-knee prostheses". In: *Control Engineering Practice* 19.8, pp. 875–882.

Xiang, Liangliang, Wang, Alan, Gu, Yaodong, Zhao, Liang, Shim, Vickie, and Fernandez, Justin (2022). "Recent machine learning progress in lower limb running biomechanics with wearable technology: A systematic review". In: *Frontiers in Neurorobotics* 16, p. 913052.

Xiong, Dezhen, Zhang, Daohui, Zhao, Xingang, and Zhao, Yiwen (2021). "Deep learning for EMG-based human-machine interaction: A review". In: *IEEE/CAA Journal of Automatica Sinica* 8.3, pp. 512–533.

Yan, Tingfang, Cempini, Marco, Oddo, Calogero Maria, and Vitiello, Nicola (2015). "Review of assistive strategies in powered lower-limb orthoses and exoskeletons". In: *Robotics and Autonomous Systems* 64, pp. 120–136.

Yang, Renyu, Zheng, Jianlin, and Song, Rong (2022). "Continuous mode adaptation for cable-driven rehabilitation robot using reinforcement learning". In: *Frontiers in Neurorobotics* 16, p. 1068706.

Young, Aaron J., Hargrove, Levi J., and Kuiken, Todd A. (2011). "The Effects of Electrode Size and Orientation on the Sensitivity of Myoelectric Pattern Recognition Systems to Electrode Shift". In: *IEEE transactions on biomedical engineering* 58.9, pp. 2537–2544.

— (2012). "Improving Myoelectric Pattern Recognition Robustness to Electrode Shift by Changing Interelectrode Distance and Electrode Configuration". In: *IEEE transactions on biomedical engineering* 59.3, pp. 645–652.

Zadziuk, Kristjan (2016). "Motion matching: The future of gameplay animation... today". In: *Proceedings of GDC*.

Zakka, Kevin, Tassa, Yuval, and MuJoCo Menagerie Contributors (2022). *MuJoCo Menagerie: A collection of high-quality simulation models for MuJoCo*. original-date: 2022-09-05T10:37:53Z. URL: `https://github.com/google-deepmind/mujoco_menagerie` (visited on 10/13/2024).

Zanghieri, Marcello, Benatti, Simone, Burrello, Alessio, Kartsch, Victor, Conti, Francesco, and Benini, Luca (2019). "Robust real-time embedded EMG recognition framework using temporal convolutional networks on a multicore IoT processor". In: *IEEE transactions on biomedical circuits and systems* 14.2, pp. 244–256.

Zeltzer, David (1982). "Motor control techniques for figure animation". In: *IEEE Computer Graphics and Applications* 2.09, pp. 53–59.

Zhai, Shumin and Woltjer, R. (2003). "Human movement performance in relation to path constraint - the law of steering in locomotion". In: *IEEE Virtual Reality, 2003. Proceedings.* Pp. 149–156. DOI: `10.1109/VR.2003.1191133`.

Zhang, Juanjuan, Fiers, Pieter, Witte, Kirby A, Jackson, Rachel W, Poggensee, Katherine L, Atkeson, Christopher G, and Collins, Steven H (2017). "Human-in-the-loop optimization of exoskeleton assistance during walking". In: *Science* 356.6344, pp. 1280–1284.

Zhang, Kaiqing, Yang, Zhuoran, and Başar, Tamer (2021). "Multi-agent reinforcement learning: A selective overview of theories and algorithms". In: *Handbook of Reinforcement Learning and Control*, pp. 321–384.

Zhang, Kuangen, Xiong, Caihua, Zhang, Wen, Liu, Haiyuan, Lai, Daoyuan, Rong, Yiming, and Fu, Chenglong (2019). "Environmental features recognition for lower limb prostheses toward predictive walking". In: *IEEE transactions on neural systems and rehabilitation engineering* 27.3, pp. 465–476.

Zhang, Xu, Wu, Le, Yu, Bin, Chen, Xiang, and Chen, Xun (2020). "Adaptive Calibration of Electrode Array Shifts Enables Robust Myoelectric Control". In: *IEEE transactions on biomedical engineering* 67.7, pp. 1947–1957.

Zhang, Yunbo, Gopinath, Deepak, Ye, Yuting, Hodgins, Jessica, Turk, Greg, and Won, Jungdam (2023). "Simulation and retargeting of complex multi-character interactions". In: *ACM SIGGRAPH 2023 Conference Proceedings*, pp. 1–11.

Zhao, Yongkun, Hodossy, Balint K, Jing, Shibo, Todoh, Masahiro, and Farina, Dario (2024). "Delayed reinforcement learning converges to intermittent control for human quiet stance". In: *Medical Engineering & Physics* 130, p. 104197.

Zhong, Boxuan, Huang, He, and Lobaton, Edgar (2022). "Reliable Vision-Based Grasping Target Recognition for Upper Limb Prostheses". In: *IEEE Transactions on Cybernetics* 52.3, pp. 1750–1762. DOI: 10.1109/TCYB.2020.2996960.

Zhou, Ping, Lock, Blair, and Kuiken, Todd A (2007). "Real time ECG artifact removal for myoelectric prosthesis control". In: *Physiological measurement* 28.4, p. 397.

Appendices

Appendix A

Motion Matching Implementation

Motion Matching is based on a nearest neighbour search, switching between animations in the motion capture dataset if they are a closer match to the current frame's feature vector than the one that would follow if the animation continues playing. As such, the choice of the feature space used for motion matching will determine the efficiency and quality of how the source data set is used. Importantly, this feature space must also capture in some form the "future outcome" of a given frame in the dataset, in addition to the current kinematic state. For example, the future position of the CoM can be used. Thanks to this, the matching process becomes directable: by constructing a query vector that uses the kinematic state features of the active animation frame, and appending novel "future outcome" features the nearest neighbour search will return matches that are similar to our intended trajectory but do not result in a large discontinuity in the kinematic state.

A.1 Feature extraction and labelling

Description of the user's current and desired kinematic state in a reduced but sufficiently information rich feature space is necessary for a mapping between the HL goals determined by the simulated intent and the associated animation. The specification of the feature set can be learnt or designed. Variational autoencoders (Higgins et al. 2016) or deep metric learning

(Kaya et al. 2019) would be two potential ways to learn a latent space of sEMG or kinematic state. Alternatively, phase space embeddings such as those used in analysis of system dynamics could also be employed (Rosenstein et al. 1993). However, designed feature sets can perform reliably, and require much less time for implementing or modifying them.

For this reason, the feature set described by (Bergamin et al. 2019) is adapted for pose description. The features describe each frame of motion capture in terms of current pose and future trajectory. These features are listed in Tables A.1 and A.2.

Symbol	Definition	$\mathbb{R}^n$
$\dot{\mathbf{p}}_{pelvis}$	Global velocity of the pelvis	$\mathbb{R}^3$
$\mathbf{p}^L_{lfoot}$	Position of the left foot, relative to pelvis (local position)	$\mathbb{R}^3$
$\mathbf{p}^L_{rfoot}$	Position of the right foot, relative to pelvis (local position)	$\mathbb{R}^3$
$\dot{\mathbf{p}}_{lfoot}$	Global velocity of the left foot	$\mathbb{R}^3$
$\dot{\mathbf{p}}_{rfoot}$	Global velocity of the right foot	$\mathbb{R}^3$
Total:		$\mathbb{R}^{15}$

Table A.1: Chosen feature set for current pose description. Matching these features reduces discontinuity when switching between frames.

It is not noted by (Bergamin et al. 2019) explicitly, however, all of these features are quantified in the right handed reference frame determined by:

1. The pelvis' frontal axis, projected on the horizontal plane.
2. The global 'up' axis.
3. A mutually orthogonal axis.

This change makes the features independent of global orientation. Other reference frames, such as one defined from the trajectory direction or the original pelvis reference frame are also viable

Symbol	Definition	$\mathbb{R}^n$
T	Future position of pelvis relative to current position. Sampled at $\frac{1}{3}$, $\frac{2}{3}$ and 1 seconds in the future. Projected onto ground plane.	$\mathbb{R}^{3\times2}$
D	Direction of pelvis trajectory, i.e. normalised velocity. Sampled at $\frac{1}{3}$, $\frac{2}{3}$ and 1 seconds in the future. Projected onto ground plane.	$\mathbb{R}^{3\times2}$
Total:		$\mathbb{R}^{12}$

Table A.2: Chosen feature set for future state description. Matching these features allows adhering to directable goals (i.e. user input in desired walking path).

options, as long as it is applied consistently during simulation as well.

The pelvis's position is used as an estimate of the CoM trajectory. While this approximation is significantly different from the true value (Eames et al. 1999), the strong correlation present is sufficient in practice and avoids the requirement for dynamically calculating the CoM.

A conversion from hierarchical kinematics to absolute position is performed by the method described by (Meredith et al. 2001). Differentiation for the feature set was performed using a 3rd order Savitzky-Golay filter. The features are also normalised and scaled according to developer defined weights which can define a trade-off between continuity of motion and responsiveness.

These features are extracted from recordings of the LaFAN (Ubisoft La Forge 2020), and the Artanim MM1 (Hodossy et al. 2024) motion capture dataset with Python scripts, and are stored in JSON files. The recordings are also converted from the human-readable, hierarchical file format (.bvh) in Blender (Community 2018) to an animation industry file format (.fbx), which are then applied to a 3D humanoid avatar (free-to-use model from Mixamo (Adobe Systems Incorporated 2019)). Each frame of extracted features has an exact recording and time point they can be mapped to. The connection between different formats and tools is shown in Figure A.1.

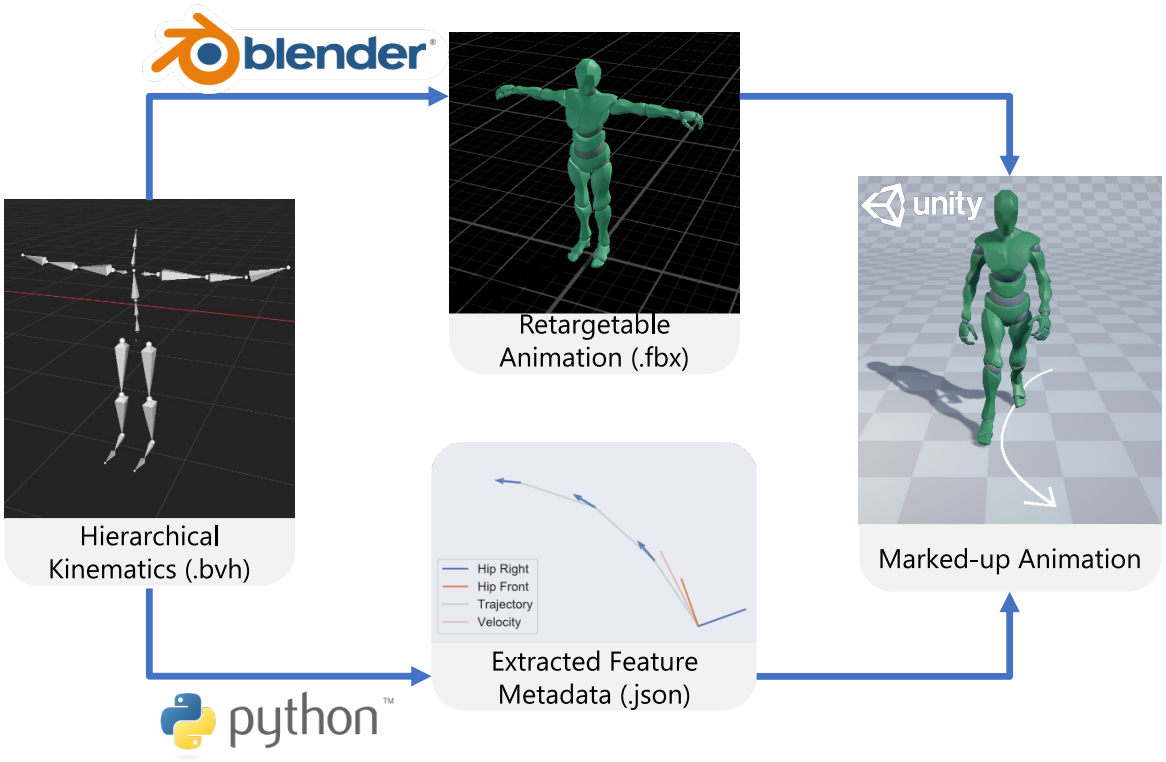


Figure A.1: Flowchart of the data processing steps, the file formats and the software used. The recorded motion capture data is exported as a BioVision Hierarchical file format (BVH). Using Blender, we convert it to the binary FBX format for efficient loading and asset management in Unity. Parallel to this, we process the BVH files with our custom scripts to label every tenth frame with the motion matching features (Hodossy et al. 2024), and export these data in a JSON or HDF5 format. Inside Unity, each FBX animation file is associated with a corresponding metadata file, meaning that the motion matching features do not need to be computed during runtime for the animation database.

A.2 Root animation

A humanoid animation is composed of two parts: firstly, the local joint angles as described in relation to the origin or root of the humanoid (usually the pelvis). Secondly the 6 DoFs motion curves of the root in global space. How the latter is defined is a matter of trade-off between fidelity to the source animation, or following the trajectories determined by the simulated intent.

The matching implemented uses displacement from animation instead of simulation (as described by (Clavet 2016)), but intentional control of trajectory is retained (Figure A.2). Furthermore, instead of constructing a query of the actual current pose, a reverse search from the current time of the animation playing is used to get the closest in time features from the data set directly. Since this means that there are exact matches along several dimensions, this could be a potential source of speedup in the search in a modified K-D tree algorithm. The object-oriented structure of the the matching process implemented is illustrated in Figure A.3.

A.3 Nearest neighbour search

A naive linear search is sufficient for simulating at 60 Hz. However, the computational costs scale linearly with the number of avatar's animated (e.g. in massively parallelised RL environments), the simulation frequency (for more accurate physical simulations) and the size of the data set (to perform more diverse motions). As such, increasing scalability of motion matching is a potential improvement in the future. Example modifications include K-D trees that increase the search speed, Product Quantization (Jegou et al. 2010) that reduce the dimensionality of the search, parallelising the search or using a learning based method to compress the motion matching process with a function estimator (Holden et al. 2020).

K-D tree based nearest neighbour algorithms have been applied with significant speed up on motion matching, as discussed by (Clavet 2016). They mention that this is a somewhat unexpected result as K-D trees do not usually perform well on high-dimensional feature spaces. A potential explanation to this could be that query vectors are expected to be quite similar to

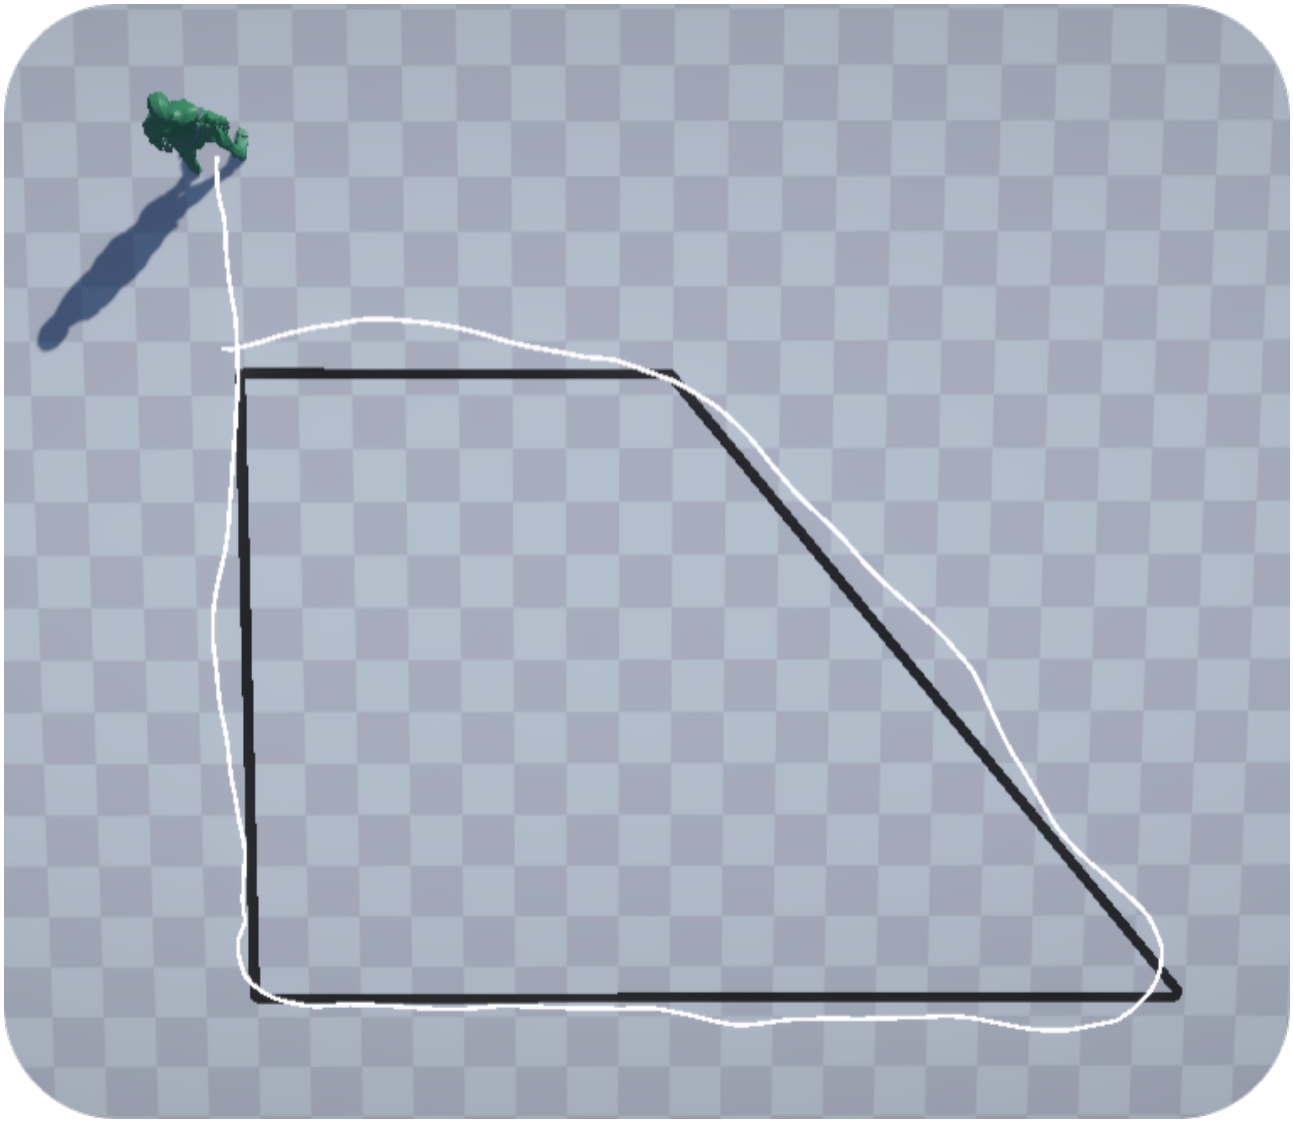


Figure A.2: Top-down view of the resulting walking path (white) when directing the motion matching process by a user attempting to guide the walk along a trapezoid track (black). Despite using displacement from animation, user based inputs with joystick allow accurate following of a pattern on the ground.

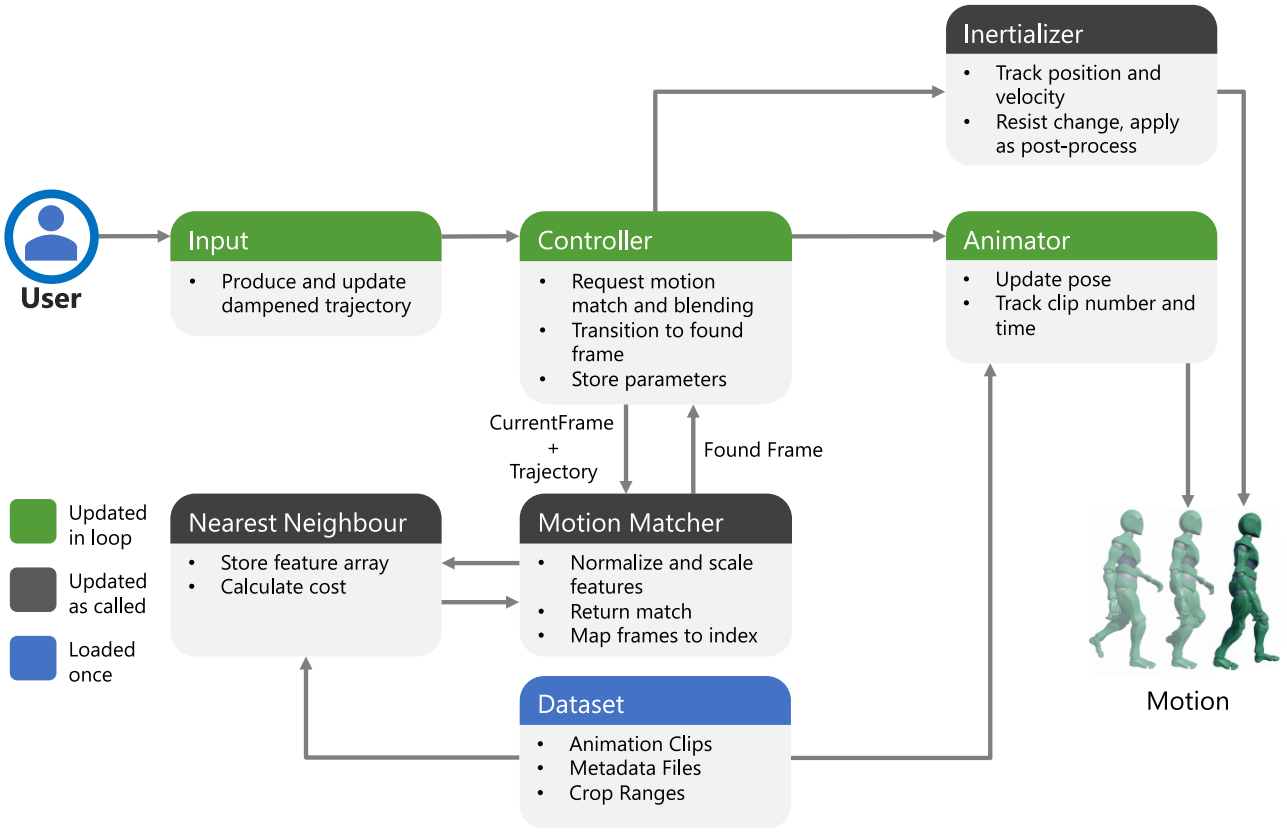


Figure A.3: Summary diagram of subsystem roles and interactions performed every frame of the simulation. The Controller waits 10 frames between constructing a query vectors, formed from the pose features of the currently active animation and the trajectory features from the user input. This is sent to the Motion Matcher object, which has processed the input dataset and paired it with the metadata features. The motion matcher optionally defines ranges of frames valid for matching, and normalises and rescales the features (based on the parameters supplied by the controller). Once the Motion Matcher receives a query vector from the Controller, it calculates a cost for each valid frame in its dataset, and returns the time and clip index of the found nearest neighbour. If the animation frame is sufficiently far away in time (determined another configurable parameter) than the currently playing one, the Controller notifies the Animator of the change, and the Inertialiser to perform the blend.

the ones in the data set. This means that it often is not necessary to check other branches of the tree for potential better matches, which is one main sources of slowdown. In a short test with K-D trees, approximately 6 times less time was needed on average for the nearest neighbour search when using all locomotion files from the LaFAN database (Ubisoft La Forge 2020). However a drawback of this is that once the tree is built it is not straightforward to modify the underlying data or cost function of matching if switching between modes of operation is desired. When using a restricted set of motion capture files of approximately 5 minutes (suitable for the locomotion tasks of Chapters 5 and 6), no speed up was gained in comparison to naive linear search. For this reason K-D tree based search was no used when training locomotion agents.

Figure A.4 illustrates how the database is navigated during motion synthesis for a steady-state straight walking scenario and a non-steady state one, when represented in the 3D PCA space of the motion matching features (with normalised features). The cyclic phase structure is apparent, and that one of the components (e.g., when decomposed with PCA aligns with the turning direction. Each turning direction has states corresponding to the different phases of the gait cycle, allowing smooth tranisitons

A.4 Inertialisation blending

The user input drives the matching process into areas of the data set that are similar in terms of pose, but a degree of discontinuity will be present since continuity is traded for a closer match in terms of future trajectory. (Bollo 2017)'s blending strategy, called inertialisation, is used to interpolate between these poses. This method is applied for motion matching by (Bergamin et al. 2019). Quaternion inertialisation is performed on all segment rotations, and scalar inertialisation is performed on the pelvis height. This blending is applied as a post-process every frame, after the regular animation. Since the upper-body's rotation is not accounted for in the matching cost function, the discontinuity is larger there. However, the exact orientation is also less critical with these segments. Therefore a stronger blending can be applied here to maintain smooth animations. The magnitude of blending can be manually tuned, with a trade-

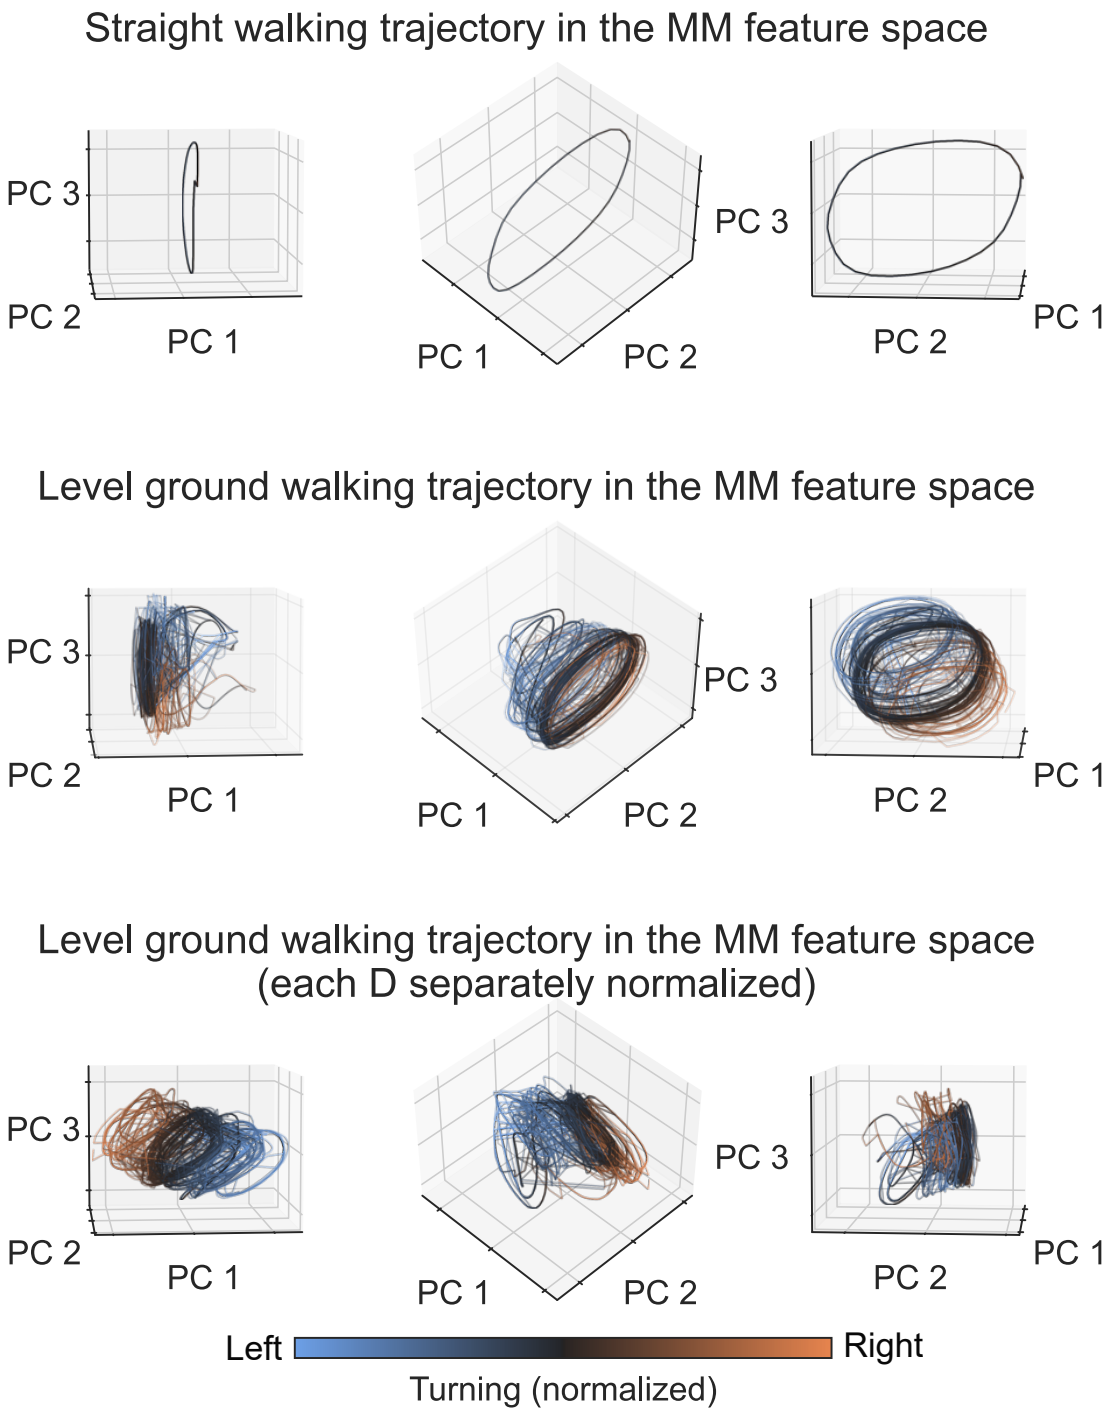


Figure A.4: Three points of view shown for each condition (rows). *Top row:* Feature space trajectory when walking only straight. *Middle row:* Trajectories when randomly generating turns and pivots (see Section 5.2.2 for details). *Bottom row:* Same except when each feature dimension is normalised separately (instead of normalising vectors based on their average norm).

off between reducing discontinuity and kinematic accuracy after a match is made. Common kinematic artefacts caused by inappropriate blending weight are foot-sliding (movement of the stance foot) if too high, or jittery motion if too low.

Appendix B

Pilot results

This pilot trajectory dataset was captured using with a Vicon Vantage (Vicon Motion Systems Ltd, UK 2021) motion capture system recording at 100 Hz, with concurrent Delsys Trigno (Delsys Incorporated, USA 2021) wireless bipolar sEMG at 2 kHz. The electrodes are placed on the lower trunk of the subject. Additionally, the 6 DoF thorax and 3 DoF foot kinematics were recorded, as different gait events and thorax orientation could hold valuable information regarding gait. This concept was not explored in the pilot in the end, but was expanded upon the follow-up study.

This collection of figures illustrates the performance of the TCN model without the dampened trajectory model using the pilot dataset. Note that only models with non-causal windows (trajectory predicted at the start of EMG window instead of its end) were successful with the pilot dataset.

The following properties have generally been observed across different models:

- As expected, the three samples are strongly correlated, but predictions do not simply scale the samples. Distributional differences and temporal patterns are apparent. For example, early samples show stronger periodicity due to the gait cycle (see sinusoidal high frequency components in the 1/3 second sample in Figure B.3).
- Predicting along the pelvis' forward (coronal) axis appears to be more accurate than along

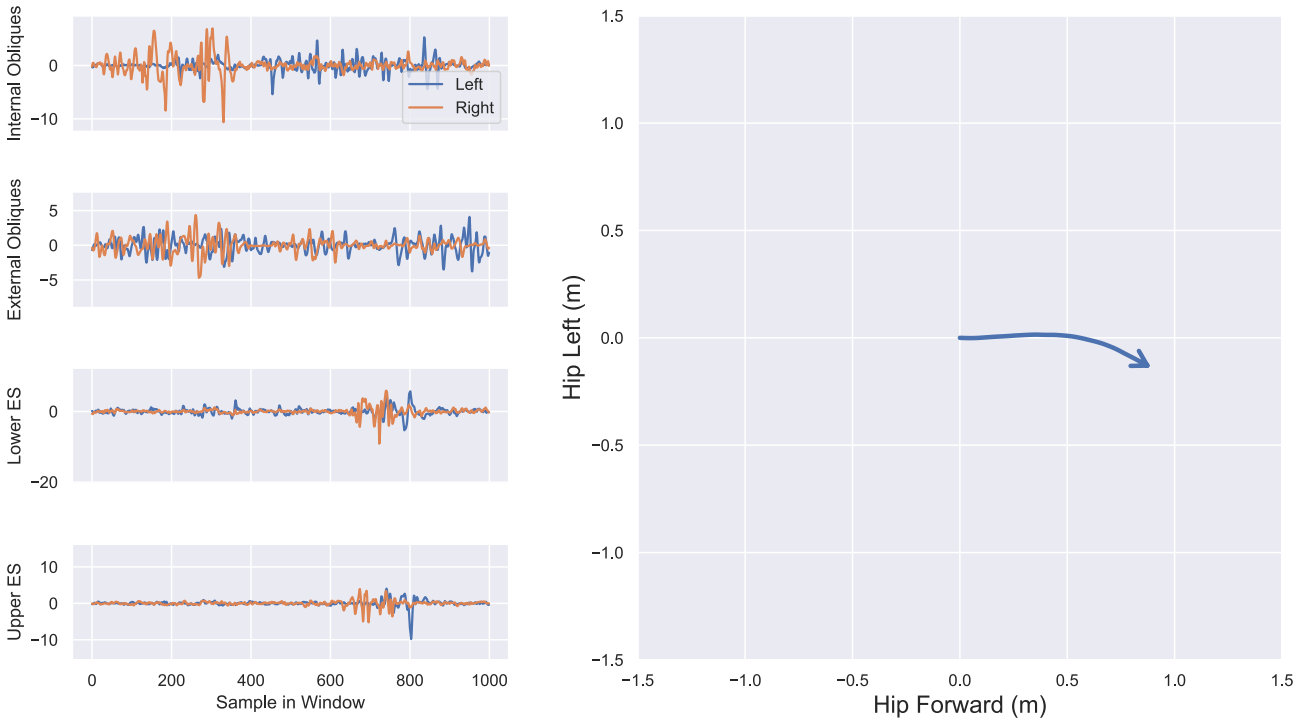


Figure B.1: Ground truth prediction frame to be approximated for intent estimation. The left columns shows the predictors: 0.5 second sized window of 8 bipolar channels recording bilateral trunk muscle groups. ES stands for Erector Spinae. The right column shows the outcome variables: 1 second long horizon of the pelvis' position, relative to the pelvis' position and orientation at 0 seconds. The event shown is an obtuse right turn.

the left/right (sagittal) axis. A contributing factor to this that the forward direction has greater range of motion, especially in early samples, therefore is a higher priority of the network to optimise. A further reason could be that the forward axis values are independent of the direction of turn.

- There are occasional spikes of error, and large variability in regions of uncertainty (e.g. turns, especially 180°). This could be partially mitigated by different smoothing filters applied as a post-process, at the cost of introducing a slight delay, which might necessitate making predictions further ahead in time. This fidelity-responsiveness trade-off is a returning concern in interactive motion control (Geijtenbeek et al. 2012).

The pilot study took place at the Musculoskeletal Biomechanics laboratory of the Michael Uren Biomedical Engineering Research Hub, London. We express our thanks to Professor Anthony Bull for allowing access to the facilities.

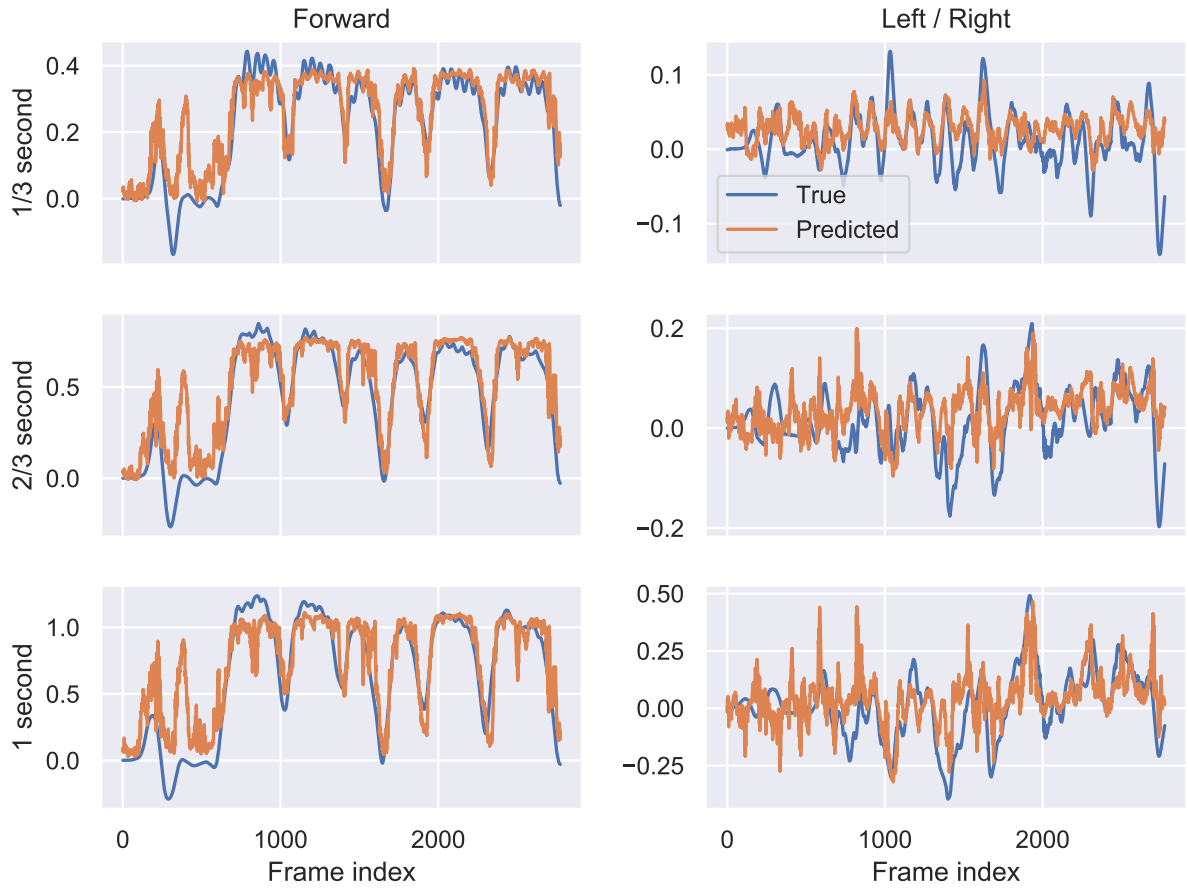


Figure B.2: Model trained and tested on angular turns. Notice at the start the user readjusted their position, walking backwards. This behaviour was not included in the training set, hence the model is unable to predict that very different behaviour.

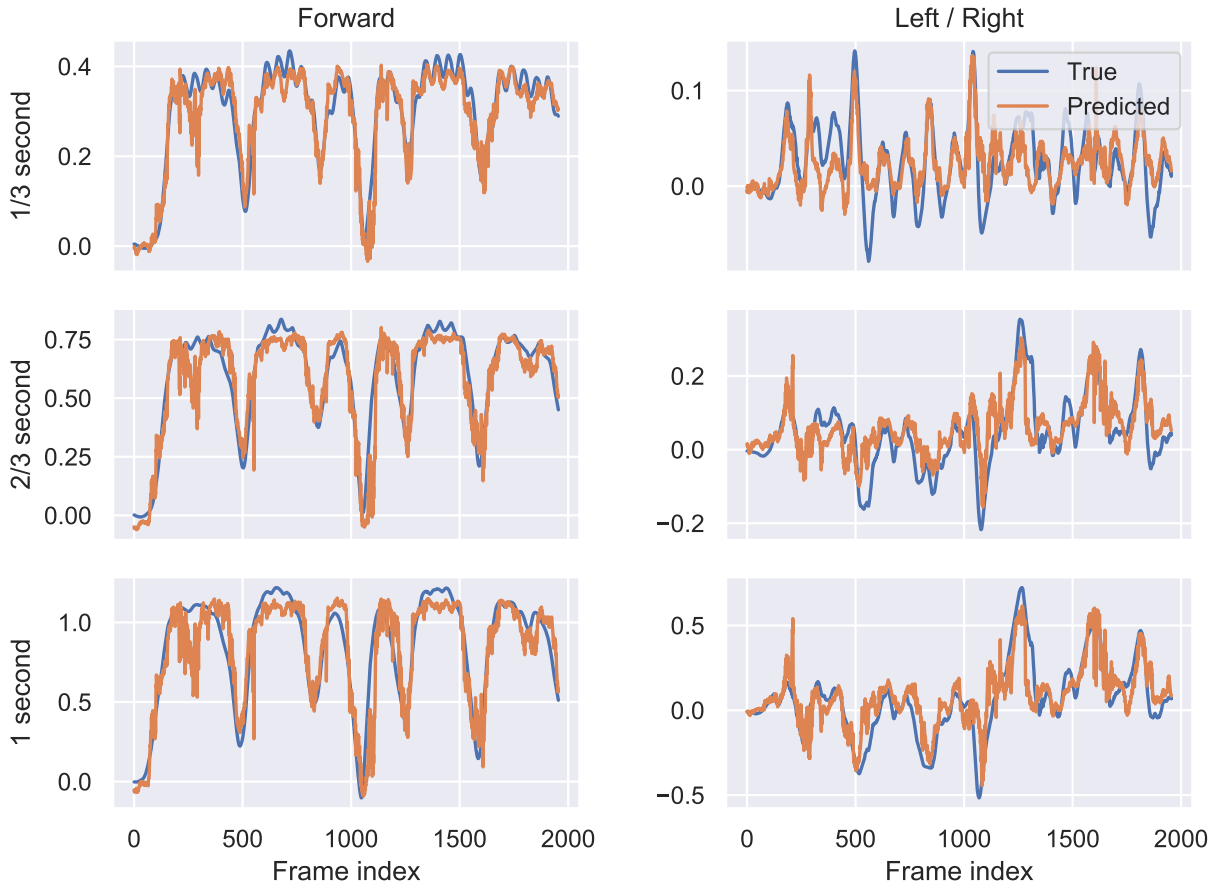


Figure B.3: Model trained and tested on angular turns. The model can not only predict the direction of turn, but also distinguish between the different angles of turning.

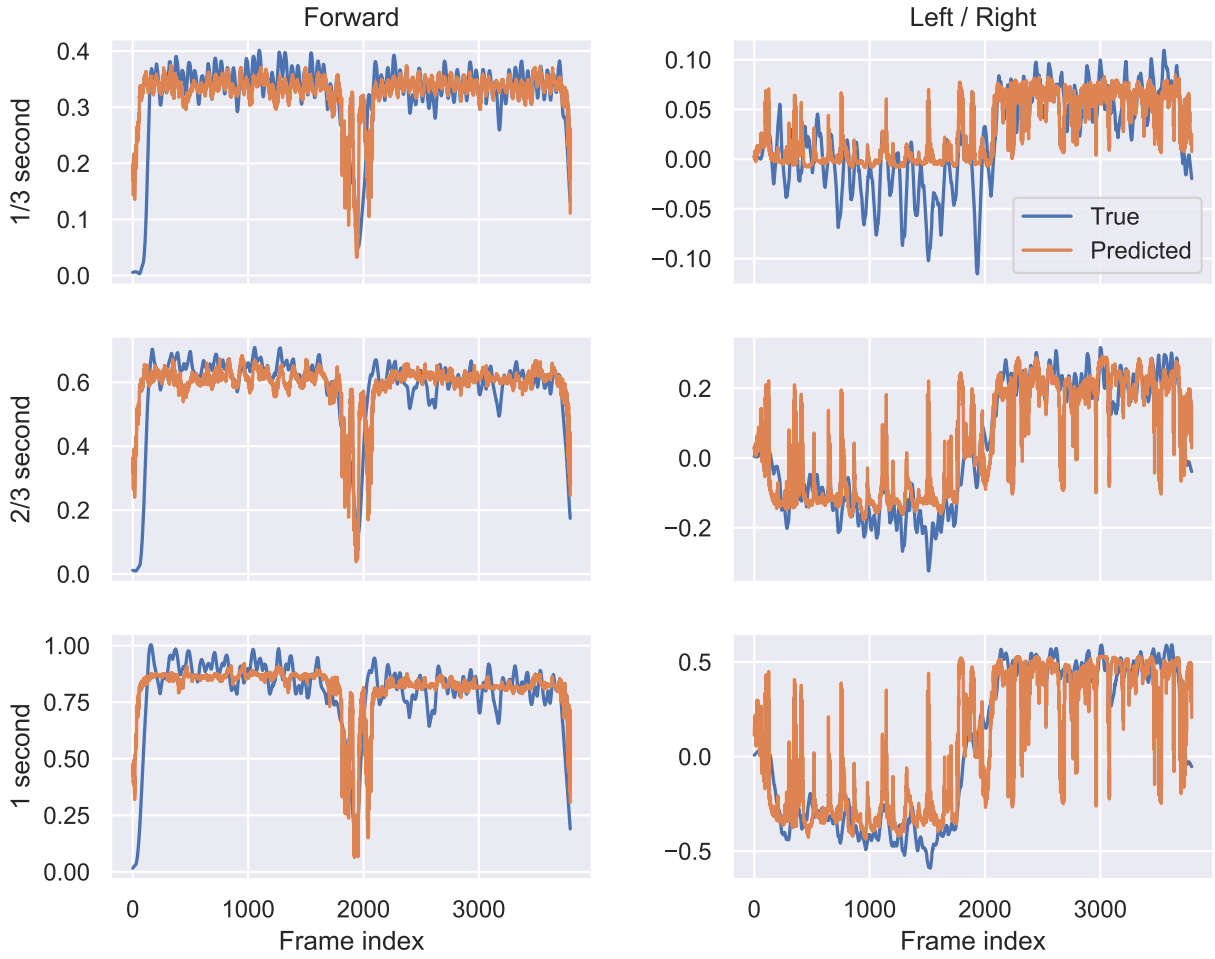


Figure B.4: Model trained and tested on circular walking. All or nothing response is shown, where the model is effectively classifying which direction the turn is going to be and predicting a stereotypical curve, with very little in between. When a mis-classification occurs a short spike in error can be observed along the left/right axis.

Appendix C

Motion Capture and sEMG Recording Protocol

Subject ID: _____ Date: _____

High Level Locomotion Intent

Estimated through motion capture experiments.

Study: "Intent-driven motion synthesis for virtual wearable robotics"
JRCO number: 22IC7765
Name of the Principal Investigator: Professor Dario Farina

I. Experiment Aims

Collect training dataset for neural control model of trajectory-based locomotion synthesizer algorithms.

Outcome variable:

- ~1s pelvis horizontal trajectory horizon

Predictors:

- Multi-electrode trunk and leg muscle bipolar sEMG time-series

6 DoF motion capture is also collected with a whole-body based markerset. To extract additional potential outcome variables (such as foot placement location), or predictors (foot contact status, thorax or head orientation).

II. Equipment

Motion Capture: Vicon Vantage V8 system
sEMG: Delsys Trigno

III. Participant information

Sex (M/F/X)	
Age (years)	
Body mass (kg)	
Height (cm)	
Knee width (cm)	
Ankle width (cm)	
Elbow width (cm)	
Wrist width (cm)	
Hand thickness (cm)	
Shoulder offset (cm)	

IV. Marker Placement

5 markers / foot, 5 markers / shank, 5 markers/thigh, 4 markers for pelvis, 4 markers for thorax, 3 markers for head, 7 markers per arm
49 markers total

Marker set from:

Leboeuf, F., Baker, R., Barré, A., Reay, J., Jones, R. and Sangeux, M., 2019. The conventional gait model, an open-source implementation that reproduces the past but prepares for the future. Gait & posture, 69, pp.235-241.

V. Electrode Placement

Electrode Location	Palpation and Orientation Guide	Checklist
Erector Spinae - lower (Left and Right)	Over palpable bulge of muscle (approx. 3cm lateral to midline) lower electrode at L1 level, vertical	2 electrodes
Abdominal External Obliques (Left and Right)	Upper electrode directly below most inferior point of costal margin, on line along pubic tubercle.	2 electrodes
Peroneus Longus (Left and Right)	25% on the line between the tip of the head of the fibula to the tip of the lateral malleolus.	2 electrodes
Soleus (Left and Right)	2/3 of the line between the medial condylis of the femur to the medial malleolus.	2 electrodes
		Total: 8 electrodes

Adapted from:

Anders, C., Brose, G., Hofmann, G.O. and Scholle, H.C., 2008. Evaluation of the EMG–force relationship of trunk muscles during whole body tilt. Journal of biomechanics, 41(2), pp.333-339.

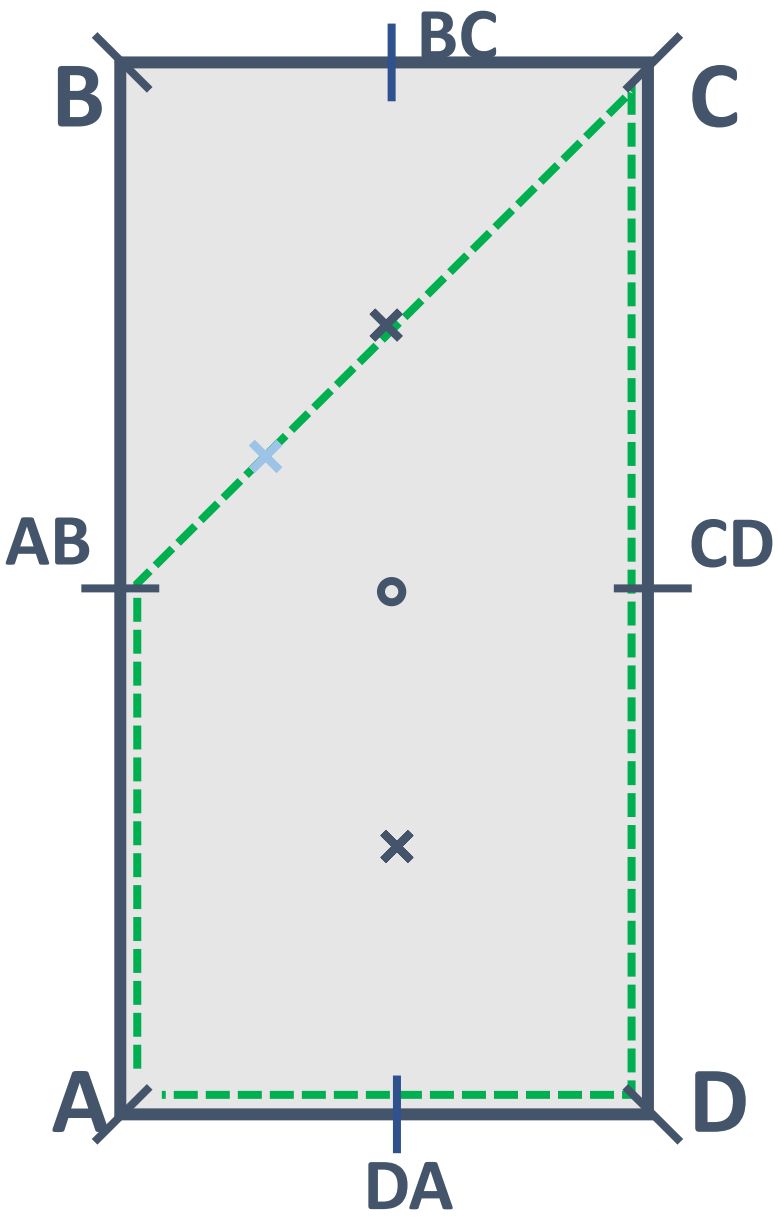
Hase, K. and Stein, R.B., 1999. Turning strategies during human walking. Journal of neurophysiology, 81(6), pp.2914-2922.

Stegeman, D. and Hermens, H., 2007. Standards for surface electromyography: The European project Surface EMG for non-invasive assessment of muscles (SENIAM). Enschede: Roessingh Research and Development, 10, pp.8-12.

VI. Calibration trials

Static and dynamic RoM trials will be captured with the full markers + electrodes

VII. Motion Trials



Ground markers for trials

Adapted from Zadziuk (2016)

First six trials are approximately 8 minutes long. Potentially repeatable, starting with different leg side (i.e. always start with right, perform repetitions, switch to starting with left, perform repetitions again).

Repetition counts include the first instance as well.

Trial 1: Angular turns

Recording start

From vertex B, follow green path at steady pace till It ends, pivot around and retrace path.

Recording stops

15 repetitions

Trial 2: Angular stops

Recording start

Same as trial 1. At each vertex come to a full stop, maintaining direction faced.

Recording stops

5 repetitions

Trial 3: Strides

Recording start

From vertex B, walk to vertex D without stopping, pivot and return to vertex A, where come to a rest.

Recording stops

5 repetitions

Trial 4: Arcs

Recording start

Starting in A walk to AB, then arc in tangent to the dark blue guide walk to CD. Carry on to D, where pivot and retrace till A is reached again. Pivot repeat for the sharper arc until DA.

Recording stops

2 repetitions

Trial 5: Slalom

Recording start

Start at DA. Using the guides in the DA-BC axis, slalom to BC. Approach first marker from left. Pivot at BC and return, approaching first marker from right

Recording stops

1 repetition

Trial 6: Circles

Recording start

Start at AB. Walk clockwise around circular guide. Cross-shaped guides tangential. After three cycles pivot at AB and do 3 counterclockwise cycles.

Recording stops

1 repetition

Trial 7: Freeform

Recording start

3 minutes of freeform walking inside capture volume.

Recording stops

Appendix D

Adversarial Imitation Learning for Policy Transfer

Motion tracking RL uses an explicit, time synchronised kinematic animation source, such as the one synthesized from motion matching. This allows training locomotion agents at a significantly faster rate, by providing direct and relevant description about the optimal solution to the learning framework (Peng et al. 2021). However, this requires careful definition and tuning of the reward function used to incentivise tracking the reference, and in some cases it is challenging to successfully learn skills beyond replicating the motion. Therefore, it would be desirable to be able to just define a style, as determined by a set of reference motions, and allow the agent to pursue its goals while adapting the provided style. This can be achieved by simultaneously training an adversarial discriminator, who is also trained by RL (Merel et al. 2017). The discriminator receives samples from either the reference or the agent's rollout; it is rewarded if it successfully labels the given steps as having come from the reference or the simulation, and the agent is rewarded on the inverse. This way no explicit loss function needs to be designed and tuned by the researcher (see Figure D.1). The locomotion agent is only influenced by the reference motion through the reward received from the discriminator, and is free to pursue other goals in the environments.

We reimplemented the Adversarial Motion Priors method from (Peng et al. 2021), and generated

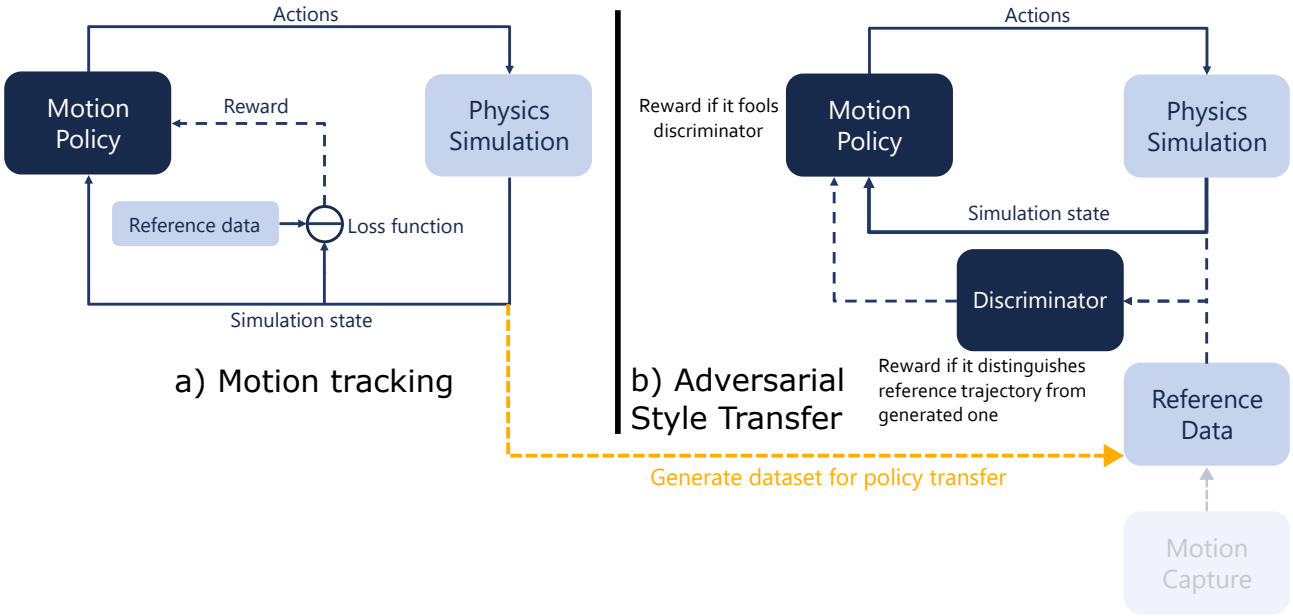


Figure D.1: Differences in the structure of motion tracking vs Adversarial Style Transfer methods. Traditionally motion capture data is used for reference in adversarial learning, here we instead use the movement generated by a motion tracking agent.

a virtual example dataset using the motion tracking methods describe in Chapter 5 for the case of steady-state straight walking. Instead of using a real motion captured dataset we used these examples to transfer the policy learned with the faster motion tracking method, to one that does not use a synchronised reference motion in its actuation as a baseline, or its reward function. Then the agent was trained using Adversarial Motion Priors to replicate the style of the motion tracking agent, while additionally trying to match an arbitrary walking direction determined at the start of the episode. This transfer is faster than trying to learn a control policy by copying the distribution from motion capture, as the motion copied from reference agent is guaranteed to be achievable and stable for the new agent (see Figure D.2). Like in (Peng et al. 2021), we also include the previous state in the observations provided for the discriminator, as the direction of how the motion is supposed to evolve guides the emergence of motion. Without this, we observe that the agent reproduces various poses from the reference, with seemingly randoml transitions between them.

The reward function for following the direction is calculated using the following method:

$$r_{direction} = \exp\left(k \cdot \left(\frac{1 - \vec{d}_{\text{des}} \cdot \vec{d}_{\text{cur}}}{2}\right)\right)$$

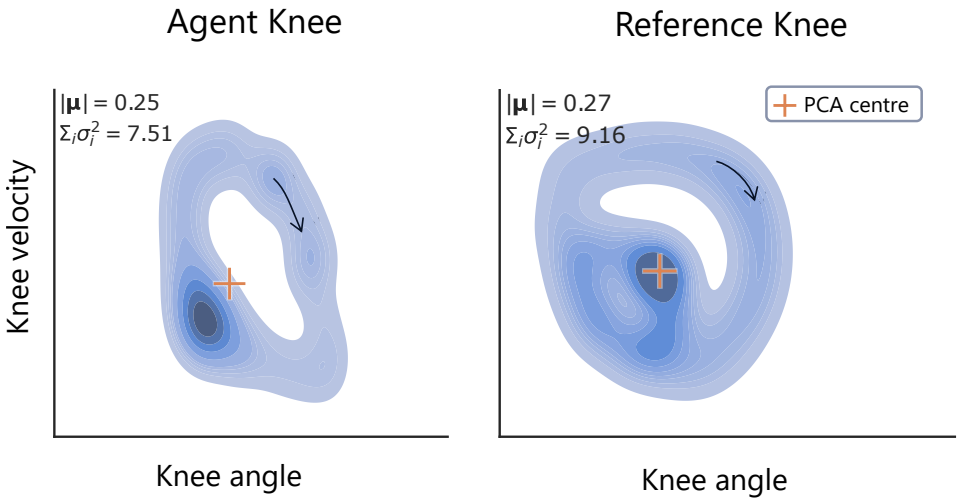


Figure D.2: Two reference phase space distributions of the left knee. Motion from another agent trained with motion tracking is shown on the left, and one from motion capture shown on the right. Small differences arising from details such as the joint centre location can bias the shape of the phase diagram distribution making it more challenging to replicate using Adversarial Motion Priors. We also illustrate the direction of the state evolution with arrows overlayed, representing the state transitions included for the discriminator. The goal of the discriminator is to guess that a given state and transition belongs to the reference distribution or the one of the trained agent.

Here $\vec{d}_{\text{des}}$ is the normalised horizontal walking velocity designated randomly, and $\vec{d}_{\text{cur}}$ is the current normalised horizontal walking velocity as measured at the pelvis. The parameter k governs the fall off rate of the reward, a value of -8 was used.

This way novel turning strategies could be learned by the new agent in less than 3 million learning steps, generating turning on command despite it not appearing in the reference dataset (example snapshots of turning in Figure D.3. We successfully demonstrate the Adversarial Motion Priors can be used for learning transfer, and indicate the use of adversarial imitation as a methods of interest for training locomotion agents in the future for modelling wearable robotics use.

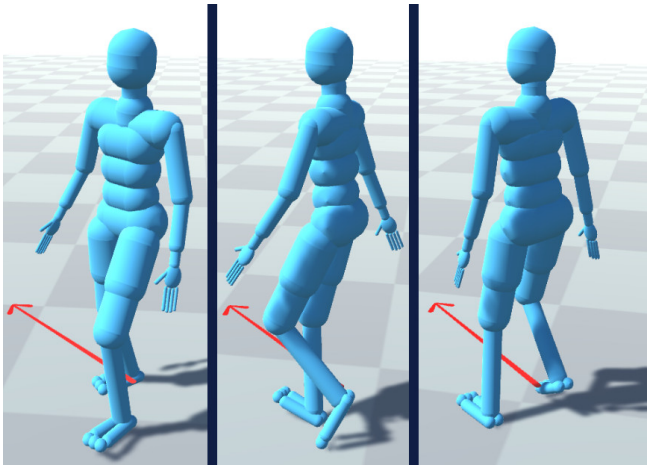


Figure D.3: When the target horizontal walking velocity changes, the agent quickly reorients the walking direction, despite no such behaviour being available in example data.

Appendix E

Motion Matching for Feedforward Control of Prosthetic Legs

Motion matching (Clavet 2016; Bergamin et al. 2019; Holden et al. 2020) is a method that is able to generate diverse full body kinematics that is directly and responsive to user input, while maintaining fidelity of the source motion capture. This can be exploited to generate motion to a part of the body if the kinematics of the rest can be measured. The generated kinematics can be used as target positions for wearable assistive robotics, to be tracked with a method such as stable PD (Tan et al. 2011).

We explored this approach in a proof of concept experiment. Two separate level ground walking motion capture datasets were used:

- *Dataset A*, constructed from (Hodossy et al. 2024).
- *Dataset B*, constructed from the non stylised locomotion trials of (Ubisoft La Forge 2020).

These used different subjects and motion patterns. Humanoid locomotion agents were trained using the methods from (Hodossy et al. 2023a), trained with the recordings of Dataset A. The control of the locomotion agent was then suspended over at three levels of the right leg for three different conditions:

1. Transtibial level, where only the ankle control was blocked.

2. Transfemoral level, where additionally control over the knee was prevented.

3. Hip disarticulation level, where hip, knee and ankle were completely blocked.

Naturally, without control, holding these in a static pose resulted in the pretrained agent to be unable to complete any gait cycles. If allowed further training and learning compensatory movement patterns, at the transtibial level the locomotion agent was able to recover ambulation, but failed to adapt to a meaningful locomotion at more proximal levels.

A parallel full body kinematic motion synthesis was then performed using Dataset B, separate from the one that is used for the control of the agent. We will refer to this synthesizer as "Motion Matcher B". Motion Matcher B generated kinematic full body motion at the same time as the dynamic agent, but it did not use its own matching features (see tables Table A.1 and A.2) as usual.

Instead, the pose features of the Dataset A corresponding to controlled bodies (left foot position and velocity, and pelvis velocity) were combined with the right foot's features from the current frame in Dataset B. This formed an altered query vector that was then used to drive the animation of Motion matcher B (hypothesised to be plausible to measure in real life through wearable systems).

In summary, kinematics that were similar in a nearest neighbour sense to the current motion of the intact side and pelvis were found in Dataset B. From these the affected side's motion could be extracted, and used as targets to stable PD actuation of the dynamic humanoid on the previously blocked DoFs, simulating a multi-jointed powered prosthesis. For a visual summary, see Figure E.1.

In a straight walking setting, even without adaptation from the locomotion agent, ambulatory skills were restored for the transtibial and transfemoral level. After permitting adaptation from the locomotion agent, even at the hip disarticulation level the agent was successful in walking. It is noted that the stability at the hip disarticulation level was heavily improved if

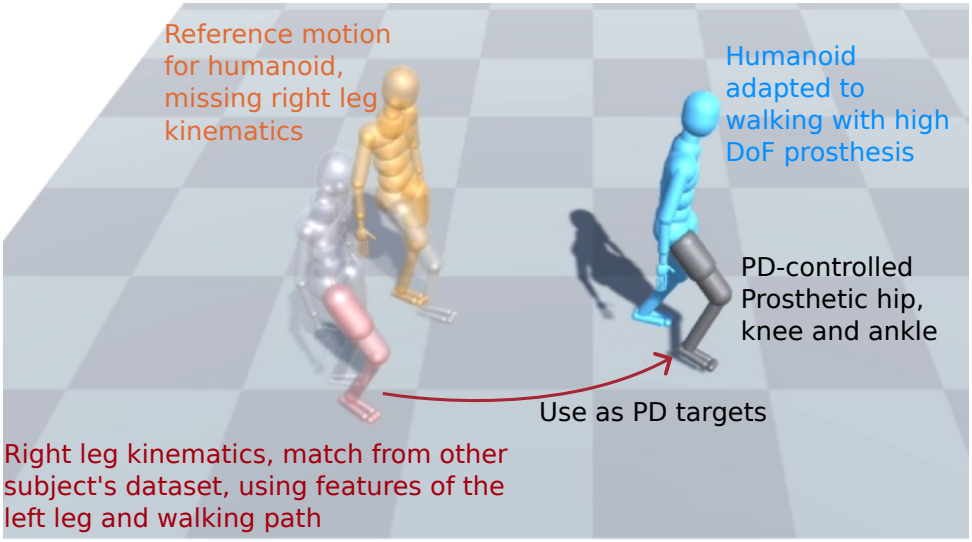


Figure E.1: Walking with high DoF prosthesis, with the missing kinematic targets recovered from a second subject's dataset using motion matching. The yellow character is the motion from Dataset A (missing the right leg), the red character is Dataset B (only right leg used) and the blue character is the physically simulated RL locomotion agent.

the trajectory features of the dynamic agent were used in addition to its pose features. These features could only be obtained if a high level intent estimator like the one in Chapter 4.

These promising results indicate that data-driven, non-deep learning methods could provide target motion trajectories in non-steady-state settings for high DoFs prosthetics.

Appendix F

Modelling Sensory Deficit in Prosthesis Use

In the studies of prosthesis use we presented so far, we have not yet modified the proprioceptive sensing of the humanoid locomotion agent. However, deafferentation on the affected side is an evident consequence of amputation. It is therefore expected to influence the kinaesthetics that would normally be guided through the sensory structures of the removed parts of the musculoskeletal system. Over time, a prosthesis user may obtain embodiment and restore proprioception through experience and well configured suspension systems but their stability and locomotion skills will be reduced until then (Alimusaj et al. 2009). This is especially true in environments where the chance of tripping is increased, such as where stair climbing or stepping up is required. Compensatory gait patterns, such as hip hiking, vaulting and amplified knee flexion are often adopted by lower limb prosthesis users on the affected side to increase the toe clearance, mitigating the risk of tripping and falls.

Artificial sensory feedback systems aim to collect relevant information about the body pose and its surroundings. They encode the relevant observations and relay it to their user through invasive (e.g., nerve stimulators) or non-invasive (e.g., vibrotactile) methods (Petrini et al. 2019). This can lead to benefits ranging from increased functional mobility, improved perception of the device and embodiment, and stability. This is especially relevant for transfemoral or more

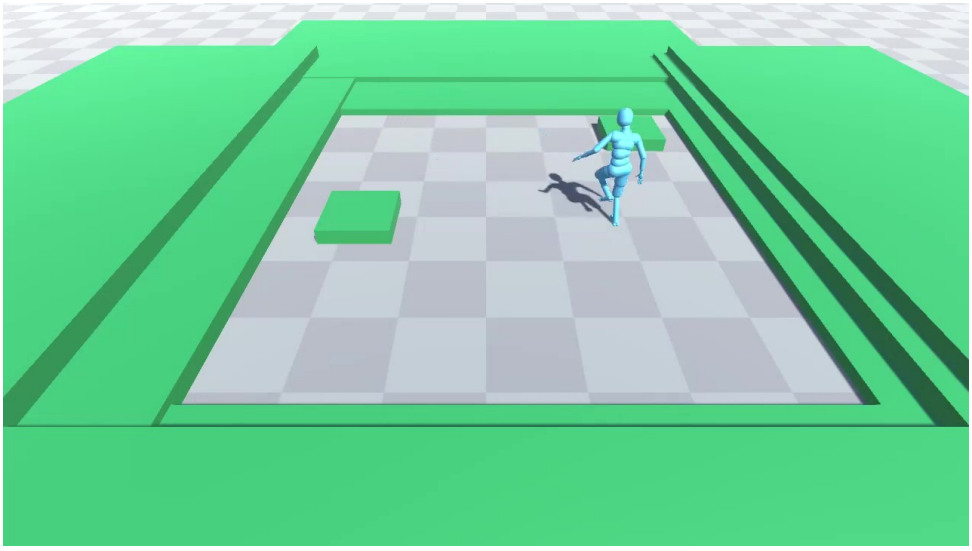


Figure F.1: Added obstacles that require learning stepping skills. Agent is performing knee hyperflexion even when no stepping up is imminent, due to the uncertainty in the environment and its pose sensing.

proximal amputees.

In this proof-of-concept study we investigate how unilateral blocking the proprioceptive signals of a RL trained locomotion agent affects the learned gait patterns in a stepping environment. We use the methods from (Hodossy et al. 2023a) to train a locomotion agent in a stepping environment shown on Figure F.1. The observation vector is augmented by the forecasting observations described in Section 6.2.3 for only the left side of the humanoid. We also model a sensory deficit on the left side. The observations regarding the position of the left foot are removed and replaced with the position of the right knee to represent deafferentiation in a transfemoral amputation case. The actuation and joint properties of the left lower limb is not altered; a single policy is controlling the whole body. This essentially assumes an ideal suspension and neural interface, as the goal of the study was to explore the effects of reduced sensory information, not the human-prosthesis communication or the prosthesis control challenge.

We compare the sensory deficit agents with agents that have the same proprioception and forecasting stepping observations on both sides. When initially trained on level ground environments, we observe no significant compensatory movement patterns emerge in the sensory deficit condition. However, once transferred to the stepping environment, the sensory deficit agent develops significantly increased toe and foot clearance on the affected side (Figure F.2),

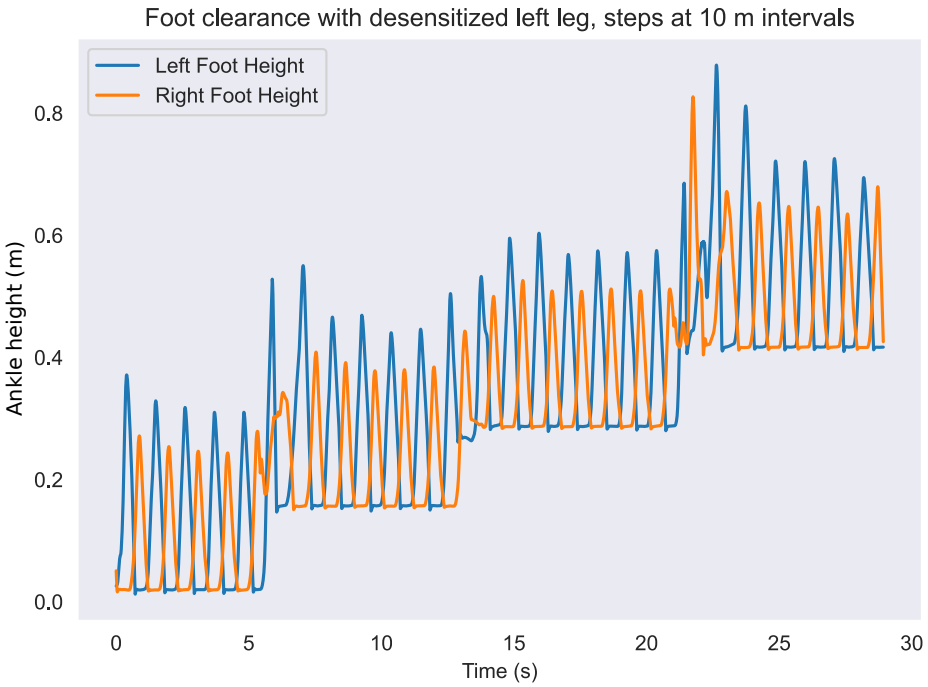


Figure F.2: Foot height in a straight walking scenario with the sensory deficit, with three non-consecutive steps. The agent selects a step height on the left side to be able to comfortably clear a potential step.

achieved through a mixture of vaulting and knee hyperflexion, when compared to the condition with symmetric sensory input. The average peak ground clearance of the ankle (across gait cycles) on the left side was 42.8 ± 7.4 cm, and the right side was 34.9 ± 6.5 cm.

These results are indicative that virtual locomotion agents may be an avenue to investigate what key information about the human pose and its environment is worth capturing and relaying to the user. For example, identifying lower dimensional representations of the trip avoidance task could mean less communication channels are needed in sensory feedback for neural bionic limbs, which imply a lower cognitive load for their users. Similarly to the control problems discussed throughout this thesis, novel sensing systems could be proposed in simulation; we, for example, show that simple 2-dimensional exteroceptive signal is able to reduce the foot and toe clearance developed in locomotion agents. The same approach is highly relevant in upper limb devices as well (Galan et al. 2015), where it has already seen some application (Kluger et al. 2019). However, in investigations where the emergence and details of compensatory gait patterns is key, musculoskeletal humanoid models may be of higher priority. Factors like range of motion

and energy usage, which musculoskeletal models capture better, may play an important role in this kind of investigation in the future.

Acknowledgement

Thank you to Shloak Mehta for his insight and discussions on foot clearance perception and gait stability, which inspired the work in this appendix.

Appendix G

RL configuration file for locomotion policy

```yaml
# These hyperparameters were used with the ML-Agents framework
# for learning device and gait policies. The configuration from
# the following source was adapted:
#   J. Booth and V. Ivanov, "Realistic physics based character controller,"
#   arXiv preprint arXiv:2006.07508, 2020.
#
# For additional information on the meaning of the configuration parameters
# the ML-Agents documentation is available at the following address:
# https://unity-technologies.github.io/ml-agents/Training-Configuration-File/
default_settings: null
behaviors:
  DReCon-Mj:
    trainer_type: ppo
    hyperparameters:
      batch_size: 8192
      buffer_size: 81920
      learning_rate: 2.0e-05
      beta: 0.0009
      epsilon: 0.2
      lambd: 0.95
      num_epoch: 3
      learning_rate_schedule: linear
      beta_schedule: linear
      epsilon_schedule: linear
    network_settings:
      normalize: true
      hidden_units: 512
      num_layers: 3
      vis_encode_type: simple
      memory: null
      goal_conditioning_type: hyper
      deterministic: false
    reward_signals:
      extrinsic:
        gamma: 0.99
        strength: 1.0
        network_settings:
          normalize: false
          hidden_units: 128
          num_layers: 2
          vis_encode_type: simple
          memory: null
          goal_conditioning_type: hyper
          deterministic: false
    init_path: null
    keep_checkpoints: 5
    checkpoint_interval: 5000000
    max_steps: 8000000
```

```
    time_horizon: 1000
    summary_freq: 50000
    threaded: false
    self_play: null
    behavioral_cloning: null
  Prosthesis:
    trainer_type: ppo
    hyperparameters:
      batch_size: 8192
      buffer_size: 81920
      learning_rate: 0.0001
      beta: 0.0005
      epsilon: 0.2
      lambd: 0.95
      num_epoch: 3
      learning_rate_schedule: linear
      beta_schedule: linear
      epsilon_schedule: linear
    network_settings:
      normalize: true
      hidden_units: 128
      num_layers: 3
      vis_encode_type: simple
      memory: null
      goal_conditioning_type: hyper
      deterministic: false
    reward_signals:
      extrinsic:
        gamma: 0.99
        strength: 1.0
        network_settings:
          normalize: false
          hidden_units: 128
          num_layers: 2
          vis_encode_type: simple
          memory: null
          goal_conditioning_type: hyper
          deterministic: false
    init_path: null
    keep_checkpoints: 5
    checkpoint_interval: 5000000
    max_steps: 8000000
    time_horizon: 1000
    summary_freq: 50000
    threaded: false
    self_play: null
    behavioral_cloning: null
env_settings:
  env_path: envs/GaitKeeperUnity.exe
  env_args: null
  base_port: 5005
  num_envs: 6
  num_areas: 1
  seed: -1
  max_lifetime_restarts: 10
  restarts_rate_limit_n: 1
  restarts_rate_limit_period_s: 60
```