





João André

**Combining Dynamical Systems with
Machine Learning for Adaptive Control of
Legged Robots**

Universidade do Minho
Escola de Engenharia





Universidade do Minho

Escola de Engenharia

João André

**Combining Dynamical Systems with
Machine Learning for Adaptive Control of
Legged Robots**

Tese de Doutoramento

Programa Doutoral em Engenharia Biomédica

Escola de Engenharia

Trabalho efetuado sob a orientação da

Professora Doutora Cristina Santos

e do

Professor Doutor Luís Paulo Reis

Direitos de Autor e Condições de Utilização do Trabalho por Terceiros

Este é um trabalho académico que pode ser utilizado por terceiros desde que respeitadas as regras e boas práticas internacionalmente aceites, no que concerne aos direitos de autor e direitos conexos.

Assim, o presente trabalho pode ser utilizado nos termos previstos na licença abaixo indicada.

Caso o utilizador necessite de permissão para poder fazer um uso do trabalho em condições não previstas no licenciamento indicado, deverá contactar o autor, através do RepositóriUM da Universidade do Minho.

Licença concedida aos utilizadores deste trabalho



Atribuição

CC BY

<https://creativecommons.org/licenses/by/4.0/>

Acknowledgements

The work described in this thesis was funded by the Foundation for Science and Technology (FCT) through the Doctoral Grant SFRH/BD/107891/2015.

Statement of Integrity

I hereby declare having conducted this academic work with integrity. I confirm that I have not used plagiarism or any form of undue use of information or falsification of results along the process leading to its elaboration.

I further declare that I have fully acknowledged the Code of Ethical Conduct of the University of Minho.

Resumo

Controlo Adaptativo de Robôs com Pernas através da combinação de Sistemas Dinâmicos com Machine Learning

Nesta tese de doutoramento é investigado o papel de estruturas primitivas em diferentes níveis de geração e controlo motor em sistemas naturais e artificiais. Através de diferentes estudos experimentais, são propostas soluções práticas para diferentes problemas de engenharia, numa abordagem multidisciplinar que explora elementos de baixa dimensionalidade de natureza cinemática, perceptual e neuromuscular. Especificamente, o 1) controlo genérico de locomoção bípede artificial é conseguido empregando apenas uma quantidade mínima de Primitivas Motoras (MPs) partilhadas em dois robôs humanóides, usando sistemas dinâmicos na representação de MPs. a 2) deteção de marcha irregular e atípica em humanos é efetuada através da construção de memórias perceptuais genéricas representando locomoção bípede normal, com diferentes Primitivas Perceptuais (MPs) partilhadas por diferentes dimensões, e levando a uma representação compacta, robusta e invariante no tempo. e a 3) identificação de fadiga muscular em movimentos repetitivos dos membros superiores, com base em variações graduais e significantivas no recrutamento de Sinergias Musculares (MSs), quantificadas através de múltiplos indicadores no domínio temporal e espectral. A validação experimental dos modelos e algoritmos introduzidos nesta tese demonstra a relevância de estruturas primitivas na construção de estratégias cognitivas artificiais, em particular dentro do paradigma de Aprendizagem Contínua ou *Lifelong Learning* (L2). Na prática, os métodos e modelos propostos são particularmente apelativos na i) construção de máquinas autónomas capazes de lidar eficientemente com tarefas e ambientes complexos e com alta dimensionalidade, podem complementar, ii) podem complementar arquitecturas cognitivas e ferramentas de raciocínio artificial existentes, e iii) contribuir no desenho de soluções de monitorização e gestão de comportamento humano, possibilitando uma mais rica coexistência entre robôs e humanos.

Palavras-chave: controlo robótico, locomoção robótica, primitivas motoras, primitivas perceptuais, sinergias musculares

Abstract

Combining Dynamical Systems with Machine Learning for Adaptive Control of Legged Robots

In this doctoral thesis, the role of primitive structures across different levels of natural and artificial motor generation and control is investigated. Through distinct experimental studies, primitive-based solutions are developed to concrete engineering problems, in a multidisciplinary approach that exploits low-dimensional primitive elements at kinematic, perceptual and neuromuscular levels. Namely, 1) generic control of artificial legged walking is achieved by employing a minimal set of Motor Primitives (MPs) shared in two humanoid platforms, using stable dynamical systems for primitive representation; the 2) detection of irregular, non-stereotypical gait in humans is conducted through the construction of generic perceptual memories of nominal legged gait, sharing Perceptual Primitives (PPs) across different input dimensions, and achieving a compact, robust and time-invariant representation; and the 3) identification of muscular fatigue in repetitive upper-limb human movement, by focusing on meaningful gradual alterations in the recruitment of different Muscle Synergies (MSs), assessed through multiple time- and frequency-domain indicators. The experimental validation of the models and algorithms introduced in this thesis attest to the relevance of primitive theory in the development of efficient artificial cognitive strategies, particularly within the paradigm of autonomous Lifelong Learning (L2). In practice, the methods proposed in this work denote particular appeal in i) the construction of autonomous machines able to efficiently deal with high-dimensional tasks and environments, ii) can complement existing cutting-edge cognitive architectures and artificial reasoning frameworks, as well as iii) aid in the design of advanced human behavioral monitoring and management solutions, enabling a richer coexistence between humans and robots.

Keywords: motion primitives, muscle synergies, perceptual primitives, robotic control, robotic locomotion

Contents

Resumo	iv
Abstract	v
List of Acronyms	x
List of Figures	xv
List of Tables	xxi
1 Introduction	1
1.1 Motivation and Objectives	6
1.2 Contributions	8
1.3 Outline	9
2 Generic Primitive-based Legged Control	11
2.1 Introduction	11
2.2 Background	13
2.2.1 A Biological Perspective on Walking	13
2.2.2 From Static to Dynamic walking	14
2.2.3 The Role of Passive Body Dynamics	16
2.2.4 Bioinspired Robotic Control	18
2.2.5 Related Work	20
2.3 Generic Primitive-based Control	24
2.3.1 Dynamic Movement Primitives	24
2.3.2 Shared Information and Primitive Combination	28
2.3.3 Motor Skill Learning and Generalization	30
2.4 Biped Locomotion with Shared Primitives	35
2.4.1 Central Pattern Generators in Closed-loop Gait Control	36
2.4.2 Generic Legged Control	39

2.4.3	Cloning Walking Behavior	41
2.5	Results	43
2.5.1	Motor Skill Learning	43
2.5.2	Dynamic Walking Performance	47
2.6	Discussion	54
2.7	Conclusion	58
3	Dynamical Perceptual Memory Model	60
3.1	Introduction	60
3.2	Background	63
3.2.1	Paradigms of Human Memory	63
3.2.2	Skill Memories and Perceptual Primitives	64
3.2.3	Perception in Artificial Cognition	65
3.2.4	Related Work	68
3.3	Perceptual Memory Model	70
3.3.1	Memory Matching	74
3.3.2	Context Identification	75
3.4	Validation in Human Locomotion	79
3.4.1	Slip-Related Disturbances	79
3.4.2	Automated Slip-Avoidance Strategies	80
3.4.3	Walking Trials	81
3.4.3.1	Participants and Equipment	81
3.4.3.2	Experimental Protocol	82
3.4.3.3	Data Processing Pipeline	83
3.4.4	Memory Consolidation and Practical Validation	84
3.5	Results	87
3.5.1	Kinematic References	87
3.5.2	Memory Consolidation	90
3.5.3	Gait Classification	91
3.5.4	Detection of Non-Stereotypical Gait	91
3.6	Discussion	98
3.7	Conclusion	102
4	Impact of Muscle Fatigue in Musculoskeletal Primitives	105
4.1	Introduction	105
4.2	Background	107
4.2.1	Electromyography	107
4.2.2	Physiological and Biomechanical Effects of Muscular Fatigue	108
4.2.3	Fatigue Indicators in Classical EMG Analysis	109
4.2.4	Muscle Synergies as Neuromuscular Primitives	110

4.3	Methodology	112
4.3.1	Participants	112
4.3.2	Instrumentation	112
4.3.3	Experimental Protocol	114
4.3.4	Data Processing	116
4.3.4.1	EMG Envelope Extraction	116
4.3.4.2	Kinematic Segmentation	116
4.3.4.3	Normalization	118
4.3.4.4	Muscle Synergy Estimation	118
4.3.4.5	Classical and Synergy Analysis	120
4.3.4.6	Statistical Significance Analysis	120
4.4	Results	121
4.4.1	Self-rated Exertion	121
4.4.2	Kinematic and Electromyographic Uniformity	121
4.4.3	General Upper Limb Synergies	122
4.4.4	Classical Analysis	125
4.4.5	Synergy Analysis	130
4.5	Discussion	132
4.6	Conclusion	136
5	Conclusion	138
5.1	Summary	138
5.2	Outcomes and Contributions	142
5.2.1	Publications	143
5.3	Limitations and Future Work	143
	Appendix A Practical Implementation	145
A.1	Core Libraries	145
A.1.1	mplibrary	145
A.1.2	mplearn	149
A.1.3	mpmemory	149
A.2	ROS Environment	150
A.2.1	roscontrol_utils	151
A.2.2	roshumanoid	152
A.2.3	roswalker	154
	Appendix B Humanoid Platforms	156
B.1	Robotis OP	156
B.2	Softbank NAO	156

Appendix C	Stereotypical Kinematics in Human Locomotion	159
C.1	Kinematic References	159
C.2	Shared Perceptual Primitives	159
Appendix D	Fatigue-related Alterations in Muscle Synergies	184
D.1	Muscle Synergy Estimation	184
D.2	Classical Analysis	184
D.3	Synergy Analysis	207
Bibliography		232

List of Acronyms

AFO Adaptive Frequency Oscillator. 80

AI Artificial Intelligence. 2, 61, 65

API Application Programming Interface. 142, 151, 152

aPitch Ankle Pitch. 35

aRoll Ankle Roll. 35, 43

ASM Associative Skill Memory. 65, 68, 69, 73, 141

BBO Black Box Optimization. 8, 24, 34, 41, 42, 44, 55, 73, 85, 138, 149

BC Behavior Cloning. 8, 36, 58, 142

BIC Biceps Brachii. xix, 112, 113, 123, 124, 125, 128, 129, 135

BO Bayesian Optimization. 34

BoS Base of Support. 81

BRR Brachioradialis. xix, 112, 113, 123, 124, 128, 129, 135

CMA-ES Covariance Matrix Adaptation - Evolutionary Strategy. xxi, 33, 34, 41, 42, 44, 55, 85, 100, 149

CMP Compliant Movement Primitive. 20, 22

CNS Central Nervous System. 80, 106, 110, 114, 134

CoM Center of Mass. xvii, 14, 15, 37, 47, 48, 50, 51, 53, 54, 57, 81

CoP Center of Pressure. 14

CoT Cost of Transport. 17, 48, 54, 57, 139

CP Capture Point. xv, 15, 16

CPG Central Pattern Generator. xvi, xvii, xxi, 4, 13, 18, 19, 20, 36, 38, 39, 40, 41, 42, 43, 44, 46, 47, 48, 49, 50, 51, 53, 54, 55, 56, 57, 58, 80, 138, 139, 151, 154, 155

CrKR Cost-regularized Kernel Regression. 22

DARwin-OP Dynamic Anthropomorphic Robot with Intelligence - Open Platform. xvi, 35, 156

DCM Divergent Component of Motion. 15

DELTA Deltoid Anterior. xix, 112, 113, 123, 124, 125, 128, 129, 130, 134, 135

DELTm Deltoid Medial. xix, 112, 113, 123, 124, 125, 128, 129, 134

DELTP Deltoid Posterior. xix, 112, 113, 123, 124, 128, 129

DL Deep Learning. 2, 5, 23, 31, 67, 69, 70

DMP Dynamical Movement Primitive. xv, xvi, xvii, xxi, 20, 21, 22, 23, 24, 26, 27, 28, 29, 30, 31, 32, 34, 35, 36, 38, 40, 41, 42, 43, 44, 46, 47, 48, 51, 52, 53, 54, 55, 56, 57, 58, 62, 68, 69, 71, 72, 76, 85, 88, 98, 99, 100, 102, 103, 138, 139, 142, 144, 145, 147, 149, 151, 154, 155

DNN Deep Neural Network. 20, 61, 69

DoF Degree-of-Freedom. xv, xvi, xvii, xviii, xix, 4, 9, 13, 19, 25, 26, 30, 35, 36, 38, 39, 40, 41, 43, 44, 46, 47, 48, 49, 50, 51, 52, 57, 58, 83, 84, 87, 88, 94, 98, 101, 117, 152, 153, 156, 171, 172, 173, 174, 175, 176, 177, 178, 179, 180, 181, 182, 183

DRL Deep Reinforcement Learning. 143

EM Expectation-Maximization. 23

EMG Electromyography. xix, xx, xxii, 23, 106, 107, 109, 110, 111, 112, 116, 119, 121, 123, 125, 126, 128, 134, 135, 136, 137, 140, 143

ES Evolutionary Strategy. 34

FS Force Sensor. xvi, 38, 40

FSR Force-Sensing Resistors. 57, 156

GD Gradient Descent. 31

GMM Gaussian Mixture Model. 20, 31, 68

GMR Gaussian Mixture Regression. 20, 31

GPR Gaussian Process Regression. 22, 31, 59

GRF Ground Reaction Force. xvi, xvii, 35, 38, 40, 41, 47, 48, 50, 51, 53, 57

GRU Gate Recurrent Unit. 106

HMM Hidden Markov Model. 20, 103

hPitch Hip Pitch. 35

hRoll Hip Roll. 35, 43

hYaw Hip Yaw. xvi, xvii, 35, 43, 44, 49, 50, 52

HZD Hybrid Zero Dynamics. 18

ID Inverse Dynamics. 15

IK Inverse Kinematics. 15, 36

IL Imitation Learning. 4, 6, 8, 9, 12, 21, 22, 23, 24, 27, 30, 36, 57, 58, 61, 138, 139, 141, 142

IMU Inertial Measurement Unit. 81, 82, 156

ISB International Society of Biomechanics. xvi, xvii, 47, 49, 51, 52, 82

KPI Key Performance Indicator. 11, 57, 60, 105, 136, 137, 139

kPitch Knee Pitch. 35

L2 Lifelong Learning. iv, v, xxiii, 6, 7, 10, 23, 24, 29, 35, 55, 57, 58, 60, 64, 69, 70, 72, 100, 102, 139, 140, 141, 142

LfD Learning from Demonstration. 30, 31, 36, 102

LIP Linear Inverted Pendulum. 15

LLM Large Language Model. 6

LoB Loss of Balance. 80, 81

LOO Leave-One-Out. xviii, 84, 86, 87, 99

LSTM Long Short-term Memory. 23, 69, 106

LVM Latent Variable Model. 23

LWR Locally Weighted Regression. 20, 31

MAV Mean Absolute Value. 109

MDF Median Frequency. 106, 109, 127

MK Mann-Kendall. xxii, xxiii, 121, 128, 130, 131, 135

ML Machine Learning. 1, 4, 7, 9, 11, 12, 20, 34, 55, 58

MNF Mean Power Frequency. 106, 109

MoMPs Mixture of Motor Primitives. 21, 22, 66

MP Motor Primitive. v, xiii, xv, xvi, 4, 5, 9, 11, 13, 19, 21, 22, 23, 24, 25, 27, 28, 29, 30, 31, 32, 33, 35, 36, 40, 41, 43, 44, 47, 54, 55, 56, 57, 58, 59, 63, 64, 65, 68, 75, 99, 101, 110, 111, 138, 139, 140, 141, 143, 144

MPC Model Predictive Control. 18

MS Muscle Synergy. v, xix, xx, xxiii, 4, 10, 32, 105, 106, 110, 111, 112, 116, 118, 119, 122, 123, 124, 125, 126, 130, 132, 133, 134, 135, 136, 137, 138, 140, 142, 143, 184

MSK Musculoskeletal. 6, 19, 22, 106, 107, 108, 110, 111, 119, 123, 125, 132, 133, 134, 136, 137, 140, 143

MVC Maximum Voluntary Contraction. 108, 112, 118

NMF Non-negative Matrix Factorization. 84, 85, 88, 99, 101, 110, 111, 119

NN Neural Network. xi, 20, 61, 67, 69, 106

OOP Object-oriented Programming. 145, 146, 147, 148, 149, 150, 152, 153, 154, 155

PECM Clavicular Head of the Pectoralis Major. xix, 112, 113, 125, 128, 129, 135

PID Proportional-Integral-Derivative. 47, 54, 151, 156

PP Perceptual Primitive. v, xviii, xxi, 4, 5, 60, 61, 62, 64, 65, 67, 69, 70, 72, 73, 79, 84, 85, 90, 91, 94, 99, 100, 101, 102, 103, 104, 138, 139, 140, 143, 159, 171, 172, 173, 174, 175, 176, 177, 178, 179, 180, 181, 182, 183

ProMP Probabilistic Movement Primitive. 20, 23, 103

PS Policy Search. 34

PSD Power Spectral Density. xxii, 109, 120

RBF Radial Basis Function. 25, 36, 38, 40, 56, 74, 147

RL Reinforcement Learning. xi, 2, 4, 12, 21, 22, 23, 24, 27, 33, 34, 59, 70, 75, 104, 143, 149

RMS Root Mean Square. 106, 109, 110, 127, 130

RMSE Root Mean Square Error. 149

RNN Recurrent Neural Network. 18

ROS Robot Operating System. 142, 145, 150, 151, 152

RQ Research Question. 7, 9, 12, 62, 107

sEMG Surface Electromyography. xix, 108, 109, 112, 114, 117

SLIP Spring-Loaded Inverted Pendulum. 17

SVM Support Vector Machine. 67

TO Trajectory Optimization. 18, 24, 31, 32, 36

TRILAT Lateral Head of Triceps Brachii. xix, 112, 113, 123, 124, 125, 128, 129, 135

TRILON Long Head of Triceps Brachii. xix, 112, 113, 123, 125, 128, 129, 135

TTP Total Power. 120

UML Unified Modeling Language. 146, 147, 148, 149, 150, 152, 153, 154, 155

UUID Universally Unique Identifier. 147

VAF Variance Accounted For. xvi, xviii, xx, xxi, xxii, 43, 44, 45, 46, 55, 56, 88, 90, 91, 93, 101, 119, 122, 123, 124, 125, 126, 133, 134

WaHP Width at Half Peak. 130

WRMSD Work-related Musculoskeletal Disorder. 105, 107, 137

XML Extensible Markup Language. 147

ZMP Zero-Moment Point. 14, 15, 18

List of Figures

1	Examples of bioinspired robots. Left: Salamandra Robotica [Crespi et al., 2013]; Center: Reproduction [Collins et al., 2001] of the 2D gravity-powered passive design introduced by McGeer [1990]; Right: Self-powered Cornell’s 3D passive dynamic walker [Collins & Ruina, 2005]	3
2	Conceptual overview of the present work; chapters 2, 3 and 4 constitute parallel independent studies on different applications of primitive theory, contextualized by Chapter 1 and with relevant outcomes summarized in Chapter 5; practical objectives (1.1) accomplished in each chapter are highlighted ($\square$).	9
3	Conceptual visualization of ZMP-based locomotion. Stability is ensured as long the estimated position of the ZMP remains within the support polygon; adapted from Kajita & Espiau [2008].	16
4	Illustration of CP-based walking; adapted from Engelsberger et al. [2011].	16
5	Conceptual diagram of the control strategy proposed in this chapter, in the case of a K DoF platform; a generic primitive library $\mathbf{P}$ provides abstract task- and platform-agnostic MPs combined in different motor skills $\mathcal{S}^s$, archived in a skill library $\mathbf{S}$; closed-loop control is achieved with task-specific feedback mechanisms that modulate ($\bullet\bullet$) prescribed trajectories y or phase ϕ, s	25
6	Top Left: Multiple basis functions Ψ_n on a periodic phase space; Top Right: DMP transformation system, converging to $g = 10$; Bottom Left: DMP rhythmic canonical system <i>i.e.</i> linear pacer, providing values of ϕ ; Bottom Right: DMP discrete canonical system, following different decay approaches: exponential (blue), linear (orange) and sigmoid (green).	27
7	Illustration of different skill learning modalities achieved with different constraints to θ ; the number of primitives P' , as well as N in extrinsic and expansive learning, are viewed as open parameters; $\mathbf{W}_0, \Lambda_0, \Delta_0$ denote preset parameterization; for simplicity, in intrinsic and sequential modalities, $\mathbf{m}$ (18) can be omitted <i>i.e.</i> $\mathbf{m} = \mathbf{I}_P$ for small P (9) <i>e.g.</i> during early cognitive stages.	33

8	Humanoid platforms employed in this work. Left: Robotis OP 1 (formerly DARwIn-OP) [ROBOTIS, 2019]; Right: Softbank Robotics NAO V5 [Robotics, 2019]. Refer to Appendix B: Humanoid Platforms for detailed information about both robots.	35
9	Conceptual diagram extending the generic MP-based control strategy in Figure 5 to bipedal walking in the Robotis OP and the SoftBank NAO (Figure 8); in both platforms, each k^{th} leg DoF is controlled by a DMP transformation system y_k (11) and force-based feedback mechanisms close the control loop.	36
10	Resulting kinematic trajectories y^{CPG} for all $K = 6$ DoFs of the right leg when employing (28 - 33) for a full period i.e. $[-\pi, \pi]$; Values of A^{motion} , $O_{i,j}$, ω and h presented in Table 1 (OP).	39
11	Kinematic overview of how CPG walking [Matos & Santos, 2012, 2014] is reproduced with the generic MP-based control strategy, integrating Force Sensor (FS) feedback loops into DMP dynamics; α , β and, notably, $f_k^*(\phi_i)$, $i \in \{\text{left}, \text{right}\}$ are shared between both legs; both hip yaw joints are mechanically coupled in the NAO; lastly, N_i^{FS} evenly distributed FSs measuring GRFs in both feet are assumed.	40
12	Conceptual diagram describing the two-stage procedure employed to teach bipedal walking to both robots; in the first, expansive learning (2.3.3) is used to populate $\mathbf{P}$ and $\mathbf{S}^{\text{OP}}$; in the latter, intrinsic learning is used to transfer skill parameterization to $\mathbf{S}^{\text{NAO}}$, reusing MPs in $\mathbf{P}$. In both robots, $K = 6$	43
13	Evolution of $J(\theta, \mathbf{y}^{\text{CPG}})$ during skill learning (2.4.3), for $P \in \{1, 2, 3, 4, 5, 6\}$, in the OP (Top) and the NAO (Bottom); final values presented in Figure 14 and Table 4.	45
14	Average final optimization cost $J(\theta^{\text{best}}, \mathbf{y}^{\text{CPG}})$ (Left) and Variance Accounted For (VAF) between DMP learned trajectories $\mathbf{y}^{\text{DMP}}$ and reference $\mathbf{y}^{\text{CPG}}$ (Right), with $P \in \{1, 2, 3, 4, 5, 6\}$ (Table 4), across 3 repetitions of the skill learning procedure (2.4.3); error bars illustrate measured standard deviation.	45
15	Reconstructed DMP trajectories $\mathbf{y}^{\text{DMP}} = \mathbf{y}^{\theta^{\text{best}}}$ for each of the $K = 6$ DoFs in the OP, with $P = 3$, as well as reference CPG curves (Figure 10).	46
16	Reconstructed DMP trajectories $\mathbf{y}^{\text{DMP}} = \mathbf{y}^{\theta^{\text{best}}}$ for each of the $K = 6$ DoFs in the NAO, with $P \in \{3, 4, 5\}$, as well as reference CPG curves (Figure 10).	46
17	Resulting set of $P = 3$ MPs and motor skill $\mathcal{S} = \{\mathbf{\Lambda}, \mathbf{\Delta}\}$ learned in the OP.	47
18	Resulting motor skills $\mathcal{S} = \{\mathbf{\Lambda}, \mathbf{\Delta}\}$ learned in the NAO with $P = \{3, 4, 5\}$	47
19	Snapshots of the simulated OP and NAO bipeds walking on flat ground.	48
20	DoF position control commands y_k^{CPG} (Figure 10) and resulting state q_k^{CPG} values in the OP during CPG walking on flat ground. Each DoF is presented according to ISB joint coordinate systems.	49
21	DoF position control commands y_k^{CPG} (Figure 10) and resulting state q_k^{CPG} values in the NAO during CPG walking on flat ground. Each DoF is presented according to ISB joint coordinate systems. Hip Yaw is omitted due to being fixed at 0° (Figure 16).	49
22	Phase diagrams of leg DoFs in the OP recorded during CPG walking on flat ground.	50

23	Phase diagrams of leg DoFs in the NAO recorded during CPG walking on flat ground. Hip Yaw is omitted due to being fixed at 0° (Figure 16).	50
24	Leg Phase ϕ_i and total GRF measured $\sum x_j^{\text{FS}}$ (Left), as well as CoM position $\{x_{\text{CoM}}, y_{\text{CoM}}, z_{\text{CoM}}\}$ and foot strike locations (Right) in the OP during CPG walking on flat ground.	50
25	Leg Phase ϕ_i and total GRF measured $\sum x_j^{\text{FS}}$ (Left), as well as CoM position $\{x_{\text{CoM}}, y_{\text{CoM}}, z_{\text{CoM}}\}$ and foot strike locations (Right) in the NAO during CPG walking on flat ground.	51
26	DoF position control commands y_k^{DMP} (Figure 15) and resulting state q_k^{DMP} in the OP values during DMP walking on flat ground. Each DoF is presented according to ISB joint coordinate systems.	51
27	DoF position control commands y_k^{DMP} (Figure 16) and resulting state q_k^{DMP} in the OP values during DMP walking on flat ground. Each DoF is presented according to ISB joint coordinate systems. Hip Yaw is omitted due to being fixed at 0° (Figure 16).	52
28	Phase diagrams of leg DoFs in the OP recorded during DMP walking on flat ground.	52
29	Phase diagrams of leg DoFs in the NAO recorded during DMP walking on flat ground. Hip Yaw is omitted due to being fixed at 0° (Figure 16).	52
30	Leg Phase ϕ_i and total GRF measured $\sum x_j^{\text{FS}}$ (Left), as well as CoM position $\{x_{\text{CoM}}, y_{\text{CoM}}, z_{\text{CoM}}\}$ and foot strike locations (Right) in the OP during DMP walking on flat ground.	53
31	Leg Phase ϕ_i and total GRF measured $\sum x_j^{\text{FS}}$ (Left), as well as CoM position $\{x_{\text{CoM}}, y_{\text{CoM}}, z_{\text{CoM}}\}$ and foot strike locations (Right) in the NAO during DMP walking on flat ground.	53
32	General overview of different <i>symbolic</i> , <i>emergent</i> and <i>hybrid</i> cognitive architectures (Bottom) and associated sensor modalities (Top); adapted from Kotseruba et al. [2016].	66
33	Conceptual flowchart illustrating the dynamical perceptual model proposed in this chapter; a library M containing perceptual memories $\mathcal{M}^m$ provides sensor-specific estimates $\hat{x}_k$ and $\hat{s}_k$ which are compared, in real time, with transient context-specific measurements x_k^* through d^m (69); phase ϕ_t and time scale τ is assumed to be inferred with sensor-based heuristic methods.	73
34	Illustrative example of contextual <i>search</i> (Algorithm 4) on a library M with $M = 6$ perceptual memories; as highlighted ($\square$), $\mathcal{M}^5$ and ϕ_t indicate maximum $d^m(\mathbf{x}_t^*)$ over $\Phi = [\phi_{\min}, \phi_{\max}]$	77
35	Illustrative example of contextual <i>bootstrap</i> (Algorithm 5) on a library M with $M = 6$ perceptual memories; Algorithm 4 is sequentially repeated for perceptual inputs $\mathbf{x}_t$, with ϕ_t ($\square$) representing maximum $d^m(\mathbf{x}_t^*)$ over Φ in each case; given the input sequence, $\mathcal{M}^3$ represents the best match in M	77
36	Conceptual workflow and hierarchical organization between the different levels of perceptual context identification described; given $\mathcal{C}_t$, Algorithm 3, 4 and 5 are sequentially chained according to the success or failure of context identification, with increasing computational overhead.	78

37	Experimental apparatus (Top) adopted in the human walking trials, with brief snapshots of a perturbation being induced by the operator (Bottom). Adapted from [Ferreira, 2022].	82
38	Conceptual flowchart of the processing steps conducted to extract fixed-length reference periods $\bar{\mathbf{x}}^{\text{REF}}, \mathbf{s}^{\text{REF}} \in \mathbb{R}^{K \times L}$ from periodic kinematic recordings $\mathbf{q}_k$, using left knee flexion/extension angle as reference for segmentation; ι_t denotes the time stamp of the t^{th} sample, $\hat{\iota}_t^i$ the interpolation time stamps in the i^{th} cycle, $q_{\min}$ the knee angle threshold, and L the number of reference samples.	84
39	Illustration of the Leave-One-Out (LOO) cross-validation scheme adopted in the creation of participant-specific perceptual libraries; kinematic patterns $\mathbf{q}$ are sourced from trials matching velocity and inclination in all remaining participants; segmented, interpolated and averaged according to the procedure in Figure 54; and the resulting reference trajectories $\bar{\mathbf{x}}^{\text{REF}}, \mathbf{s}^{\text{REF}}$ fed, separately, to the skill learning procedure described in Chapter 2 (2.3.3) (23 - 24). S2 and S6 were discarded due to data corruption, resulting in a total of 9 folds for the remaining participants. As highlighted, in each fold, all trials conducted with the skipped participant, with and without perturbation, compose the test subset.	86
40	Kinematic references $\bar{\mathbf{x}}^{\text{REF}} \pm \mathbf{z}^{\text{T}} \mathbf{s}^{\text{REF}}$, with $\mathbf{z} = 1.96\mathbf{1}_K$, obtained for 6 lower limb DoFs of S3	87
41	Global $\bar{\mathbf{x}}^{\text{GLB}} \pm \mathbf{z}^{\text{T}} \mathbf{s}^{\text{GLB}}$ and participant-specific $\bar{\mathbf{x}}^{\text{REF}} \pm \mathbf{z}^{\text{T}} \mathbf{s}^{\text{REF}}$ kinematic references, with $\mathbf{z} = 1.96\mathbf{1}_K$, obtained for 3 lower limb DoFs of S9 in 1.8 km/h walking at 10° incline.	88
42	Illustration of Spearman's rank correlation r_s between different dimensions in different walking modalities; values closer to 1 and -1 imply greater correlation between different dimensions.	89
43	Variance Accounted For (VAF) between original and reconstructed patterns obtained when factorizing $\mathbf{q}$ with $r_{\text{NMF}} \in \{1, 2, 3, 4, 5, 6\}$ components; higher values imply better reconstruction quality; $\text{VAF} \geq 0.9$ was assumed to represent an accurate reconstruction.	90
44	Evolution of $J(\theta, \bar{\mathbf{x}}^{\text{REF}})$ (Top) and $J(\theta, \mathbf{s}^{\text{REF}})$ (Bottom) (23) during memory consolidation with the skill learning procedure (Chapter 2, 2.3.3) with S7 walking at 5.4 km/h and a 0° incline under different learning strategies (Table 7); in the isolated learning case, dimension-specific costs $J_k(\theta, \bar{\mathbf{x}}^{\text{REF}})$ are sequentially compounded; final values presented in Table 9.	92
45	Comparison between original references $\{\bar{\mathbf{x}}^{\text{REF}}, \mathbf{s}^{\text{REF}}\}$ (Figure 40) and trajectories learned with the isolated (no shared PP, $P = 18$) and two-stage ($P = 5$) strategies (Table 7), for 6 lower limb DoFs of S3 during walking at 1.8 km/h on a 0° incline.	94
46	Resulting set of $P = 5$ PPs and skill $\mathcal{S}^*$ (16, 66) describing $\bar{\mathbf{x}}^{\text{REF}}$ learned with the two-stage strategy (Table 7) for walking at 1.8 km/h on a 0° incline with S3 (Top) and S7 (Bottom) ; dimension labels are provided in Table 8.	94

47	Illustration of the gait classification procedure (3.4.4) in a trial with S4 at 5.4 km/h on a 10° incline, with two perturbations during left heel strike ($\square$, 3.4.3.2); captured lower limb kinematics $\mathbf{q}$ are used as perceptual input $\mathbf{x}^*$; upper and lower edges indicate perceptual memory in $\mathcal{C}_t$ (70) when $\mathcal{C}_t \neq \emptyset$ (Algorithm 4); detection matches at P0 and P1 , respectively 16 ms <i>after</i> and 16 ms <i>before</i> ground truth, corresponding to 1 sample, at a sampling frequency of 60 Hz.	95
48	Conceptual illustration of the underlying principle behind Muscle Synergies, wherein EMG envelope $\mathbf{E}_m$ is factorized into synergies $[\mathbf{w}^1, \mathbf{w}^2, \mathbf{w}^3]$, recruited according to $[\mathbf{h}^1, \mathbf{h}^2, \mathbf{h}^3]$, often representing functional kinematic primitives e.g. elbow flexion, elbow extension and shoulder extension.	111
49	Participant instrumented with the sEMG sensors (<i>Trigno Avanti</i> [®]) on: 1) PECM, 2) DELTA, 3) DELTM, 4) DELTP, 5) TRILON, 6) TRILAT, 7) BIC, and 8) BRR (Left); and the inertial motion capture system (<i>MVN Awinda</i> [®]) in addition to the sEMG sensors (Center, Right).	113
50	Conceptual flowchart illustrating the two separate parts of the experimental protocol: P1 , comprising 3 uni-articular exercises (A, B, C); and P2 , where a composite motion is repeated until exertion (D). Generic synergies $\mathbf{W}_{ABC}$ are extracted from P1 ; In P2 , classical and synergy analysis (4.3.4.5) is conducted on each individual repetition of D	114
51	P1 exercises A, B and C (Top) and P2 exercise D (Bottom).	115
52	Sequence of movements composing exercise D , as depicted in Figure 51 (Bottom), including uni-articular motions from A, B and C ; shoulder extension (C) is only performed after the first repetition.	115
53	Conceptual flowchart describing the extraction of linear envelope $\mathbf{E}_m$ for the m^{th} muscle under study, from raw sEMG $\mathbf{e}_m$ recorded at $f_s = 2048$ Hz; in P2 (–), classical EMG analysis is conducted on filtered $\mathbf{e}_m$ within a fixed window Δt <i>clipped</i> around the global maximum of $\mathbf{E}_m$	117
54	Conceptual flowchart of the kinematic segmentation procedure; $\mathbf{q}^{\text{REF}}$ denotes an exercise-specific reference kinematic DoF, used to estimate I sets of L time values $\hat{t}_l^i$, relating to the l^{th} sample of the i^{th} repetition executed during the trial, $l \in \{1, 2, \dots, L\}$ and $i \in \{1, 2, \dots, I\}$	117
55	Illustrative example of the kinematic segmentation procedure, for exercise C with S8	118
56	Cascade of processes applied to the EMG envelope in order to extract the synergistic organization of underlying muscular behavior with and without fatigue; in P2 (–), each repetition is processed individually (bypassing concatenation) and NNMF is constrained to the global synergies $\mathbf{W}_{ABC}$	119
57	Illustration of Spearman's rank correlation magnitude $ r_s $ across EMG data recorded with different participants; data from individual muscles is concatenated sequentially; values closer to 1 imply greater correlation.	122

58	Illustration of Spearman's rank correlation magnitude $ r_s $ across kinematic data recorded with different participants; all axis of the <i>RightT4Shoulder</i> , <i>RightShoulder</i> , <i>RightElbow</i> and <i>RightWrist</i> joints are concatenated sequentially; values closer to 1 imply greater correlation.	122
59	Subject and average VAF obtained in P1 for $n \in \{1, \dots, M\}$; values ≥ 0.9 were assumed to represent an accurate reconstruction of the original EMG envelope E	123
60	Synergies W_{ABC} (<i>left</i>) and activation profile H_{ABC} (<i>right</i>) obtained in P1 ; for each of the $N_{\text{ABC}} = 4$ synergies, average and inter-participant variance ($\pm$ std. deviation) is presented.	124
61	Comparison between relevant right-arm kinematics q_k , EMG envelope E_m and MSs recruitment H_{ABC} during P1 with S14 , around exercise-specific maxima in E_m (:).	125
62	VAF obtained with each participant during exercise D , as well as resulting second-order polynomial trend (over i) approximation.	126
63	Comparison, for the first repetition ($i = 1$) of D for S12 (Left) and S14 (Right) , of arm kinematics q_k with EMG envelope E_m and MS activation H_D ; in P2 , W_D = W_{ABC}	126
63	Scale- and time-normalized (interpolated) variation of RMS (Top) and MDF (Bottom) (4.2.3) during P2 ; average and inter-participant variation ($\pm$ std. deviation) presented for each muscle.	127
64	Scale- and time-normalized (interpolated) variation of RMS (Left) and WaHP (Right) values computed from H_D during P2 across all subjects. Average and inter-participant variation ($\pm$ std. deviation) presented for each MS.	130
65	Synergy recruitment H_Dⁱ for $i \in \{1, 2, \dots, 64\}$ observed during D in the case of S5	131

List of Tables

1	Hand-tuned CPG parameters (28 - 33); same parameterization used in both legs;	44
2	DMP parameters employed in the generic control strategy (2.4.2).	44
3	CMA-ES parameters employed in the motor skill learning procedure (2.4.3).	44
4	Average $J(\boldsymbol{\theta}^{\text{best}}, \mathbf{y}^{\text{CPG}})$ and Variance Accounted For (VAF) (Figure 14), as well as estimated error per sample $\hat{e}_y = 1/KT \times J(\boldsymbol{\theta}^{\text{best}}, \mathbf{y}^{\text{CPG}})$, across 3 repetitions of the skill learning procedure (2.4.3).	46
5	Gait metrics obtained during CPG and DMP simulated walking with both robots.	53
6	Overview of experimental conditions in walking trials conducted with each participant, showing all combinations of velocity, treadmill incline, perturbed leg - left (L) and right (R) - and gait event - heel strike (HS) and toe off (TO); each trial was repeated with and without perturbations; a total of 32 trials were conducted per participant, and executed in a blind randomized order.	83
7	Different learning strategies (2.3.3) adopted during memory consolidation, applied separately to $\bar{\mathbf{x}}^{\text{REF}}$ and $\mathbf{s}^{\text{REF}}$, to populate all $\mathbf{M}^{\mathbf{s}^j}$, and corresponding number of PPs, CMA-ES parameters λ and σ , learning dimensionality n (given $N = 20$ and $K = 18$) as well as the total number of repetitions (#) required throughout the 9 cross-validation folds, to fully consolidate $\{\bar{\mathbf{x}}^{\text{REF}}, \mathbf{s}^{\text{REF}}\}$ into $\mathcal{M}$ (66).	85
8	Average magnitude of Spearman's rank correlation coefficient computed for each dimension in different walking modalities; higher values imply greater average pairwise rank correlation magnitude $ r_s $	89
9	Memory consolidation results across all 9 cross validation folds, including convergence time I^{conv} , final costs $J(\boldsymbol{\theta}^{\text{best}}, \bar{\mathbf{x}}^{\text{REF}})$ and $J(\boldsymbol{\theta}^{\text{best}}, \mathbf{s}^{\text{REF}})$, estimated errors per sample $\hat{e}_{\mathbf{x}} = 1/KL \times J(\boldsymbol{\theta}^{\text{best}}, \bar{\mathbf{x}}^{\text{REF}})$ and $\hat{e}_{\mathbf{s}} = 1/KL \times J(\boldsymbol{\theta}^{\text{best}}, \mathbf{s}^{\text{REF}})$, and measured $\text{VAF}_{\mathbf{x}}$ and $\text{VAF}_{\mathbf{s}}$, with different learning strategies (Table 7).	93
10	Classification metrics (<i>sensitivity</i> , <i>specificity</i> and <i>accuracy</i>) of context identification using Algorithm 4 with participant-specific perceptual libraries $\mathbf{M}^{\mathbf{s}^j}$, as well as global averages ($\pm$ std. deviation).	95

11	Classification metrics (<i>sensitivity</i> , <i>specificity</i> and <i>accuracy</i>) and <i>detection time</i> (s) obtained in the detection of perturbed walking using Algorithm 4 with participant-specific perceptual libraries $\mathbf{M}^{\mathbf{s}^j}$, as well as global averages ($\pm$ std. deviation), grouped by walking velocity and terrain incline.	96
12	Classification metrics (<i>sensitivity</i> , <i>specificity</i> and <i>accuracy</i>) and <i>detection time</i> (s) obtained in the detection of perturbed walking using Algorithm 4 with participant-specific perceptual libraries $\mathbf{M}^{\mathbf{s}^j}$, as well as global averages ($\pm$ std. deviation), grouped by type of perturbation.	97
13	Time- and frequency-domain metrics commonly extracted from EMG [Chang, 2018; Najarian & Splinter, 2012; Phinyomark et al., 2012]; e_i denotes the raw EMG signal, E_t the extracted envelope curve, $S_{ee}[f]$ the estimated Power Spectral Density (PSD) of e_i , f_s the sampling frequency, and T the number of samples in e_i	109
14	Overview of the participant pool recruited for the study, including weekly exercising habits, sleep hours and stimulant consumption on the date of the trials.	113
15	Metrics comprising the analysis of the synergy activation $\mathbf{H}_{\mathbf{D}}^i$ in the i^{th} repetition, during P2 ; $h_l^{n,i}$ represents the activation of the n^{th} synergy at the l^{th} sample and $S_{hh}^i[f]$ the estimated PSD of $\mathbf{h}_{\mathbf{D}}^{n,i}$	120
16	Self-perceived fatigue level [Borg, 1998] reported by each participant after each exercise.	122
17	Average rank correlation magnitude $ r_s $ of EMG and kinematic data across different participants <i>i.e.</i> Figure 57 and Figure 58 column/row averages; higher values imply greater pairwise correlation.	123
18	Number of executed repetitions (I) and average VAF obtained in exercise D by each participant.	124
19	p -values obtained when conducting the MK test over classical metric values computed during P2 ; results presented for each muscle, and for all participants under a one-tailed test (any significant trend) and two-tailed test (for same -direction trends); values colored according to the rejection (green <i>i.e.</i> trend) or non-rejection (red <i>i.e.</i> no trend) of H_0 under a significance level of $\alpha = .05$	128
20	Average Mann-Kendall normalized test statistic Z across all participants for each muscle and classical metric (Table 13); higher absolute magnitude (green) implies higher variation, while values closer to 0 (red) suggest no significant trend; t-tests where $p > .05$ (Table 19) greyed out.	129
21	Average determination coefficient R^2 obtained for the polynomial models (4.3.4.6) for all classical metrics; higher values (green) imply a better model fit; values where $p > .05$ (Table 19) greyed out.	129

22	p -values obtained when conducting the MK test over synergy metric values computed during P2 ; Results presented for each synergy, and for all subjects under a one-tailed test (any significant trend) and two-tailed test (for same -direction trends); Values colored according to the rejection (green <i>i.e.</i> trend) or non-rejection (red <i>i.e.</i> no trend) of H_0 under a significance level of $\alpha = 0.05$	131
23	Average Mann-Kendall normalized test statistic Z across all subjects for each MS and metric (Table 15); higher absolute magnitude (green) implies higher variation, while values closer to 0 (red) suggest no significant trend; t-tests where $p > .05$ (Table 22) greyed out.	132
24	Average determination coefficient R^2 obtained for the polynomial models (4.3.4.6) for all synergy metrics; higher values (green) suggest a better model fit; values where $p > .05$ (Table 22) greyed out.	132
25	Overview of hallmark features in Lifelong Learning (L2) systems [Kudithipudi et al., 2022] and where each is addressed in this work; task-agnostic learning, while not directly addressed, is achievable with the outcomes of Chapter 2 and Chapter 3.	142

Introduction

Over the last century, the relevance of automation in society has increased significantly, to the point where robotic systems have become ubiquitous in both domestic and the gross majority of industrial settings. This modern dependency on robots has, unsurprisingly, seeped into the daily life of the average citizen, who deals with these types of automated machines on a daily basis, from self-driven vacuum machines, to kitchen appliances, to the very vehicles they use while commuting. As a result, current robots are often required to share environments, whether domestic or industrial, with humans and/or other robotic agents. This brings about extra potential for the development of collaborative robotic interaction, while at the same time requiring adaptation to account for dynamic perturbations, obstacles, or different task demands [Calinon et al., 2012]. Particularly in domestic environments - by nature more chaotic than predictable industrial scenarios - it is both impractical and time-consuming to curate a selection of motor behaviors to be executed in all possible situations; nevertheless, for a long time this straightforward, albeit limited, approach was the choice strategy behind the design of robotic control policies, resulting in *pseudo*- or semi-autonomous systems that could rarely function outside a limited number of well-structured workspaces and within very constrained task parameters [Sigaud & Stulp, 2019].

For robots to be able to safely and productively operate *in tandem* with humans, they must be furnished with autonomous controllers able to decide *what to do* in the most varied scenarios, without the need of supervisory input [Ajoudani et al., 2018; Peters et al., 2011; Sigaud & Stulp, 2019]. To that end, providing an agent with the capability to incrementally build its own behavior database, and update it as it gathers information on - and interacts with - the surrounding world, represents a more powerful and versatile approach to adaptive robotic control. While *adaptation* can broadly encompass different degrees of flexibility in robotic movement, at top-level it implies *lifelong learning*, whereby robots are able to continuously acquire new skills without jeopardizing existing ones, and extrapolate from similar tasks [Hospedales et al., 2021; Kudithipudi et al., 2022]. Over the last few years, scientific research in the field of Machine Learning (ML) has focused in developing techniques to generalize preset behaviors across different scenarios while complying with task demands [He et al., 2023; Kramberger et al., 2017; Kroemer et al., 2020; Kupcsik et al., 2017; Singh et al., 2017; Zeng et al., 2022]. Ultimately, this can allow robots to share, in a *care-free* manner, environments with humans and other robots, as well as cooperate with both in tasks that would

be cumbersome or time-consuming otherwise [Ajoudani et al., 2018].

Correspondingly, the field of Artificial Intelligence (AI) has progressed to the point where truly autonomous robots are expected to become a trivial component of daily life in the near future. Notwithstanding, the practicality of designing artificial systems capable of continuous self-improvement remains a significant hurdle, with current solutions for robotic skill learning still exhibiting difficulty in scaling to more elaborate platforms e.g. industrial manipulators or humanoid robots, due to their high-dimensional action- and state-spaces [Peters et al., 2011]. This increase in the complexity of robotic control, known as the *curse of dimensionality* [Bellman, 1957], often presents an inescapable control challenge; in essence, the dimensionality of any control strategy becomes critical to - or rather, has relevant implications in - any learning process, which is a cornerstone of autonomous robotics and a hard requirement in real-time adaptive systems [Peters et al., 2011]. While modern Reinforcement Learning (RL) solutions have made long strides towards numerical tractability in such platforms [Kober, 2014; Kroemer et al., 2020; Li, 2018; Mnih et al., 2015; Sigaud & Stulp, 2019], assisted by modular and parametric reduced-order movement representations [Neumann et al., 2014; Saveriano et al., 2021], as well as recent breakthroughs in Deep Learning (DL) that rekindled the interest of the research community into robotic problems [Gu et al., 2017; Ibarz et al., 2021; Levine et al., 2018; Mnih et al., 2015; Silver et al., 2016], these techniques still require a prohibitively high number of interactions for the majority of real robotic tasks [Chatzilygeroudis et al., 2020; Levine et al., 2016]; as such, their applicability to real robotic systems, particularly high-dimensional platforms, is either greatly reduced or downright impractical outside a research context [Kamthe & Deisenroth, 2018]. The bottom line is that in the real world - where adaptation is most important - robots lack the luxury of access to computational clusters able to simulate thousands of interactions, extensive datasets to extrapolate from, as well as time and safety assurances to actively and preemptively attempt trial-and-error exploration; and therefore must interact with their environment in real-time with limited resources and quickly adapt their behavior while extracting the most information from each interaction [Calinon, 2016; Chatzilygeroudis et al., 2020; Kroemer et al., 2020].

In the pursuit of alternative strategies to deal with this complexity, roboticists have historically borrowed inspiration from nature, namely how vertebrate animals move and behave, and extensively analyzed their movements from a cognitive and biomechanical standpoint, which has yielded valuable research outputs [Arena et al., 2004; Ijspeert, 2014; Menciassi & Laschi, 2014; Miller & Clark, 2015], such as replicating running, jumping, swimming and even climbing in robots [Chevallier et al., 2008; Degallier et al., 2011; Miller & Clark, 2015]. Yet, when compared to the musculoskeletal systems of vertebrate animals, modern robots are still greatly lacking in versatility and adaptability. Not only are natural movements fluid and energy-efficient, they can intuitively be adjusted between different scenarios and in response to perceptual stimuli, in real-time, and without compromising task success, which is particularly prevalent in behaviors that require direct interaction with - and adaptation to - the surrounding environment. This flexible and adaptive nature, a result of morphological and physiological adaptation of animals in response to the evolutionary requirement to deftly navigate complex, dynamic, erratic or otherwise unreliable environments, has been under the spotlight of the scientific community for several decades [Ijspeert, 2008], leading to neural and musculoskeletal models with special focus in different locomotive modalities, ranging from aquatic



Figure 1: Examples of bioinspired robots. **Left:** Salamandra Robotica [Crespi et al., 2013]; **Center:** Reproduction [Collins et al., 2001] of the 2D gravity-powered passive design introduced by McGeer [1990]; **Right:** Self-powered Cornell's 3D passive dynamic walker [Collins & Ruina, 2005]

and amphibious animals, such as lampreys [Ekeberg & Grillner, 1999] and salamanders [Ijspeert et al., 2007], to legged insects [Büschges et al., 2008], cats [Ivashko et al., 2003] and ultimately humans [Hase et al., 2003; Kim et al., 2011; Ogihara & Yamazaki, 2001; Paul et al., 2005]. These models are notably useful when artificially and/or empirically - *i.e.* through numerical simulations - replicating animal behaviors, crucial to understand not only the high-level control strategies behind biological movement, but also the role of the body's dynamical and the mechanical properties [Geyer & Herr, 2010; Pearson et al., 2006]. The field of *bioinspired* robotics or *biorobotics* borrows these insights and integrates them into the design of robotic platforms and in their control (Figure 1), of which practical applications are present across our society [Craddock et al., 2022; Ijspeert, 2014; Karakasiliotis et al., 2016; Menciassi & Laschi, 2014; Wang et al., 2011; Zhou & Bi, 2012]. Considering learning as a hallmark feature of biological organisms, these provide a valuable template for a new kind of machine intelligence, able to robustly navigate dynamic settings, deal with unforeseen disturbances, and continually improve their own skill set for better performance [Hospedales et al., 2021; Kudithipudi et al., 2022].

While replicating natural behaviors in robotic systems, particularly humanoid platforms, still presents multiple challenges beyond the aforementioned high dimensionality, such as non-linearity, model inaccuracies and the frequent need for limb coordination [Ijspeert, 2014; Reher et al., 2020], there are multiple levels on which biorobotics can prove valuable to scientific research. On one hand, they constitute appealing testbeds for working theories in neuroscience and biomechanics, due to sheer practicality: robots provide an objective morphology, and allow high repeatability within safety constraints; in addition, they can complement numerical simulations by bypassing the complexity of nonlinear dynamics resulting from the interaction and contact with the environment, which is notoriously hard to implement in compliant surfaces *e.g.* sand, mud or gravel [Ijspeert, 2014]. Pushing this notion even further, robots have been shown to be able to mirror muscular behavior with significant accuracy, whether by passive actuation [Verrelst

et al., 2005] or active impedance control [Semini et al., 2016]. In the opposite direction, biology can provide multiple ways to simplify robotic design and control: 1) by influencing robot design at morphological and/or electromechanical level - such as musculotendon actuators [Ananthanarayanan et al., 2012; Craddock et al., 2022], self-stabilizing pantograph-like limb design [Spröwitz et al., 2013], or other compliant muscle-like control [Semini et al., 2016; Verrelst et al., 2005]; and, alternatively, 2) by extrapolating from the high-level control architecture of musculoskeletal systems, from which concepts like Central Pattern Generators (CPGs) and Motor Primitives (MPs) have emerged.

It's precisely the notion that, despite the complex morphologies found in vertebrate musculoskeletal systems - with multiple redundant Degrees-of-Freedom (DoFs) and continuous action-spaces [Pastor et al., 2013; Rückert et al., 2013], animals are able to segment and sequence complex movements into simpler adaptable functional units which has over the last few decades gathered interest both from the field of neurobiology, which seeks to better understand the cognitive process behind this control structure [Flash & Hochner, 2005; Mussa-Ivaldi, 1999], and robotic control, which aims to harness that same structure in an effort to simplify sophisticated motions into a combination of low-dimensional modular components [Saveriano et al., 2021]. These functional units, known as Motor Primitives (MPs)¹ [d'Avella & Bizzi, 2005; Giszter et al., 2001; Schaal, 1999], can provide a versatile and abstract representation for robotic motor skills, which is particularly useful, for instance, when teaching and generalizing skills between different contexts and platforms, as well as in autonomous trial-and-error applications [Pastor et al., 2013; Schaal et al., 2005]; moreover, such a modular approach allows decoupling high-level planning from the low-level command trajectories [Degallier et al., 2011]. Accordingly, MP-based control constitutes today a popular and widespread approach for trajectory control across different robotic applications, commonly relying on dynamical systems to encode kinematic [Ijspeert et al., 2013, 2003; Samant et al., 2016; Stulp et al., 2012] and/or dynamic [Deniša et al., 2016; Petrič et al., 2018] control patterns. Moreover, as a dimensionality reduction strategy, it facilitates the development of Machine Learning (ML) solutions in high-dimensional tasks, with particular relevance to Imitation Learning (IL) wherein robots are shown, kinesthetically, optimal behaviors through demonstration [Ravichandar et al., 2020; Schaal, 1999], and RL, wherein robotic skills can be fine-tuned through interaction [Chatzilygeroudis et al., 2020; Stulp & Sigaud, 2013], both key traits observed in humans [Nakanishi et al., 2004].

MPs have been observed at both kinematic [Moro et al., 2012; Rückert & d'Avella, 2013] and neuromuscular² [Dominici et al., 2011; Grillner & El Manira, 2020; Zandvoort et al., 2022] levels, hinting at a general simplification strategy found in natural processes; beyond motor control, the existence of Perceptual Primitives (PPs) has long been theorized to be part of the cognitive toolset employed by humans (and animals) to extrapolate structure from raw sensor stimuli [Aslin & Smith, 1988; Frégnac & Bathellier, 2015; Palmer, 2015], namely in visual perception [Chen, 1982, 2005; Huang et al., 2010; Mausfeld, 2003; Mayeli, 2019] - with contributions at the root of computer vision and image processing [Huang et al., 2009, 2010; Irani & Ware, 2003], but also explored in other perceptual modalities [Klatzky & Reed,

¹Also referred to as *movement schemas*, *basis behaviors* and *units of action* [Schaal, 1999], and, more recently, as *movement templates* [Kober et al., 2010], *subpolicies* [End et al., 2017] and *motor programs* [Shankar et al., 2020].

²Primitive structures in muscular activation patterns are generally known as Muscle Synergy (MS) [Zandvoort et al., 2022]

2016; Poeppel & Idsardi, 2011; Stiles & Shimojo, 2015b] and linguistics [Endress et al., 2007, 2009]. Notwithstanding, the use of such PPs in robotic control is, to this date, lacking. In a broader sense, despite decades of research in cognitive science, only relatively recently did engineering started to extend the role of perceptual information beyond low-level feedback to higher-level *grounded* cognitive architectures and reasoning frameworks [Kotseruba & Tsotsos, 2018; Vernon, 2022], namely through different *emergent* approaches to artificial cognition [Beetz et al., 2010; Crawford et al., 2016; Kumar, 2021; Raman & Maybank, 2016; Rothfuss et al., 2018; Sun, 2007; Verschure & Althaus, 2003; Yang et al., 2022]. Under this emergent view of cognition, *cognitive processes are deeply rooted in the real-world interaction* [Wilson, 2002], and consequently, perception - regardless of modality - is a crucial component of any cognitive architecture. By extension, its integration into episodic and procedural memory models - inherently tied to experience - is equally important [Kotseruba & Tsotsos, 2018; Tulving, 1984; Vernon, 2022]. In practice, not only can motion planning be enhanced by persistent memory [Rothfuss et al., 2018], but the ability to avoid long-term forgetting of learned skills has been described as a target feature of autonomous artificial machines [Kudithipudi et al., 2022]. Nonetheless, existing cognitive frameworks are heavily oriented towards limited or specific perceptual inputs (e.g. vision data) - generally processed in isolation and with *off-the-shelf* solutions - or tasks (e.g. human-robot interaction or activity recognition), or otherwise exploit a *priori* problem-specific assumptions [Kotseruba & Tsotsos, 2018; Vernon, 2022]; the underlying structure is generally either opaque or overly focused on abstract, *symbolic* or *semantic* declarative representations [Kotseruba & Tsotsos, 2018; Kumar, 2021; Rothfuss et al., 2018]; and lastly, not only has episodic memory a reduced role in these architectures, they fail to conceptually integrate crossmodal interaction e.g. shared PPs into the model [Kotseruba & Tsotsos, 2018; Vernon, 2022]. Modern DL-based alternatives, on the other hand, still have difficulty scaling outside simulated environments [Sünderhauf et al., 2018; Vernon, 2022]. Conversely, modern findings have disputed a purely abstract and amodal approach to generalized knowledge representation, highlighting instead the *grounded* and *compositional* nature of human cognition, whereby sensorimotor patterns and contextual information gain additional relevance, and simpler elements (*i.e.* PPs) provide building blocks for richer semantics [Akkaladevi et al., 2021; Barsalou, 2016; Kumar, 2021].

In the scope of robotic control, these notions can be explored through associative memories holding proprioceptive information in the form of both motor and perceptual patterns, while also modeling stochasticity [André et al., 2016b; Knowlton & Schorn, 2024; Pastor et al., 2013, 2012; Willingham, 1998]. In episodic *stereotypical* experiences, these structures couple perception and action, and facilitate both outcome prediction (e.g. failure detection) and high-level generalization of motor skills to similar contexts [André et al., 2016b; Pastor et al., 2013, 2012; Tulving, 1984; Vernon, 2022]. Notably, repurposing existing MP-based frameworks for perceptual representation would allow structuring cognitive models around shared primitive elements, and potentially lead to richer, more informative state descriptors for contextual learning [Klink et al., 2020; Kober, 2014; Kupcsik et al., 2013]. Nevertheless, despite several MP applications in spatial pattern recognition with humanoid robots [Ijspeert et al., 2002, 2003; Pastor et al., 2013; Samant et al., 2016], such primitives in robotic perceptual patterns remain, for the most part, an unexplored concept.

Analogous strategies have gathered practical interest beyond direct application in movement generation or high-level cognitive strategies; in particular, human behavior estimation, when performed through detailed neuromuscular models or derived surrogates [Eskinazi & Fregly, 2018; Messeri et al., 2022; Peternel et al., 2019], has been shown to be greatly facilitated by exploiting the *synergistic* nature of vertebrate Musculoskeletal (MSK) systems [Chen et al., 2019; Meyer et al., 2016; Qiao et al., 2022; Rückert & d’Avella, 2013], which are known to self-organize through primitive-like functional groupings [d’Avella et al., 2006; Zandvoort et al., 2022]; this is appealing, for instance, in the case of collaborative robotics or in shared environments, where recognizing contextual cues or behavioral changes in human agents (e.g. task-related activity, physical impairments, muscular fatigue) can prompt robots to autonomously adjust their behavior accordingly [Akkaladevi et al., 2021; Oguz et al., 2017; Yang et al., 2022], going beyond passive self-adaptation to the environment, and proactively supporting human effort for deeper collaboration and better societal integration.

Taking into account this emergent role of primitive theory in the design and control of artificial machines and automated systems, the present thesis presents multiple studies on primitive-based applications, focused on - but not limited to - locomotive and industrial robotics. From the perspective of primitive elements as a valuable instrument to bypass complexity and redundancy, not only in high-dimensional robots but also in the analysis of human behavior, it proposes general solutions that not only validate their relevance, but ultimately aid in the development of autonomous systems capable of continual self-improvement [Kudithipudi et al., 2022]. In a multidisciplinary approach, primitives are initially 1) exploited by a generic trajectory-based controller, which is parametrized to replicate bipedal locomotion in different platforms by IL; secondly, the same primitive structure is 2) repurposed for a generic and episodic perceptual memory model, which is used in a gait monitoring application with human subjects; and lastly, the 3) timely detection of muscle fatigue (a documented challenge in industrial settings [Bevan, 2015; de Kok et al., 2019]) is investigated in upper limb repetitive movements through cumulative changes in neuromuscular primitives.

1.1 Motivation and Objectives

While the high-functioning cognitive capabilities of modern Large Language Models (LLMs) constitute impressive milestones in computer science and systems engineering, the construction of a machine capable of seamless self-improvement over its lifespan remains an aspiration for the future [Silver et al., 2021; Sutton, 2019]. Approaching the subject from a bioinspired perspective, Kudithipudi et al. [2022] briefly summarize the key features in autonomous systems capable of Lifelong Learning (L2)³, mapping them to several documented mechanisms of biological organisms; notably, the authors list 1) *transfer and adaptation*; 2) *exploiting task similarity*; 3) *overcoming catastrophic forgetting*; 4) *performing task-agnostic learning*; and exhibiting both 5) *noise tolerance* and 6) *resource efficiency* as hard requirements for machines to match the level of behavioral adaptability found in nature. Within this paradigm, this work argues

³Also referred to as *continual* or *incremental* learning [Kudithipudi et al., 2022; Saveriano et al., 2021; Thrun, 1995].

that integrating primitive elements into the design different control strategies, whether at motor generation or perceptual level, can - directly or indirectly - help in artificially reproducing each of these faculties.

Furthermore, despite accurate and precise robotic movement often being the product of detailed kinematic and dynamic models imbued into the control process at the cost of simplicity [Reher & Ames, 2021], generic solutions have, as broadly noted by Sutton [2019], historically outperformed sophisticated hand-tailored algorithms when leveraging ML over task-oriented planning. As such, *generality* and *modularity* occupy a central role in this work: instead of highly performant *cutting edge* approaches to different scientific and engineering challenges, generic primitive-based alternatives where few or no task- or platform-specific assumptions are made (e.g. locomotive strategy, sensor hardware) are introduced, relying on ML for systematic parameterization within the bounds of the problems under study. Notwithstanding the focus on rhythmic movements, and legged locomotion in particular, prompted both by its relevance and pervasiveness in human mobility and humanoid robotics, as well as its appeal as a valuable testbed for control strategies [Ijspeert, 2014], it is the general aspect of the documented methods that is highlighted, with potential applications in domestic, clinical and industrial settings. In order to bring about such general solutions to the problems tackled in this thesis, the following practical objectives are proposed:

- 01** *Learn model-based trajectories from observed behaviors by exploring combinations of primitives;*
- 02** *Provide a common representation for movement generation and recognition, and implement primitive libraries;*
- 03** *Generalize locomotion behavior to different anthropomorphic robots with ML methods;*
- 04** *Construct episodic skill memories by storing sensor patterns in a compact and generalizable representation;*
- 05** *Establish contextual relationships by means of shared perceptual primitives;*
- 06** *Identify high-level contextual alterations by means of objective primitive-related features or markers;*

In practice, the realization of these goals implies meeting the referenced practical requirements for L2 systems; to that end, three distinct scientific studies are presented that exploit primitive structures in motor control, perceptual patterns, and muscular recruitment, where the following Research Questions (RQs) are, respectively, explored:

- RQ1** *How can a generic primitive-based controller generate functional locomotion patterns in different platforms?*
- RQ2** *How can dynamical systems encode stereotypical episodic information associated with different perceptual contexts, while accounting for primitive elements?*
- RQ3** *What cumulative effect do low-force repetitive movements have in the activation profile of muscle synergies?*

Answering these RQs not only provides diverse scientific contributions, including practical, problem-oriented methodologies with appeal not limited to engineering and robotics, but also highlights - at a conceptual level - the valuable role primitive elements can have in the development of robust, adaptive

and intelligent machines and systems.

1.2 Contributions

This work builds upon existing research, borrowing inspiration from different scientific fields while contributing to the research effort through theoretical concepts, practical algorithms or otherwise novel applications of existing tools to new challenges, namely:

- **Generic primitive-based trajectory control strategy** defined on top of multi-level motor libraries and capable of replicating stable biped-locomotion in different simulated high-dimensional 3D platforms, while integrating low-level feedback mechanisms into the control loop;
- **Imitation Learning strategy for bipedal locomotion**, employing Black Box Optimization (BBO) methods in a Behavior Cloning (BC) approach from reference kinematic patterns, which can be fostered for generalized skill learning and generalization, demonstrated through the transfer of functional walking between platforms and locomotion modalities, as well as memory acquisition;
- **Generic episodic memory model** encoding stereotypical perceptual data (discrete or rhythmic) in a compact, data-driven, amodal, time-invariant representation, modeling stochasticity in data patterns and accounting for shared perceptual primitives across sensor sources and tasks;
- **Context identification and tracking algorithms** derived from the episodic perceptual model, and which are used to identify gait type and failure scenarios, such as **slip detection** with reasonable accuracy, in a prerecorded experimental dataset of human walking;
- **Metrics for muscle fatigue detection**, objectively quantifying the extent to which primitive elements are affected over time, validated through functional and statistical analysis of repetitive low-force non-isometric human movement;
- **Experimental kinematic and electromyographic dataset** of human participants in multiple uni-articular and composite exercises until self-perceived exertion, as an outcome of fatigue analysis, which can be used to comprehensively study the development of muscle fatigue in healthy individuals with low-force tasks over time.

Beyond these contributions, it is relevant to note that the methods described in this thesis, in multidisciplinary fashion, focus on different types of motions, from legged locomotion to upper limb movement, both in simulation and real experiments, as well as robots and human subjects, while tackling different scientific and engineering challenges; this offers a unifying perspective about the relevance of primitive elements across different scientific branches, and how they can be harnessed in the development of intelligent autonomous machines.




Figure 2: Conceptual overview of the present work; chapters 2, 3 and 4 constitute parallel independent studies on different applications of primitive theory, contextualized by Chapter 1 and with relevant outcomes summarized in Chapter 5; practical objectives (1.1) accomplished in each chapter are highlighted (☐).

1.3 Outline

The relevant scientific contributions of this work are achieved through a set of three separate studies that exploit primitive structures in motor control, perceptual patterns, and muscular recruitment, for the generation and analysis of robotic and human movement, while tackling practical issues in different areas. Accordingly, Chapters 2, 3 and 4 can be read separately, with each presenting specific motivations and formulating different Research Questions (RQs) that drive the research effort. Figure 2 presents the conceptual outline of this thesis, along with how each contributes to achieving objectives **01** to **06**.

The present chapter, 1: Introduction, provided a brief introduction to the overarching notions and practical issues explored within this work; it contextualizes the need for generic and straightforward methods for robotic control and in the design of intelligent automated systems, with generality and adaptability in mind, and highlights motor and perceptual primitives as an appealing instrument to achieve them.

The following chapter, 2: Generic Primitive-based Legged Control, provides an historical overview of significant advances in robotic locomotion, and describes the development of a generic and adaptable MP-based robotic controller, which is applied to legged control of humanoid robots, achieving functional closed-loop locomotion patterns sharing MPs across different DoFs. Functional parameterization and generalization is conducted using ML methods using demonstrative examples from bioinspired robotic locomotion, in a IL scheme. The controller is validated and benchmarked in different simulated legged platforms.

In Chapter 3: Dynamical Perceptual Memory Model, an episodic memory model base on dynamical systems is proposed, repurposing the primitive mixing strategy employed in Chapter 2 and accounting for stochasticity in sensor data. Practical algorithms for memory acquisition, context identification and failure detection are proposed and validated in an experimental dataset acquired from human walking trials. An objective comparison with alternative bioinspired approaches to slip detection is presented.

Chapter 4: Impact of Muscle Fatigue in Musculoskeletal Primitives, moves away from bipedal locomotion to rhythmic upper limb movements, through the analysis of primitive elements in human muscular

recruitment patterns, widely known as Muscle Synergy (MS). Constructing an experimental dataset with simple and composite motions of the right arm repeated until full exertion, different objective indicators are studied - by means of detailed functional and statistical analysis - regarding their ability to detect muscular fatigue in primitive patterns, providing insights into the compositional nature of muscular behavior.

Lastly, Chapter 5: Conclusion, scrutinizes the findings, contributions and major takeaways presented in the previous chapters, with focus on the achievement of the proposed objectives and discussing how each contributes to meet the practical requirements of L2 systems (1.1). A broad general discussion is presented, emphasizing the role primitive-based control and assessment can have on the path towards truly autonomous systems in the real world, while approaching each from a robotic viewpoint and highlighting the contributions in, as well as beyond, the field of robotics.

Generic Primitive-based Legged Control

This chapter describes the development and validation of a generic bioinspired controller within the context of bipedal locomotion in legged robots, employing dynamical systems and Motor Primitives (MPs) to achieve a parametric low-dimensional trajectory representation. It contextualizes walking as a complex skill that requires, by its very nature, the ability to generalize and adapt motor policies between different environments, and which can benefit from bioinspired notions for generality and dimensionality reduction. Using previously validated gait patterns as reference, Machine Learning (ML) is used to teach locomotive patterns to two simulated humanoid robots using a shared MP library. The generic nature of the implemented control is highlighted, with learned open-loop joint trajectories requiring no locomotion-specific knowledge or model, leveraging instead *learning* to achieve functional gait; instead, different gait-specific feedback mechanisms provide added robustness and stability in the movement.

In practice, objectives **01**, **02** (partially) and **03** (1.1) are accomplished. Reproducing walking behaviors with $>90\%$ similarity to reference demonstrations is seen as a Key Performance Indicator (KPI) of the proposed approach.

2.1 Introduction

Automated human-like locomotion has been pursued even before the word *robot* was first introduced by Karel Čapek, going back to the late 15th century and Da Vinci's *Automa Cavaliere*, a pulley/cable-operated mechanical humanoid automaton. Since then, artificial walking machines have, much like other branches of robotics, evolved from conservative over-actuated solutions offering rudimentary walking patterns to near-passive robust and versatile controllers generating fluid human-like gait [Reher & Ames, 2021].

In truth, there is an endless untapped potential for autonomous robotics in our world, in the most varied scenarios, from health care to service industries, within our homes or in outer-space [Kroemer et al., 2020]. On the other hand, as exposed in Chapter 1, for productive interaction and collaboration between automated systems (e.g. robots) and humans, whether in work or domestic environments, control strategies with autonomy, versatility and adaptability beyond those of current robotics need to be developed [Peters et al., 2011; Sigaud & Stulp, 2019], ultimately paving the way towards ever-learning

artificial intelligent and self-sufficient machines [Kudithipudi et al., 2022]. This is all the more evident - and necessary - in walking robots, whose increasingly complex tasks often place them in treacherous and challenging environments which they are required to safely navigate while complying with task demands; in practice, legged robots are particularly suited to navigate rough and uneven terrains as they can afford limited ground contact (*i.e.* minimal foothold requirements) over dangerous or unstable terrains [Reher & Ames, 2021]. Overall, this represents a significant increase in versatility and overall maneuverability - in contrast with conventional wheeled platforms - useful in a large number of applications; when considering the dynamic nature of the real world and the ever-changing features it displays, reliable and efficient ways to locomote are of special interest in robotic research and industry [Ijspeert, 2014; Reher & Ames, 2021].

On the other hand, implementing legged locomotion is often challenging as it 1) requires dynamic interaction between multiple elements, including the environment, and is 2) high-dimensional, 3) highly nonlinear and 4) multidisciplinary [Ijspeert, 2014; Reher et al., 2020]. Nonetheless, locomotion has long gathered the interest of robotics, biomechanics and neuroscience, with artificially-generated walking approached from different scientific and engineering perspectives, including purely mathematical and robotic methods based on different notions of stability and bioinspired methodologies aiming to replicate natural gaits, eventually leading to modern cutting edge methods with impressive walking capabilities, exploiting notions such as dynamic stability, passive dynamics, reduced-order models, and state-of-the-art Reinforcement Learning (RL) [Apgar et al., 2018; Faraji et al., 2014; Ijspeert, 2014; Niu et al., 2024; Peng et al., 2020; Reher & Ames, 2021; Smaldone et al., 2022].

This chapter provides an alternative strategy to achieve functional gait in legged robots: instead of pursuing highly-performant and efficient locomotion, a generic control structure is proposed that requires little to no knowledge of the robotic platform, but is versatile enough to be taught, with automated ML tools, walking patterns from preset behaviors or demonstrations. Building upon the underlying idea explored in this thesis, the reduced dimensionality of high-level primitive-based control is exploited, discarding the need for detailed kinematic or dynamic models commonly required in robot walking [Ijspeert, 2014; Reher & Ames, 2021; Reher et al., 2020], simplifying generalization of robotic walking to different humanoid robots; this is in line with the modern paradigm of harnessing *learning* solutions over integrating task-oriented and human-sourced *knowledge* into the control strategy [Silver et al., 2021; Sutton, 2019], and provides an answer to the following Research Question (RQ):

RQ1: How can a generic primitive-based controller generate functional locomotion patterns in different platforms?

To that end, Objectives **01**, **02** and **03** (1.1) are pursued. Relevant scientific contributions (1.2) include a 1) *generic primitive-based trajectory control modality* and a 2) *practical IL strategy for bipedal locomotion*. This chapter is organized as follows: Section 2.2 provides a brief summary of the relevant biological principles behind human and animal locomotion, as well as an historical review of major milestones in robotic walking and bioinspired control. In Section 2.3, the generic control strategy developed in this work is extensively described from a general and task-agnostic perspective, including practical ap-

proaches to motor skill learning. A practical application of this general strategy is presented in Section 2.4, wherein closed-loop control of lower limbs in humanoid robots is studied by teaching functional locomotion based on Central Pattern Generator (CPG) to two different legged platforms, which is extensively profiled in Section 2.5, with particular emphasis on key performance indicators. A discussion on the achieved results is provided in Section 2.6, highlighting the main features-advantages and drawbacks of the proposed approach, with focus on the role of MPs in movement generation. Closing the chapter, Section 2.7 highlights relevant insights gained from this study, including how they contribute to the general goals and objectives of this dissertation, and possible subsequent lines of research that can build upon it in the short term.

2.2 Background

2.2.1 A Biological Perspective on Walking

From a biological perspective, while, at first glance, some types of gait may seem simple and straightforward, studies in neuroscience have shown that locomotion are the result of a complex network wherein the central and peripheral nervous systems, as well as their musculoskeletal counterpart, constantly interact, in the form of coordinated muscular activation and articulation around different Degrees-of-Freedom (DoFs) [Dzeladini et al., 2019; Ijspeert, 2014]. In particular, two core concepts are hypothesized to be the key to understanding animal locomotion, 1) rhythmic activation patterns originating on the spinal cord; and 2) discrete forces resulting from interaction with the environment [Geyer & Herr, 2010; Hogan & Sternad, 2013; Ijspeert, 2014]. The interplay between these notions has been observed across different modes of locomotion, from swimming to legged walking, highlighting the *hybrid* nature of the process [Reher & Ames, 2021].

For a long time, following experiments on transected and deafferented animals wherein regular stepping motions were observed both *with* and *without* perceptual input, the existence of localized neural circuits where periodic patterns emerged, in a self-sustained and semi-autonomous fashion, was hypothesized [Dzeladini et al., 2019; Hooper, 2001]. The major breakthrough behind this idea is precisely in this localized nature, that bypasses top-level cognitive structures like the cerebellum, and suggests the presence of these patterns even in the absence of high-level control and sensor feedback [Arena et al., 2004; Buchli & Ijspeert, 2008; Geyer & Herr, 2010; Orlovsky et al., 1999]. In contrast, concurrent research, focusing on the role that peripheral proprioceptors have in the generation of these motor patterns, postulated that gait-specific limit cycle behaviors arose naturally from perceptual reflex mechanisms, discarding central neural pattern generation [Dzeladini et al., 2019; Geyer & Herr, 2010]. Modern studies in neuroscience and biology, however, suggest a hierarchical structure wherein the contribution of both these aspects are relevant. On one hand, the genesis of rhythmic muscle activation patterns is hypothesized to occur in the spinal cord by Central Pattern Generators (CPGs), as a response to either reflexive mechanisms or high-level cognitive processes [Grillner & El Manira, 2020; Pearson et al., 2006]. CPGs produce self-sustained rhythmic activation patterns forwarded - in a subject- and movement-specific syner-

gistic structure - to the musculoskeletal system¹, notably without the need for any kind rhythmic high-level input or sensory feedback [Dzeladini et al., 2019; Geyer & Herr, 2010; Ijspeert, 2014]. While there is still active debate as to what extent these structures are present or relevant in primates - whose locomotion holds specific particularities when compared to simpler musculoskeletal mechanisms, such as the need for upright posture - which add complexity to the underlying control mechanisms, it has been shown that pattern generation does indeed occur locally in spinal cord [Dimitrijevic et al., 1998]. On the other hand, sensor input is no less significant to this process as it helps shape and modulate the output of these rhythmic patterns through *spinal reflexes*, which establish a bridge between proprioceptor feedback and leg activation mechanics through alpha motoneurons, circumventing central cognitive centers [Geyer & Herr, 2010]. This is, for instance, supported by scientific evidence that highlights the relevance of positive force feedback in extensor muscles in load bearing and compliant behavior [Prochazka et al., 1997].

2.2.2 From Static to Dynamic walking

Early legged robots approached the problem of locomotion from a static perspective, looking at the movement as a sequence of motionless postures and ensuring the classical notion of stability *i.e.* keeping the projection of the Center of Mass (CoM) in the ground plane within the convex hull formed by the outermost ground contacts or the support base [Kajita & Espiau, 2008]. While successful at achieving rudimentary locomotion, these methods were introduced on the assumption that velocities and accelerations were small, and as such CoM displacement was very constrained and resulting gait relatively slow.

In order to escape this conservative interpretation of gait stability, alternative locomotion strategies were developed which employed the concept of *dynamic* stability *i.e.* the notion of stable behavior over time, while accounting for the body's momentum and inertia [Ijspeert, 2014]. This spans the family of methods known as *dynamic walking*, which comprises current modern high-performant walking strategies [Reher & Ames, 2021]. The main insight behind this idea is to extend the classical geometric notion of static stability by taking into account the system's dynamic wrench resulting from the gravitational and inertial forces acting upon it. Under the assumption that posture changes on robot are only possible through ground contacts, dynamic balance is achieved when the forces at the contact surface nullify gravity-inertia forces felt by the robot [Kajita & Espiau, 2008; Sardain & Bessonnet, 2004]. This entails identifying the coordinates of the so-called Zero-Moment Point (ZMP), using the Newton-Euler equations of motion (from multi-body dynamics) and screw theory *i.e.* the point where all tangential forces nullify each other, and therefore no lateral momentum or *tilt* is felt by the robot [Reher & Ames, 2021; Vukobratović & Stepanenko, 1972]. The ZMP is closely related to the Center of Pressure (CoP), with the notable caveat they are derived from different aspects of locomotion, with the CoP defined locally in relation to a contact surface, as the point *within* said surface where the resulting interaction forces can be virtually replaced by a normal or tangential component without any friction. During single support phase (single ground contact), CoP and ZMP are essentially the same [Sardain & Bessonnet, 2004]. In practice, while a quasi-static measure in theory, the ZMP represents a stability criterion with a broader scope and applicability than conventional

¹In the case of legged locomotion, flexor and extensor muscles in the legs.

static stability analysis; constituting a major breakthrough in the field of legged robotics, it sat at the center of research for several decades [He et al., 2019; Kajita & Espiau, 2008; Reher & Ames, 2021], as it accommodates a greater range of movements than previously achieved by legged robots at its introduction, allowing faster and more flexible gaits [Reher & Ames, 2021].

From a control perspective, conventional ZMP-based locomotion (Figure 3) commonly employs reduced-order models, such as Linear Inverted Pendulum (LIP), to map target CoM coordinates that ensure stability, relying on Inverse Kinematics (IK) or Inverse Dynamics (ID) to generate control inputs; this was implemented, for instance, on Honda’s ASIMO [Chestnutt et al., 2005] and Kawada’s HRP [Kaneko et al., 2002] humanoids. Notwithstanding, the ZMP is currently seen as overly restrictive within the context of modern dynamic walking, since 1) estimating ZMP coordinates implies a support polygon, and is therefore not applicable to robots with line feet [Vukobratović & Borovac, 2004] such as Agility Robotics’ Cassie [Gong et al., 2019]; in addition, it assumes flat-footed and quasi-static contacts with the floor [Reher et al., 2020], in direct contrast to known properties of human gait *i.e.* multiphasic, underactuated and, as exposed in the previous section, heavily reliant on impact *i.e.* reflexes [Dzeladini et al., 2019; Geyer & Herr, 2010; Sutherland, 1994]; 2) conventional ZMP-based locomotion is not robust against perturbations to its movement nor accounts for recovery behaviors in unstable conditions [Collins et al., 2005; Van der Noot & Barrea, 2014]; 3) the need for a detailed kinematic and dynamic model of the platform is often associated with extensive parameterization, hindering the generality of the approach; on one hand, simple reduced-order models such as LIP assume a constant CoM height, which is both unintuitive from a biological standpoint and restrictive performance-wise; on the other, more detailed or realistic models may overcome this limitation may add complexity *e.g.* nonlinearity [Reher & Ames, 2021; Seyde et al., 2018]. And finally, 4) walking is energetically inefficient due to overactuation, as these methods typically require full controllability *i.e.* continuously ensuring dynamic stability through the walking process, which has been shown to not be a hard requirement for stable locomotion [Dallali, 2012; Van der Noot & Barrea, 2014]. As a result, current robotic locomotion is no longer solely driven by ZMP dynamics; nevertheless, it remains widely used as a stability constraint complementing modern techniques [Beranek et al., 2022; He et al., 2019; Kaymak et al., 2023; Reher et al., 2020; Smaldone et al., 2022].

A complementary approach to the ZMP was later introduced in the form of the Capture Point (CP)² [Hof, 2008; Pratt et al., 2006]; initially developed as a method to handle sudden unstable conditions (*e.g.* push recovery), it eventually gained relevance in gait planning, as is illustrated in implementations in Honda’s ASIMO [Takenaka et al., 2009], the DLR Bipid [Englsberger et al., 2011, 2014] and Boston Dynamic’s Atlas [Johnson et al., 2015; Seyde et al., 2018]. In practice, a CP designates the target foot placement that halts a robot’s current forward or lateral momentum within a finite time horizon; when employed sequentially at each stepping motion, it can help ensure dynamic stability of the CoM while accommodating transient instability during walking, which is not only more in line with biological gait, but provides additional robustness against unforeseen perturbations [Reher & Ames, 2021; Van der Noot & Barrea, 2014]. On the other hand, much like conventional ZMP approaches, CP methods are often coupled with reduced-order models (*e.g.* LIP) and may be limited by their suboptimal assumptions [Englsberger et al., 2014]; moreover,

²Also referred to as *extrapolated* Center of Mass [Hof, 2008] and Divergent Component of Motion [Takenaka et al., 2009].




Figure 3: Conceptual visualization of ZMP-based locomotion. Stability is ensured as long the estimated position of the ZMP remains within the support polygon; adapted from Kajita & Espiau [2008].




Figure 4: Illustration of CP-based walking; adapted from Engelsberger et al. [2011].

the assumption of null momentum at each step can lead to higher overall energy expenditure [Reher & Ames, 2021]. While the conservation of momentum is addressed in later works that extend the definition of the CP, these methodologies are still often considered inefficient in terms of energy consumption [Reher & Ames, 2021; Seyde et al., 2018].

It was this last aspect that motivated the shift in focus of research towards the role of *passive* and *hybrid* dynamics in the locomotive processes found in nature, framing walking as an intrinsically nonlinear task where stability is not guaranteed continuously [Reher & Ames, 2021], as evidenced, for instance, in human legged gait being easily disrupted mid-swing, potentially resulting in falls.

2.2.3 The Role of Passive Body Dynamics

Despite the versatile walking capabilities of early dynamic walking solutions, as reviewed by Kajita & Espiau [2008] and partly by Zhou & Bi [2012], a significant challenge remains the significant energetic inefficiency associated with them, generally as a result of overactuation and ignoring the body's natural or

passive dynamics, which animals exploit passive dynamics to better utilize the limited energetic resources available e.g. non-elastic interaction with the terrain for kinetic energy dissipation [Hogan & Sternad, 2013], and which have been shown to be heavily ingrained in natural locomotion [Ijspeert, 2014; Reher & Ames, 2021; Renjewski & Seyfarth, 2012; Sutherland, 1994].

While approaching the problem from a biomechanical standpoint, the field of robotics delved into *passive* locomotion, namely through the design of legged robots and control strategies that exploited the passive (albeit nonlinear) dynamical properties of the system to generate movement e.g. leg swinging caused by gravity and inertia, instead of actively powered joints [Ijspeert, 2014; Reher & Ames, 2021]. This entails moving away from the conservative nature of previous methods, and looking at locomotion as a steady state process [Kajita & Espiau, 2008], which, in practice, entails favoring *global* over *local* stability measures and requirements i.e. from foot strike to foot strike [Reher & Ames, 2021; Van der Noot & Barrea, 2014]. The underlying idea, later to be dubbed as *limit cycle* walking [Hobbelen & Wisse, 2007], was demonstrated in the 80s in works by Raibert [1986] and Hurmuzlu & Moskowitz [1986] with the development of compliant machines wherein stable periodic orbits were used to achieve hopping and dynamic locomotion at high velocities, and would not only highlight the potential of reduced order models such as the Spring-Loaded Inverted Pendulum (SLIP), but also serve as inspiration for eventual fully passive designs, such as those introduced by McGeer [1990] in the early 90s, which relied solely on gravity to achieve dynamically stable locomotion (Figure 1, *center*); McGeer’s designs brought about the field of *passive dynamic walking* and were notable for approaching the problem from a dynamical systems perspective (using nonlinear systems analysis), and evaluating the orbital stability using Poincaré maps [Ijspeert, 2014; Kajita & Espiau, 2008; Reher & Ames, 2021]

Under this periodic view of stability, locomotion is seen as *controlled falling*, where transitory instability is required - and exploited - to achieve higher velocities and a reduced energetic footprint [Reher & Ames, 2021]. Relevant scientific output includes the development of *passive* dynamical walkers, reaching more energy-efficient and natural gait in artificial systems, through fully passive and/or underactuated systems (Figure 1, *right*), with Cost of Transport (CoT)³ similar to that estimated for human walking, representing a decrease in energy consumption of $\geq 90\%$ when compared to alternatives [Collins et al., 2005].

Generally, this represents a notable shift in the design of walking robots; on the other hand, it is precisely the underactuated nature of such passive implementations that poses a practical challenge, as it limits the extent to which such systems can be controlled, and in turn how adaptable and robust against unforeseen obstacles or gait perturbations the resulting gait can be [Reher & Ames, 2021]. In fact, up to the last decade, modern walking machines could be classified as either *versatile* robots able to replicate different kinds of locomotion with reasonable performance, or, alternatively, *efficient* walking machines that are able to reproduce a fluid, natural-looking gait with a decreased energetic footprint; the former often requires intricate and demanding control algorithms, while the latter is achieved through careful bioinspired morphologies, semi-passive walking control and energy-recapturing electronics [Ijspeert, 2014]. On the surface, this suggests *versatility* and *efficiency* to be competing aspects of the broader robotic walking

³Measure of energy efficiency of a mobile platform, and is defined as the ratio between power consumption and the product of weight and speed [Seok et al., 2015]. For reference, the estimated CoT of human walking is 0.2 [Ijspeert, 2014]

problem. In order to bridge this gap, researchers focused on the *hybrid* nature of gait - with both a continuous component in the form of periodic motions of the legs, and with discrete interaction with the ground at foot contact⁴ - and employed nonlinear control theory to simplify the dynamics of legged robots to low-dimensional dynamical systems through which stable behaviors can be generated. The underlying concept behind this approach, Hybrid Zero Dynamics (HZD) [Grizzle et al., 2001; Westervelt et al., 2003], was initially derived for and validated in *planar* 2D walking in the RABBIT [Chevallereau et al., 2003] and later in the MABEL [Sreenath et al., 2013] bipeds; these implementations were notably on 2D pointed foot robots - on which ZMP walking, for instance, isn't applicable. One of the main attractive features of HZD is its versatility, allowing the definition of *virtual* constraints that modulate and/or coordinate leg motions, such as integrating ZMP as a stability margin [Grizzle et al., 2014]. These constraints are akin to holonomic constraints imposed through the physical morphology of the robot, but defined at control level and enforced through linear feedback control [Reher & Ames, 2021]. For an in-depth review of the formulation behind HZD and dynamic walking methodologies in general, the reader is referred to Reher & Ames [2021]; within the context of this work, HZD is framed as a cutting edge approach to legged locomotion, overcoming the limitations of previous strategies, and capable of impressive walking performance, as demonstrated on Softbank's NAO [Ames et al., 2012], SRI's DURUS (with human-like foot placement motions⁵) [Hereid et al., 2018; Ma et al., 2017; Reher et al., 2016] and Agility Robotics' Cassie [Gong et al., 2019]. In a practical sense, HZD defines a way to plan dynamically consistent and efficient periodic motions, having been used in mid- or low-level control in conjunction with high-level optimization-based methodologies such as Model Predictive Control (MPC) and Trajectory Optimization (TO) providing reference patterns [Gong et al., 2019; Hereid et al., 2018; Nguyen & Sreenath, 2015; Reher et al., 2019].

2.2.4 Bioinspired Robotic Control

An alternative approach to implementing robotic locomotion entails exploiting bioinspired notions such as Central Pattern Generators (CPGs) to achieve limit cycle behaviors on leg joints. In contrast with methodologies such as HZD that excel in precise and efficient dynamic control of legged machines, bioinspired high-level approaches constitute a more accessible, biologically plausible, adaptable and generally computationally less intensive solution to generate locomotive patterns, without requiring the detailed knowledge of the robot's dynamics [Dzeladini et al., 2019; Reher & Ames, 2021].

From an engineering perspective, the rhythmic nature of CPG patterns can be seen as a type of limit cycle behavior [Buchli & Ijspeert, 2008], commonly implemented with synchronized oscillator-driven networks of nonlinear dynamical systems, mirroring leg symmetries, and which can be parametrized to generate different gait modalities [Arena et al., 2004; Matos & Santos, 2014]. Bio-inspired designs such as the Matsuoka half-center oscillators [Cristiano et al., 2014; Liu et al., 2008], neural models such as Recurrent Neural Networks (RNNs), and abstract mathematical approaches such as Hopf bifurcations or Van der Pol oscillators, among others [Dzeladini et al., 2019; Nachstedt et al., 2017; Watanabe et al., 2008], have been

⁴Simplified as a sequence of continuous systems whose state is reset at each foot strike [Reher & Ames, 2021]

⁵It has been shown that *multi-contact* stepping *i.e.* heel-toe foot planting has a relevant role in biological gait, reducing joint torques (increasing energy efficiency) and increasing walking velocity [Inman, 1966; Reher et al., 2020; Sutherland, 1994].

used by roboticists. In particular, abstract oscillator modalities, despite lacking biological relevance, are notably convenient regarding to ease of shape modulation and explicit phase control⁶ [Dzeladini et al., 2019]. Practical applications range from aquatic [Wang et al., 2011; Yu et al., 2016] and amphibious locomotion e.g. in salamanders [Chevallier et al., 2008] to both quadruped [Buchli & Ijspeert, 2008; Chevallier et al., 2008] and biped [Matos & Santos, 2014; Nakanishi et al., 2004] walking; additionally, the same design principles have been used beyond robotic locomotion, such control of drumming robots [Degallier et al., 2011] and ball juggling [Rizzi & Koditschek, 1994; Schaal et al., 1996]. In general, these implementations are attractive because 1) their parametric nature presents advantages regarding dimensionality reduction - valuable when considering optimization and learning strategies; 2) their modularity makes them suitable for application in distributed systems such as humanoid platforms, and implementation of gait transition mechanisms that facilitate multimodal locomotion; 3) limit cycle design allow dynamical coupling or mechanical entrainment with the robotic platform, as well as robustness against perturbations due to the attractor dynamics; and lastly, as mentioned, 4) minimal knowledge about the robot is required [Dzeladini et al., 2019]. On the other hand, these applications are mostly on stiff, fully actuated robots, while, as discussed, in nature most musculoskeletal systems exhibit a significant degree of complacency, influenced by their passive dynamics, and which have a relevant role in the energetic efficiency of the gait [Arena et al., 2004; Buchli & Ijspeert, 2008]. Moreover, CPG is inherently an open-loop concept and practical applications commonly require task-oriented feedback mechanism for closed-loop control [Matos & Santos, 2014; Nassour et al., 2014].

Another relevant bioinspired notion employed in robotics is that of Motor Primitives (MPs). As briefly introduced in Chapter 1, MP theory emerged from observations of biological organisms coping with the complexity of the real world by breaking down processes into simpler components, and can be traced back to the 50s in works by Lashley and Keele [Rosenbaum et al., 2007; Shankar et al., 2020]. Under this paradigm, unique primitives represent elemental motions in their most basic form, constituting building blocks that can be modulated, combined and sequenced to generate a diverse collection of composite movements; assuming that these primitives can be shared and reused in different tasks e.g. running and walking [Barroso, 2016; Saito et al., 2018], the whole motor repertoire can be encoded through minimal set of MPs. This represents a significant simplification strategy in behavioral control, which helps bypass the natural complexity and redundancy of Musculoskeletal (MSK) systems, and which can be exploited by roboticists when designing similarly versatile and robust robots: adaptive controllers can benefit from a modular primitive scheme where motor policies result from the combination of simpler functional units. Beyond decoupling low-level control from the high-level planning stage [Degallier et al., 2011], this has the advantage of simplifying adaptation by sharing primitives across different environments, behaviors, platforms and ultimately, between different DoFs in the same agent. Ultimately, from an adaptation point of view, finding new skills or adjusting current motions in response to environmental stimuli can amount to exploring new combinations of individual primitives instead of fully re-parametrizing behaviors from scratch.

General approaches to primitive representation - *i.e.* rigorous mathematical formulations of arbitrarily-

⁶Under such a control scheme, the periodic patterns for each joint are anchored into a phase variable that can be freely manipulated, providing greater flexibility regarding the generated trajectories.

defined control patterns - have been suggested in the literature. A noteworthy example are Dynamical Movement Primitives (DMPs) [Ijspeert et al., 2003; Schaal et al., 2005], which build upon contributions from nonlinear dynamical systems theory and optimal control to achieve a parametric, stable and versatile framework for robotic control. Employing well-structured attractor landscapes to encode both discrete and periodic trajectories, and displaying properties such as parameter linearity, numerical tractability and time invariance, DMPs particularly suited are for control applications where movement patterns require online adjustment [Saveriano et al., 2021] and their parametric nature, much like with CPGs, is particularly attractive in ML [Li et al., 2018; Neumann et al., 2014; Stulp & Sigaud, 2013]. DMPs have been extensively used in robotic applications, constituting a widely popular method for motor policy representation [Chatzilygeroudis et al., 2020; Saveriano et al., 2021], with applications in object manipulation [Ginesi et al., 2021; Kober et al., 2012; Pastor et al., 2009; Samant et al., 2016], bimanual coordination [Gams et al., 2014], and, particularly relevant in this work, legged locomotion [Nakanishi et al., 2004; Rückert & d’Avella, 2013; Wang & Wang, 2020]; modular approaches where individual primitives are combined to generate control patterns are also highlighted [Deniša et al., 2013; Mülling et al., 2013; Pastor et al., 2009]. Furthermore, the notable versatility of DMPs has led to several extended formulations addressing particular control requirements, such as Compliant Movement Primitives (CMPs) [Deniša et al., 2016; Petrič et al., 2018, 2014], targeting compliant behavior by encoding kinematic trajectories with the associated torque dynamics; Probabilistic Movement Primitives (ProMPs) [Paraschos et al., 2013], integrating stochasticity into the representation, bringing forth several advantages from probability theory, such as co-activation and conditioning; and modern variants employing Deep Neural Networks (DNNs) in DMP generation [Li et al., 2023; Pahič et al., 2020; Pervez et al., 2017]. For a comprehensive overview of theoretical foundations, practical applications, strengths and shortcomings, the reader is referred to Saveriano et al. [2021].

In practice, CPGs implementations can be seen as a case of rhythmic DMPs [Degallier et al., 2011; Dzeladini et al., 2019; Ijspeert et al., 2013; Nakanishi et al., 2004; Peters & Schaal, 2008]⁷, often discretizing and compounding sub-motions as primitive elements [Matos & Santos, 2012; Zaier & Kanda, 2007]. Alternatively, Hidden Markov Models (HMMs) and Gaussian Mixture Models (GMMs) have also been proposed as a trajectory representation tool for robotic control [Calinon et al., 2010, 2012, 2011; Guan et al., 2021; Guenter et al., 2007; Khansari-Zadeh & Billard, 2011; Rozo et al., 2016; Samant et al., 2016], leveraging robust optimization techniques such as Locally Weighted Regression (LWR) and Gaussian Mixture Regression (GMR) for parameterization. Reportedly, however, both HMM and GMM lack scalability to high-dimensional systems [Mülling et al., 2013].

2.2.5 Related Work

As a hallmark methodology in the representation of motor policies, DMPs - along with varied extensions to the classical framework - have exhibited proficiency in and outside a research context. Beyond the works referenced so far, notable applications include varied object manipulation and bimanual coordination tasks e.g. pancake-flipping [Kormushev et al., 2010], sweeping [Pervez et al., 2017], door and drawer

⁷One can look at CPGs as a conceptual notion while DMPs provide the mathematical tools to implement CPGs.

opening [Magyar et al., 2019], drawing and writing [Ijspeert et al., 2013; Kulvicius et al., 2012; Lioutikov et al., 2017; Pahič et al., 2020; Zhou et al., 2019], pick-and-place [Deniša et al., 2016; Magyar et al., 2019; Stulp et al., 2012], grasping [Forte et al., 2012; Li et al., 2018; Pastor et al., 2009] and collaborative assembly [Gams et al., 2014; Huang et al., 2019]; sports-specific motions e.g. dart throwing [Kober et al., 2011], ball paddling [Kober & Peters, 2009], and playing pool [Pastor et al., 2011a], hockey [Kupcsik et al., 2013] and table tennis [Calinon et al., 2010; Mülling et al., 2013, 2010; Samant et al., 2016]; generic reaching [Gams et al., 2014; Koutras & Doulgeri, 2020; Peternel et al., 2016; QueiBer et al., 2018; QueiBer & Steil, 2018; Weitschat & Aschemann, 2018]; rhythmic motions e.g. drumming [Calinon et al., 2010; Deniša et al., 2016; Matsubara et al., 2011; Ude et al., 2010]; path planning [Ginesi et al., 2021; Lee et al., 2018; Mei et al., 2017]; movement recognition [Ijspeert et al., 2013, 2003; Samant et al., 2016]; and, lastly, in legged locomotion [André et al., 2015; Hogan & Sternad, 2013; Nakanishi et al., 2004; Nunes et al., 2020; Rückert & d’Avella, 2013; Teixeira et al., 2014; Wang & Wang, 2020]. Across the extensive literature on this particular subject matter, a heavy focus on IL from human and/or kinesthetic demonstrations, and the prevalence of articulated robotic arms as robotic agents constitute common trends. By contrast, research where 1) behavior or motor databases are constructed and 2) individual MPs - with a noted interest in DMPs - are combined during motor generation, is deemed particularly relevant within the context of this thesis. While researchers often pursue these goals within the context of RL [Guan et al., 2021], the focus herein is primarily on the structural organization of motor policies.

Prior to the original DMP formulation, Schaal [1999] discussed the role of MPs and behavior libraries within the context of bioinspired IL; building on top of some of Schaal’s ideas, Pastor et al. [2009] constructed a simple MP database containing different basic motions (e.g. *reaching*, *grasping*, *releasing*) learned from human demonstrations and generated composite behaviors by sequencing multiple MPs, modulating them with task-specific metaparameters⁸ and employing perceptual input for robust closed-loop obstacle avoidance. Similarly, Stulp et al. [2012] and Kulvicius et al. [2012] also focused on sequencing or *joining* different MPs; the first employed RL in reaching and grasping motions, while the latter extended the classical DMP formulation and used IL to learn handwriting, codifying single letters as separate discrete DMPs. Nevertheless, in line with the *competing* aspect described by Schaal [1999], in these works each MP encodes a single functional motion, and their combination is limited to sequential chaining.

From a RL perspective, Mülling et al. [2010] introduced the notion of Mixture of Motor Primitives (MoMPs), later extended in Mülling et al. [2013], wherein a gating network combined the outputs of multiple DMP systems in response to external stimuli, validated in hitting motions in table tennis with an anthropomorphic robot arm. An MP library, constructed from kinesthetic demonstrations, refined with RL, and storing perceptual *augmented* states, provided basic striking templates whose contribution was proportional to the perceptual similarity and cumulative success rates. On the other hand, sequential combination of MPs is not addressed, and the resulting mixture is defined as a linear combination of second-order dynamics⁹; moreover, these are limited to policies linear in the parameters [Deisenroth et al., 2014], and require significant *a priori* knowledge about the task [Kupcsik et al., 2017]. Notwithstanding, the same

⁸Start and goal positions, obtained/inferred from external vision systems, and time scaling.

⁹In contrast with the time-independent forcing-term element f within DMP dynamics (2.3.1).

underlying idea was later adopted by Gomez-Gonzalez et al. [2016] and Samant et al. [2016], with the latter using an alternative piecewise linear canonical system; and extended by End et al. [2017] and Celik et al. [2022], within the bounds of hierarchical *contextual* RL and stochastic motor policies, validated in simulated reaching motions in addition to the original table tennis task, respectively. In practice, MoMPs represented an early approach into effective perceptuomotor coupling and contextual adaption of motor control patterns with DMPs [Kober et al., 2011; Kober & Peters, 2011]¹⁰.

In a more generic approach, Rückert & d’Avella [2013] bridged the classical DMP formulation to naturally occurring primitive structures in animal MSK systems, developing the idea of *DMPSynergies* than could be linearly scaled and phase-shifted, in a more flexible combination scheme than the previous alternatives *i.e.* accommodating for both superposed and sequential MPs. Despite the authors only validating the proposed methodology in reaching motions and 2D walking in simulated platforms, as well as detailed neuromuscular modeling, a similar approach would be followed by Krug & Dimitrov [2015] and Nunes et al. [2020] with real platforms, namely in the control of grasping motions with anthropomorphic robotic arm, and in gait rehabilitation with a lower limb exoskeleton, respectively. However, these works did not explore MP sharing between different platforms and tasks, or formalize the construction and update of motor databases holding shared MPs and their higher-level combinations.

Inspired by works from Yamane et al. [2013, 2011] wherein binary trees and transition graphs were used to construct a general-purpose database describing human movements from a motion capture dataset, Deniša et al. [2013] delved into compliant control by exploring Compliant Movement Primitives encoding coupled kinematic and dynamic trajectories, and constructing a similar hierarchical and graph-inspired motor library with supervised kinesthetic demonstrations. The practical search efficiency of the underlying structure was harnessed when extrapolating to new behaviors, synthesizing new control patterns by sequencing useful states and segments found in the library. A notable downside, however, was the need for archiving system states and torque commands at different time scales¹¹ directly into the motor library (vs parameterization), decreasing the data-efficiency and generality of the approach.

Alternatively, Forte et al. [2012], extending the work of Ude et al. [2010], employed Gaussian Process Regression (GPR) to estimate a generalization mapping of DMP policies from multiple demonstrations, which were stored in memory and encoded into DMP systems for data efficiency and numerical tractability, allowing practical real-time usage. Both works validate their methodologies in reaching and grasping movements with a humanoid robot, generalizing them to new goal positions, time scales and movement shapes. This statistical approach would later be adopted by Deniša et al. [2016] and Petrič et al. [2018] in the particular case of compliant MPs; in the latter, expansion of the motor database, when extrapolated motions are successfully executed, is discussed. Notwithstanding, generally, the incremental construction of this database is not directly addressed, and, similarly to the referenced works [Kulvicius et al., 2012; Pastor et al., 2009; Stulp & Sigaud, 2012], its entries are assumed to constitute complete, functional behaviors and not simpler, low-dimensional primitive elements.

Regarding the segmentation of demonstrated behaviors in IL, Meier et al. [2011] proposed a frame-

¹⁰e.g. Cost-regularized Kernel Regression (CrKR) combines MPs by a weighted average of parameter sets [Kober, 2014].

¹¹Torques are usually not linearly scalable [Petrič et al., 2018].

work based on Expectation-Maximization (EM) to identify MP sequences from partially observed trajectories, and applied it to character recognition in simulation. However, the authors assume a MP library to be previously learned and readily available to the agent. Later, Lioutikov et al. [2017] developed a similar probabilistic algorithm to generate multiple sequential MPs, generating a MP library in a data-efficient way; employing Probabilistic Movement Primitives (ProMPs) [Paraschos et al., 2013] for motor representation, the methodology is validated on two different real robots, in composite writing, chair assembly and table tennis tasks. Using DMPs to encode compliant control trajectories, Yang et al. [2019] focused on transferring manipulation skills from humans to robots using motion capture systems to record kinematics and Electromyography (EMG) to estimate joint stiffness, segmenting demonstrations into simpler MPs. In these works, however, constructed libraries are task- and platform-specific, lacking generality, and only sequential MPs combinations are explored.

More recently, accompanying the increasing interest of roboticists in Deep Learning (DL) models, Shankar et al. [2020] framed IL as a *recomposition* problem, and employed Long Short-term Memory (LSTM) networks and Latent Variable Models (LVMs) to identify MPs from a set of observed behaviors in a total of 20 object manipulation and bimanual operation tasks. Although results were promising and validated in a real robot, no parametric motor representation e.g. DMPs was used; moreover, despite the unsupervised nature of the proposed methodology, training required an extensive preconstructed dataset (8000 kinesthetic demonstrations). In contrast, Dalal et al. [2021] constructed a general-purpose library of MPs e.g. *lift* and *push* motions, employing *deep* RL to identify optimal combinations in a variety of tasks in simulated environments. Despite constraining the motor library to a fixed set of manually defined MPs, the notable aspect of this work was the primitive library being shared across different tasks, environments, and robots. Other recent research using DL to extrapolate and generalize motor policies, such as works from Levine et al. [2016, 2018], Pahič et al. [2020] and Kalashnikov et al. [2021], generally build and curate extensive vision-based datasets not unlike the previous behavior databases. However, in addition to overlooking MPs as basic behavior components, the large-scale data collection required (commonly in tens of thousand of episodes [Chatzilygeroudis et al., 2017; Levine et al., 2018]) can often be prohibitive from the standpoint of autonomous L2 machines in the real world [Dalal et al., 2021].

The construction of motor libraries in robotic control implies a practical compromise between generality and tractability, often facilitated with heuristics [Lioutikov et al., 2017], but nonetheless limited by dimensionality and *a priori* information available [Droniou et al., 2014]. In addition, one common drawback of the referenced works is failing to learn multi-level motor databases, focusing on either generalizing preexisting behaviors, identifying MPs in specific tasks, or addressing both problems separately [Guan et al., 2021; Lioutikov et al., 2017; Ravichandar et al., 2020]. Alternatively, this work focuses on establishing a generic, modular and flexible structure for motor generation, that can simultaneously be harnessed by learning techniques and task-parametrized in straightforward fashion. The mathematical formulation is inspired by Rückert & d’Avella [2013], but extended to integrate high-level motor libraries, with practical validation in more complex applications *i.e.* bipedal locomotion in 3D humanoid robots. While generalization is only explored across different platforms, initialization and expansion of behavior libraries are addressed.

2.3 Generic Primitive-based Control

This work proposes a generic multi-level modular approach to robotic control built on top of the DMP framework, with a particular focus on MP sharing and combination. To that end, the action space is decomposed into mid-level unidimensional MPs (*i.e. what*), and high-level flexible combinations of MPs (*i.e. how*), from which open-loop control patterns are obtained, and tracked through low-level linear feedback control, while accounting for task-specific feedback mechanisms in closed control loops.

The conceptual overview of this control strategy is depicted in Figure 5: within a modular architecture, a *primitive library* **P** gathers generic MPs (defined within DMP dynamics) in platform- and task-agnostic fashion, and a *motor skill* is defined as a combination of MPs in a specific platform - effectively mapping them to a finite set of controls - and stored in a *skill library* **S**. This implies separate *extrinsic* and *intrinsic* parameterization, denoting, respectively, generic MPs and their combination in a motor skill; as previously observed, this can facilitate behavior adaptation and generalization by prioritizing the exploration of different combinations of existing MPs over extending **P** - in theory a more complex process.

Furthermore, as a form of Imitation Learning, a simple *offline* Trajectory Optimization (TO) strategy for motor skill learning is proposed employing Black Box Optimization (BBO), used to populate both primitive and skill libraries from functional preset behaviors.

The generic nature of this approach is highlighted: it is both *versatile*, not tied down to any specific task, platform or control topology *e.g.* kinematic or dynamic trajectories in joint or task space; and *extensible*, accommodating for alternative low-level techniques to track prescribed motions *e.g.* joint compliance for higher efficiency, real-time learning strategies to acquire and refine motor skills with RL, and practical task-oriented feedback mechanisms.

In contrast with works where generalization is approached with task-parameterized MPs [Calinon, 2016; Matsubara et al., 2011; Pervez et al., 2017] or contextual RL [Celik et al., 2022; End et al., 2017; Kober et al., 2012; Mülling et al., 2013], herein MPs constitute task-agnostic elements that can be recruited as building blocks for intricate motor control, representing basic components of a multi-tiered motor database that can be harnessed for generalization and adaptation in Lifelong Learning (L2) machines, exploiting the open-ended nature of **P** and **S** for continual and incremental growth. Under this control structure, motor skills can be seen as latent representations of motor behavior using MP semantics.

2.3.1 Dynamic Movement Primitives

Per the most common formulation [Ijspeert et al., 2013; Saveriano et al., 2021], DMPs rely on second-order spring-damper dynamical systems to ensure stable convergence of control variable y towards a unidimensional goal state g *i.e.* a fixed setpoint or oscillatory baseline:

$$\tau \dot{z} = \alpha_z(\beta_z(g - y) - z) + f \tag{1}$$

$$\tau \dot{y} = z \tag{2}$$




Figure 5: Conceptual diagram of the control strategy proposed in this chapter, in the case of a K DoF platform; a generic primitive library $\mathbf{P}$ provides abstract task- and platform-agnostic MPs combined in different motor skills $\mathcal{S}^s$, archived in a skill library $\mathbf{S}$; closed-loop control is achieved with task-specific feedback mechanisms that modulate ($\cdots$) prescribed trajectories y or phase ϕ, s .

which, for simplicity, can be rewritten, discarding the auxiliary variable z , as:

$$\tau^2 \ddot{y} = \alpha_y (\beta_y (g - y) - \dot{y}) + f \quad (3)$$

where τ denotes a temporal scaling factor, and $\{\alpha_y, \beta_y\}$ define the system damping¹²; the open-loop term f , in its turn, defines a perturbing force that shapes the trajectory of y towards g , defined as a linearly-weighted normalized combination of N Radial Basis Functions (RBFs) or *kernels* Ψ_n :

$$f(\phi) = \frac{\sum_{n=1}^N w_n \Psi_n(\phi)}{\sum_{n=1}^N \Psi_n(\phi)} \quad \vee \quad f(s) = \frac{\sum_{n=1}^N w_n \Psi_n(s)}{\sum_{n=1}^N \Psi_n(s)} \quad (4)$$

anchored on proxy *phase* variables ϕ and s in case of rhythmic¹³ and discrete (point-to-point) motions, respectively. The individual basis Ψ_n are in turn implemented using radial bell-shaped curves *i.e.* von Mises¹⁴ functions in the former and Gaussian (Figure 6) functions in the latter:

$$\Psi_n(\phi) = \exp(h_n(\cos(\phi - c_n) - 1)) \quad \vee \quad \Psi_n(s) = \exp\left(-\frac{(s - c_n)^2}{2h_n^2}\right) \quad (5)$$

¹²Per the generic equation of motion for a damped oscillator, $c = \alpha_y$ is the damping and $k = \alpha_y \beta_y$ the spring constant; in critically damped systems *i.e.* $\zeta = 1$, $\alpha_y = 4\beta_y$. β_y can also be seen as a stiffness factor [Scherzinger et al., 2023]

¹³Simplified formulation *i.e.* unitary amplitude in f [Peternel et al., 2016].

¹⁴von Mises functions are particularly suited due to their periodic nature, which avoids discontinuities in f at the start/end of each execution period and therefore leads to smoother transitions; [Saveriano et al., 2021].

where c denotes the center and h scales the width of the bell shape around c ; each kernel Ψ_n is thus fully parametrized by $\{w_n, c_n, h_n\}$, $\forall n \in \{1, 2, \dots, N\}$.

In practice, ϕ and s are provided by a secondary system that guides execution of encoded motion, removing explicit time dependency from (4); while the authors note that these are flexible and should be selected to fit the application¹⁵, Ijspeert et al. [2013, 2003] provide bounded linear fixed-pace oscillators and quadratic decaying first-order systems as reference for rhythmic and discrete motions, respectively:

$$\tau \dot{\phi} = 1 \quad \vee \quad \tau \dot{s} = -\alpha_s s \quad (6)$$

with both systems are also temporally scaled with τ as in (1 - 3). In both cases, modulating τ leads to topologically equivalent trajectories, effectively defining the movement frequency in rhythmic systems¹⁶. In discrete applications, the overall shape is further influenced by α_s i.e. how fast s decays.

Phase values should be consistent with the domain of the basis functions in (5); commonly, for simplicity and consistency, homogeneous regularly-spaced kernels are convenient in order to ensure the whole phase domain is equally controllable [Saveriano et al., 2021]:

$$c_n = \phi_{\min} + \frac{(\phi_{\max} - \phi_{\min})}{N+1}n \quad \vee \quad c_n = \exp\left(-\alpha_s \frac{i-1}{N-1}\right) \quad (7)$$

$$h_n = \frac{1}{(c_{n+1} - c_n)^2} \quad (8)$$

with $\phi_{\min}$ and $\phi_{\max}$ as, respectively, the lower and upper bounds of rhythmic phase ϕ . An alternative heuristic to (8) is, when dealing with equally-spaced basis functions, to scale h_i so that the midpoint between consecutive kernels roughly matches half-scale ($\Psi(\phi) \sim 0.5$) and minimum amplitude at the center of neighboring kernels; given this simplification, a weight vector $\mathbf{w} = [w_1, w_2, \dots, w_N]^T$ is sufficient to define f , with N directly influencing its spatial resolution and overall dimensionality.

Under conventional DMP terminology, (1 - 3) describe the *transformation* system, providing the underlying attractor landscape of the system dynamics, and (6) the *canonical* system, acting as a sort of *pacemaker* for the system's evolution, often shared between multiple transformation systems for synchronization and coupling in e.g. multi-DoF platforms or task-space control. Figure 6 illustrates the temporal evolution of both systems in a discrete DMP. In practice, despite Ijspeert et al. [2013, 2003] highlighting the underlying control structure over the specific dynamics in (1 - 6) as the main scientific contribution of DMPs, the presented formulation exhibits several attractive properties from a control perspective:

- **Stability:** Due to stable and well-understood attractor dynamics in (1 - 3), as long as $\tau > 0$ and the system is critically damped ($\alpha_z > 0$ and $\alpha_z = 4\beta_z$), then y is bound to *monotonically* converge to g , even if momentarily perturbed [Ijspeert et al., 2013; Saveriano et al., 2021];
- **Invariance:** Scaling parameter in (1 - 3) implies temporal invariance i.e. updating $\tau' = k\tau$ with

¹⁵Alternatives to (6) such as exponential and sigmoid curves have been proposed in the literature [Saveriano et al., 2021].

¹⁶Also described with $\Omega \equiv w \equiv \dot{\phi} = 1/\tau$ [Ijspeert et al., 2013; Saveriano et al., 2021].




Figure 6: **Top Left:** Multiple basis functions Ψ_n on a periodic phase space; **Top Right:** DMP transformation system, converging to $g = 10$; **Bottom Left:** DMP rhythmic canonical system *i.e.* linear pacer, providing values of ϕ ; **Bottom Right:** DMP discrete canonical system, following different decay approaches: exponential (blue), linear (orange) and sigmoid (green).

$k > 0$ preserves the overall shape of the trajectory, with subsequent works extending the original formulation addressing spatial invariance [Ijspeert et al., 2013; Pastor et al., 2009];

- **Online Modulation:** The lack of explicit time dependency in (4) allows adjusting, in real time, the evolution of the system, whether by updating g , manipulating ϕ (or s) directly, or using advanced goal switching [Ijspeert et al., 2013] and phase stopping mechanisms [Ijspeert et al., 2002];
- **Parameter Linearity:** Beyond their parametric nature which provides a way to deal with high-dimensional systems, the linearity in w_n (4) improves numerical tractability and makes DMPs particularly appealing for learning applications, including IL and RL approaches;
- **Versatility:** By encoding arbitrarily complex trajectories, providing discrete/rhythmic variants and compatibility with notions such as viapoints [Weitschat & Aschemann, 2018] and spatiotemporal coupling [Ijspeert et al., 2013], DMPs fit a wide range of applications in and beyond robotic control.

Exploiting these features, the control strategy in Figure 5 employs the DMP framework to encode open-loop control patterns in generic form; in particular, the nonlinear component f in (3) and (4) is understood to uniquely describe individual MPs in a primitive library $\mathbf{P}$:

$$\mathbf{P} = [f^1, f^2, \dots, f^P]^\top \quad (9)$$

such that $\mathbf{P} : \mathbb{R} \rightarrow \mathbb{R}^P$, and each element is parametrized by $\mathbf{w}^p \in \mathbb{R}^N, \forall p \in \{1, 2, \dots, P\}$ in task- and platform-agnostic form; conversely, the transformation (3) and canonical (6) systems incorporate platform-specific control. From (9), the complete set of parameters defining $\mathbf{P}$ is:

$$\mathbf{W} = [\mathbf{w}^1, \mathbf{w}^2, \dots, \mathbf{w}^P]^\top \quad (10)$$

which is expanded to a $P \times N$ matrix if $\Psi_n, \forall n \in \{1, 2, \dots, N\}$ are constant across the whole library; otherwise, $\mathbf{W} \in \mathbb{R}^{\sum_{p=1}^P N_p}$, with N_p as the number of kernels in the p^{th} primitive in $\mathbf{P}$. From the perspective of the proposed control structure (Figure 5), $\mathbf{W}$ is seen as form of *extrinsic* parameterization.

Throughout the remainder of this work, the classical rhythmic formulation is considered; notwithstanding, unless expressly stated otherwise, relevant insights and observations extend to the discrete case, as well as to the referenced extensions to DMP dynamics. Moreover, in line with conventional applications, $\{y, \dot{y}, \ddot{y}\}$ is viewed as an open-loop trajectory reference, whose low-level tracking with e.g. linear feedback control falls outside the scope of this study.

2.3.2 Shared Information and Primitive Combination

By and large, practical DMP applications adopt a separate transformation system for each individual dimension, which commonly employ separate primitives f (3); while this mostly straightforward approach is adequate in simple tasks where generalization and scalability are not critical concerns, control complexity can increase significantly in higher-dimensional tasks and platforms [Saveriano et al., 2021].

In an alternative holistic strategy, dimensionality reduction can be achieved - beyond the parametric nature of DMPs - by exploiting the often redundant nature of articulated platforms, and introducing *shared information* across different control dimensions. In fact, one of the core ideas behind MP theory lies in the reuse of information - in the form of elemental motions - across different dimensions, tasks, and potentially even different platforms. In an objective formulation of this idea, (3) is updated:

$$\tau^2 \ddot{y} = \alpha_y (\beta_y (g - y) - \dot{y}) + f^*(\phi) \quad (11)$$

with $f^*(\phi)$ denoting a combination of P primitives f^p (4):

$$f^*(\phi) = \sum_{p=1}^P \lambda^p f^p(\phi + \delta^p) \quad (12)$$

where λ^p and $\delta^p, \forall p \in \{1, 2, \dots, P\}$, linearly scale and phase shift f^p , respectively. When considering K coupled dimensions $y_k, \forall k \in \{1, 2, \dots, K\}$ sharing a single phase reference ϕ , for each k :

$$f_k^*(\phi) = \sum_{p=1}^P \lambda_k^p f^p(\phi + \delta_k^p) \quad (13)$$

which can be expanded, given primitive library $\mathbf{P}$, to:

$$f_k^*(\phi) = \boldsymbol{\lambda}_k^{*\top} \mathbf{P}(\phi + \boldsymbol{\delta}_k^*) = \begin{bmatrix} \lambda_k^1 \\ \lambda_k^2 \\ \vdots \\ \lambda_k^P \end{bmatrix}^\top \times \begin{bmatrix} f^1(\phi + \delta_k^1) \\ f^2(\phi + \delta_k^2) \\ \vdots \\ f^P(\phi + \delta_k^P) \end{bmatrix} \quad (14)$$

with $\boldsymbol{\lambda}_k^* \in \mathbb{R}^P$ and $\boldsymbol{\delta}_k^* \in \mathbb{R}^P$ representing, respectively, scale and shift vectors for the k^{th} dimension. Globally, for the whole system, matrices $\boldsymbol{\Lambda} \in \mathbb{R}^{K \times P}$ and $\boldsymbol{\Delta} \in \mathbb{R}^{K \times P}$ aggregate individual parameters:

$$\boldsymbol{\Lambda} = \begin{bmatrix} \boldsymbol{\lambda}_1^{*\top} \\ \boldsymbol{\lambda}_2^{*\top} \\ \vdots \\ \boldsymbol{\lambda}_K^{*\top} \end{bmatrix} = \begin{bmatrix} \lambda_1^1 & \lambda_1^2 & \cdots & \lambda_1^P \\ \lambda_2^1 & \lambda_2^2 & \cdots & \lambda_2^P \\ \vdots & \vdots & \ddots & \vdots \\ \lambda_K^1 & \lambda_K^2 & \cdots & \lambda_K^P \end{bmatrix} \quad \boldsymbol{\Delta} = \begin{bmatrix} \boldsymbol{\delta}_1^{*\top} \\ \boldsymbol{\delta}_2^{*\top} \\ \vdots \\ \boldsymbol{\delta}_K^{*\top} \end{bmatrix} = \begin{bmatrix} \delta_1^1 & \delta_1^2 & \cdots & \delta_1^P \\ \delta_2^1 & \delta_2^2 & \cdots & \delta_2^P \\ \vdots & \vdots & \ddots & \vdots \\ \delta_K^1 & \delta_K^2 & \cdots & \delta_K^P \end{bmatrix} \quad (15)$$

which comprise the *intrinsic* parameterization of the generic control strategy (Figure 5), effectively mapping P generic primitives f^p into K dimensions. (11 - 14) represent a reformulation of the primitive combination scheme proposed by Rückert & d'Avella [2013], revised for clarity and consistency; in a practical sense, $\boldsymbol{\Lambda}$ matches the role of gating networks in the literature [Mülling et al., 2013, 2010], while $\boldsymbol{\Delta}$ allows sequential chaining of MPs [Kulvicius et al., 2012; Samant et al., 2016; Stulp et al., 2012].

Building upon this foundation, given $\mathbf{P}$, a *motor skill* is defined solely through intrinsic parameters:

$$\mathcal{S} = \{\boldsymbol{\Lambda}, \boldsymbol{\Delta}\} \quad (16)$$

which can be coupled to perceptual input or contextual metaparameters in high-level planning stages for generalization [Kober, 2014; Pervez et al., 2017]. Given τ , goal $\mathbf{g} \in \mathbb{R}^K$, and an initial state $\mathbf{y}_0 \in \mathbb{R}^{K^{17}}$, the complete set of control trajectories $\mathbf{y} \in \mathbb{R}^K$ is updated numerically using (11), exploiting the stability of DMP dynamics for smoother transitions between different motor skills in real time.

At a conceptual level (Figure 5), analogously to $\mathbf{P}$ (9), a motor skill library $\mathbf{S}$ is defined:

$$\mathbf{S} = [\mathcal{S}^1, \mathcal{S}^2, \dots, \mathcal{S}^S]^\top \quad (17)$$

From a Lifelong Learning perspective, $\mathbf{P}$ and $\mathbf{S}$ are understood to be dynamically sized, growing as new primitives and skills are acquired; conversely, a fixed number of K dimensions are controlled with constant damping $\{\alpha_{y,k}, \beta_{y,k}\}^{18}$. In order to bound intrinsic parameterization and avoid incrementally updating all $\mathcal{S}$ as new primitives are learned, motor skills can be defined over a finite subset of the primitive library containing the $P' \leq P$ MPs required to codify control patterns; this entails updating (14) to:

$$f_k^*(\phi) = \boldsymbol{\lambda}_k^{*\top} \mathbf{mP}(\phi + \boldsymbol{\delta}_k^*) \quad (18)$$

¹⁷In practice, only relevant in discrete motions.

¹⁸ $\alpha_{y,k}$ and $\beta_{y,k}$ are usually hand-tuned to fit the application [Scherzinger et al., 2023].

where $\mathbf{m} \in \{0, 1\}^{P' \times P}$ is a binary selection matrix identifying P' relevant primitives¹⁹ in $\mathbf{P}$, e.g. :

$$\mathbf{mP} = \underbrace{\begin{bmatrix} 1 & 0 \\ 0 & 0 \\ \vdots & \vdots \\ 0 & 1 \end{bmatrix}}_{P'}^\top \times \begin{bmatrix} f^1 \\ f^2 \\ \vdots \\ f^P \end{bmatrix} \begin{matrix} \checkmark \\ \\ \\ \checkmark \end{matrix} \quad (19)$$

and which can also be used to mask parameterization i.e. $\mathbf{W}' = \mathbf{mW}$, and appended to $\mathcal{S}$ (16):

$$\mathcal{S} = \{\mathbf{\Lambda}, \mathbf{\Delta}, \mathbf{m}\} \quad (20)$$

with $\mathbf{\Lambda} \in \mathbb{R}^{K \times P'}$ and $\mathbf{\Delta} \in \mathbb{R}^{K \times P'}$. Overall, this implies a fixed number of $n_{\text{int}} = 2 \times K \times P'$ intrinsic parameters, in contrast with $K \times N$ in the conventional approach²⁰ where only primitive-specific i.e. extrinsic parameterization is considered, herein seen instead as task- and platform-agnostic (Figure 5). Fully describing a motor skill therefore requires a total of $n = n_{\text{int}} + n_{\text{ext}} = 2 \times K \times P' + P' \times N$ parameters; when compared to conventional strategies, the increase in complexity brought about by the additional n_{int} parameters can be offset by sufficiently small P' values; in practice, global dimensionality n is reduced as long as $P' \leq \lfloor \frac{KN}{2K+N} \rfloor$. This highlights one of the major appeals of MP theory in control; from the practical standpoint of joint-space control, a significant number of multi-DoF systems e.g. industrial manipulators or legged platforms lend themselves to this approach through their morphological redundancy, which often allows for P' to be much smaller than K .

2.3.3 Motor Skill Learning and Generalization

Generally, in applications wherein MP combination is not considered, codifying a K -dimensional motor policy into a DMP system implies learning K nonlinear terms f (3) by parametrizing the underlying $K \times N$ kernels Ψ_n which, under a simplified scheme, requires finding K weight vectors $\mathbf{w}$ (4). A widely popular methodology is to exploit *expert* or *reference* trajectories obtained either from kinesthetic demonstrations - with human agents guiding the robot remotely or in a compliant manner²¹ - or kinematic trajectories recorded with motion capture systems; under the broader scope of IL methods, this kind of supervised learning represents a practical approach of Learning from Demonstration (LfD), wherein a motor policy is learned by interactive teaching, and which was found to be successful in a large variety of robotic applications [Ijspeert et al., 2013; Ravichandar et al., 2020; Saveriano et al., 2021]. At its simplest, in DMP learning with LfD, each of the K transformation systems (3) is rearranged to:

$$f_{\text{target}} = \tau^2 \ddot{y}^{\text{REF}} - \alpha_y (\beta_y (g - y^{\text{REF}}) - \tau \dot{y}^{\text{REF}}) \quad (21)$$

¹⁹Each row containing one non-zero element i.e. $\sum_{p=1}^P m_{ip} = 1, \forall i \in \{1, 2, \dots, P'\}$; if $\mathbf{m} = \mathbf{I}_P$, entire library $\mathbf{P}$ is used (14).

²⁰Assuming constant primitive complexity i.e. N fixed-width equidistant kernels; generally, N may vary with f^p (10).

²¹Appealing as movement is constrained by the robot kinematics and dynamics, simplifying optimization [Deniša et al., 2016]

where y^{REF} denotes the reference values to approximate. For its simplicity, Locally Weighted Regression (LWR) [Atkeson et al., 1997] is commonly adopted by roboticists when learning DMP policies; it frames the problem as standard function approximation wherein a localized - *i.e.* linearly weighted by kernel Ψ_n - measure of the quadratic error J_n , for all N basis in f , is minimized [Ijspeert et al., 2013]:

$$J_n = \sum_{n=1}^N \Psi_n(t) (f_{\text{target}} - w_n)^2 \quad (22)$$

which, given a T -sample reference $\mathbf{y}^{\text{REF}} \in \mathbb{R}^T$, has complexity of $\mathcal{O}(N^2T)$ [Gams & Ude, 2009] in each of the K dimensions, and can be solved in *batch* or *incrementally*²² [Ijspeert et al., 2013]. Alternative methods to minimize (22) include Gaussian Mixture Models (GMMs)/Gaussian Mixture Regression (GMR) [Calinon et al., 2010, 2012; Chen et al., 2017], Gradient Descent (GD) [Samant et al., 2016], Deep Learning (DL) [Pahič et al., 2018, 2020] and, notably, Gaussian Process Regression (GPR) [Deniša et al., 2016; Matsubara et al., 2011]. In particular, GPR does not require N to be set prior to learning [Ude et al., 2010], and, unlike LWR, is not limited to learning weights w_n [Samant et al., 2016]; on the other hand, it has been shown to be both more computationally demanding and less accurate than LWR [Forte et al., 2012]. Nonetheless, a common practical issue of LfD methodologies is that N often grows exponentially as K increases, adding complexity in dealing with high-dimensional inputs *e.g.* demonstrations from highly-articulated platforms [Saveriano et al., 2021].

Beyond fully parametrizing motion shape f (3) from scratch, new behaviors may also be learned by adapting or generalizing existing skills to new contexts or tasks. While the invariance properties of DMPs streamline extrapolating motions to different time horizons, frequencies or spatial targets *i.e.* adjusting τ and g (2.3.1) in response to external stimuli [Ijspeert et al., 2003; Li et al., 2018; Pastor et al., 2009; Scherzinger et al., 2023], generalization of learned behaviors is often treated as a separate optimization problem, tackled with methods ranging from GPR [Forte et al., 2012; Ude et al., 2010] and contextual learning strategies [Klink et al., 2020; Kober, 2014; Queißer & Steil, 2018] to flexible combination schemes *i.e.* overlapping and/or sequencing individual MPs [Kulvicius et al., 2012; Mülling et al., 2013; Pastor et al., 2009; Samant et al., 2016; Stulp et al., 2012] (2.2.5).

By contrast, the modular multi-tiered organization proposed (Figure 5) allows a higher-level approach to motor skill acquisition, which accommodates both learning from scratch and behavior generalization. In a broader formulation, given the full parameter vector $\boldsymbol{\theta} = [\mathbf{W}, \boldsymbol{\Lambda}, \boldsymbol{\Delta}]^\top$ (2.3.2), a straightforward Trajectory Optimization (TO) strategy is to minimize the cumulative residual error $J(\boldsymbol{\theta}, \mathbf{y}^{\text{REF}})$ between the states of a $\boldsymbol{\theta}$ -parametrized DMP system and a given reference trajectory until a finite time horizon T :

$$J(\boldsymbol{\theta}, \mathbf{y}^{\text{REF}}) = \sum_{k=1}^K J_k(\boldsymbol{\theta}, \mathbf{y}_k^{\text{REF}}) = \sum_{k=1}^K \sum_{t=1}^T (y_{k,t}^{\boldsymbol{\theta}} - y_{k,t}^{\text{REF}})^2 \quad (23)$$

where y_t^{REF} and $y_{k,t}^{\boldsymbol{\theta}}$ denote, respectively, reference values and associated DMP state (11) in the k^{th} dimension, at the t^{th} sample, $\forall t \in \{1, 2, \dots, T\}$. In order to ensure that $\mathbf{y}^{\boldsymbol{\theta}} \in \mathbb{R}^T$, a reference phase

²²By employing a custom forgetting factor $\lambda_{\text{LWR}} < 1$ [Ijspeert et al., 2013; Petrič et al., 2018].

vector $\phi^{\text{REF}} \in \mathbb{R}^T$ can optionally be estimated given finite phase bounds $\Phi = [\phi_{\min}, \phi_{\max}]$, replacing the DMP canonical system (6). The goal of TO is thus to find θ^{best} such that:

$$\theta^{\text{best}} = \arg \min_{\theta} [J(\theta, \mathbf{y}^{\text{REF}})] \quad (24)$$

which presents a generic least-squares optimization problem that can be solved with a large number of techniques - and whose benchmark ultimately goes beyond the scope of this work. In most practical scenarios, however, (23) can potentially be extended to a multi-objective²³ approach, e.g. :

$$\begin{aligned} J(\theta, \mathbf{y}^{\text{REF}}) = \sum_{k=1}^K J_k(\theta, \mathbf{y}_k^{\text{REF}}) = & \underbrace{\alpha_1 \sum_{k=1}^K \sum_{t=1}^T (y_{k,t}^{\theta} - y_{k,t}^{\text{REF}})^2}_{\text{track reference}} \\ & + \underbrace{\alpha_2 \sum_{k=1}^K \sum_{t=1}^T (\ddot{y}_{k,t}^{\theta})^2}_{\text{minimize jerk}} \\ & + \underbrace{\alpha_3 (\|\theta\|_2)^2}_{\text{regularization}} \end{aligned} \quad (25)$$

favoring minimum-jerk²⁴ trajectories for smoother and human-like motions with reduced energetic footprint [Flash & Hogan, 1985; Schaal et al., 2007]; and scaling down θ with ℓ_2 regularization [Boyd & Vandenberghe, 2004]²⁵. Additional soft constraints can e.g. enforce transient viapoints [André et al., 2013; Queißer & Steil, 2018; Werner et al., 2014], however it should be noted that this often comes at the cost of simplicity, as careful tuning of weights α_1 , α_2 and α_3 becomes critical.

The generic formulation in (23), in conjunction with the modular organization of the proposed control structure (Figure 5), enable the implementation of different learning modalities exclusively by constraining parameterization, regardless of how optimization is performed; as illustrated in Figure 7, given $\mathbf{P}$ and $\mathbf{S}$, θ can be manipulated²⁶ to reflect different types of motor skill learning, namely:

- **Intrinsic**, extending $\mathbf{S}$ through a new or updated combination of known MPs in $\mathbf{P}$; current extrinsic parameterization $\mathbf{W}_0$ is preserved;
- **Extrinsic**, where MPs in $\mathbf{P}$ are updated under a fixed combination scheme²⁷ defined by intrinsic parameters $\{\Lambda_0, \Delta_0, \mathbf{m}_0\}$; conventional DMP control if $P' = K$, $\Lambda_0 = \mathbf{I}_K$ and $\Delta_0 = \mathbf{0}_K$;
- **Expansive**, compounding intrinsic and extrinsic learning to expand both $\mathbf{P}$ and $\mathbf{S}$ with new MPs and motor skills, respectively; the full parameter set θ is learned;

²³ $\alpha_1 \gg \alpha_2$ and $\alpha_1 \gg \alpha_3$ in order to ensure reference trajectory is efficiently tracked;

²⁴ Applicable when y represents joint- or task-space position control commands.

²⁵ Illustrative case; in practice, DMP dynamics inherently prevent significant jerk during the movement (2.3.1).

²⁶ Parameter-specific constraints in θ e.g. binary values in $\mathbf{m}$ are assumed to be enforced on an algorithmic basis.

²⁷ Analogous to the organization of musculoskeletal systems through Muscle Synergy (MS) [d'Avella & Bizzi, 2005].




Figure 7: Illustration of different skill learning modalities achieved with different constraints to θ ; the number of primitives P' , as well as N in extrinsic and expansive learning, are viewed as open parameters; $\mathbf{W}_0, \mathbf{\Lambda}_0, \mathbf{\Delta}_0$ denote preset parameterization; for simplicity, in intrinsic and sequential modalities, $\mathbf{m}$ (18) can be omitted i.e. $\mathbf{m} = \mathbf{I}_P$ for small P (9) e.g. during early cognitive stages.

- **Sequential**, as a special case of intrinsic learning, where MPs in $\mathbf{P}$ are sequenced according to a vector $\mathbf{i} \in \mathbb{N}^{P'}$ specifying an arbitrary MP order;

Accordingly, parameter dimensionality differs in each of these learning modalities; generally, solving (23) or (25) iteratively with numerical methods is directly impacted by n , and thus by K, P (or P') and N , with computational complexity and convergence time often growing exponentially²⁸; on the other hand, in the particular case of classic extrinsic learning where dimensions are codified independently i.e. $P' = K$, $\mathbf{\Lambda}_0 = \mathbf{I}_K$ and $\mathbf{\Delta}_0 = \mathbf{0}_K$, the problem can alternatively be approached as K separate optimization steps of dimension N (in contrast with $K \times N$), which can, in theory, greatly reduce total convergence time. Another simplification strategy is to provide initial guesses or *priors* for θ , which can be straightforward in cases of conceptual similarity between motor skills, and will be explored in this work.

Beyond the examples of guided learning through optimization constraints presented in Figure 7, more intricate learning strategies can be formulated in similar fashion. The flexibility of the approach is highlighted, which can be harnessed to meet different task or platform requirements, or achieve different learning goals. In practice, not only are motor skills implicitly generalizable as low-dimensional combinations or sequences of generic MPs - facilitating e.g. online RL - but by restricting their definition to $f^*(\phi)$ in (11), generalization strategies that adapt the underlying attractor dynamics e.g. metaparameters τ and g [Dragan et al., 2015; Forte et al., 2012; Koutras & Doulgeri, 2020; Kulvicius et al., 2012; Pastor et al., 2009; Ude et al., 2010; Weitschat & Aschemann, 2018] are still applicable.

²⁸In the case of e.g. CMA-ES, complexity scales according to $\mathcal{O}(n^3)$ [Hansen et al., 2003; Hansen & Ostermeier, 2001].

Overall, this represents a high-level approach to motor skill learning, where exploring the parameter space is favored over problem-specific knowledge [Sutton, 2019]. While the choice of optimization strategy is left open to suit diverse platform and task particularities, given the heterogeneous nature of θ , the use of Black Box Optimization (BBO) is assumed throughout the remainder of this work. BBO is a branch of ML that approaches the underlying problem empirically *i.e.* without making any assumption about - or attempt to model - its objective function, treating it instead as a *black box* that maps an n -dimensional parameter set to a single scalar *i.e.* $J : \mathbb{R}^n \rightarrow \mathbb{R}$, relying only on practical evaluations for parameter updates [Stulp & Sigaud, 2012, 2013]. It is particularly appealing when little to no information is available about J - whether continuous, linear or differentiable - and in contexts where those traits prevent the application of alternatives such as gradient-based or convex optimization techniques [Bajaj et al., 2021]. On the other hand, BBO is often associated with suboptimal performance and slower convergence, requiring a large number of evaluations, and not being guaranteed to find a global optimum, particularly in non-convex and high-dimensional problems [Alarie et al., 2021; Bajaj et al., 2021]; this becomes problematic outside simulated environments, where evaluations are generally expensive and time-consuming.

Nevertheless, by not imposing requirements on the nature of J and θ , BBO becomes applicable to a wider range of problems than other ML alternatives, having gained popularity in solving real-world problems in a variety of scientific fields [Alarie et al., 2021]. In particular, strong parallels can be traced between episodic RL and iterative BBO, by viewing RL algorithms that explore directly on policy parameters as a form of iterative BBO [Rückstieß et al., 2010]. Incidentally, cutting edge methods such as Covariance Matrix Adaptation - Evolutionary Strategy (CMA-ES) and Bayesian Optimization (BO) were successfully employed in RL as a form of direct Policy Search (PS) or in contextual behavior generalization [Calandra et al., 2014; Chatzilygeroudis et al., 2017, 2020; QueiBer & Steil, 2018; Stulp & Sigaud, 2012, 2013].

The high-level and parameter-agnostic nature of BBO methods can be harnessed by general approaches to optimization, being particularly suited to problems with heterogeneous parameter sets and multimodal objective functions, such as (23) [Heidrich-Meisner & Igel, 2008; QueiBer & Steil, 2018; Stulp & Sigaud, 2012]. In this work, CMA-ES [Hansen & Ostermeier, 2001], is used; as a local²⁹ direct-search BBO method derived from classical ES, CMA-ES computes parameter updates based only on the total cost given by J [Bajaj et al., 2021; QueiBer & Steil, 2018], requiring only minimal parameterization. CMA-ES was shown to be successful in different complex nonlinear optimization problems including, notably, parametrizing DMP systems [QueiBer et al., 2018; QueiBer & Steil, 2018; Rückert & d’Avella, 2013; Stulp & Sigaud, 2013]. Moreover, while not pursued in this study, its black-box nature allows extending θ to parameters beyond $[\mathbf{W}, \mathbf{\Lambda}, \mathbf{\Delta}]^\top$, such as goal states $\mathbf{g} = [g_1, g_2, \dots, g_K]^\top$, system damping $\{\alpha_k, \beta_k\}$, center and width of each basis Ψ_n , and potentially N and P' .

As part of this study, custom code libraries were developed in C++ implementing the control structure described in this section. These libraries were tailored to address the specific challenges posed by the proposed method, ensuring an efficient and accurate computational framework for motor skill representation and acquisition. For a detailed description of the implementation, including the underlying structure and

²⁹In contrast with *global* search methods *e.g.* Bayesian Optimization [Bajaj et al., 2021; Calandra et al., 2014].



Figure 8: Humanoid platforms employed in this work. **Left:** Robotis OP 1 (formerly DARwIn-OP) [ROBOTIS, 2019]; **Right:** Softbank Robotics NAO V5 [Robotics, 2019]. Refer to Appendix B: Humanoid Platforms for detailed information about both robots.

technical specifications, please refer to Appendix A: Practical Implementation.

2.4 Biped Locomotion with Shared Primitives

In order to demonstrate the practicality of the MP-based control strategy described in 2.3, bipedal locomotion is replicated in two different simulated humanoid bipeds, the Robotis OP 1³⁰ and the SoftBank Robotics³¹ NAO V5 (Figure 8). As illustrated in Figure 9, exploiting the generality of task- and platform-agnostic MPs, a single shared primitive library $\mathbf{P}$ is constructed for both robots, and robot-specific skill libraries $\mathbf{S}^{\text{OP}}$ and $\mathbf{S}^{\text{NAO}}$ populated with a *walking* motor skill $\mathcal{S}^1$ that generates stable bipedal locomotion by combining a fixed number of MPs. Conceptually, this provides a foundation for Lifelong Learning (L2) solutions by kickstarting behavior databases in robotic agents, which can be expanded upon by repeating the skill acquisition process for other types of gait as well as, potentially, other general-purpose leg motions.

Both robots are roughly similar in size, hardware and morphology³², with detailed information provided in Appendix B: Humanoid Platforms; when adopting a fixed upper body configuration and replicating motions in both legs and in both robots, control is limited to a total of $K = 6$ DoFs: 1) Hip Pitch (hPitch), 2) Hip Roll (hRoll) and 3) Hip Yaw (hYaw), 4) Knee Pitch (kPitch), 5) Ankle Pitch (aPitch) and 6) Ankle Roll (aRoll). In a practical sense, in line with DMP design (2.3.1) and multidimensional control, each leg DoF is prescribed a rhythmic trajectory y_k through a transformation system (11), driven with coupled leg-specific canonical systems $\phi_i, i \in \{\text{left}, \text{right}\}$ (6), which enforce a half-cycle phase delay between both legs; gait-specific feedback mechanisms modulate y_k and ϕ_i in real time according to measured Ground Reaction Forces (GRFs), adding stability and robustness to the resulting gait through closed-loop control.

³⁰Rebranded from Dynamic Anthropomorphic Robot with Intelligence - Open Platform (DARwIn-OP).

³¹Formerly Aldebaran Robotics.

³²The NAO has both Hip Yaw DoFs symmetrically coupled, effectively representing a single DoF;




Figure 9: Conceptual diagram extending the generic MP-based control strategy in Figure 5 to bipedal walking in the Robotis OP and the SoftBank NAO (Figure 8); in both platforms, each k^{th} leg DoF is controlled by a DMP transformation system y_k (11) and force-based feedback mechanisms close the control loop.

As legged walking, by its complex and dynamic nature, does not lend itself to interactive kinesthetic teaching as commonly performed in LfD, and given the substantial body of research on the development of robotic walking (2.2), functional control patterns of CPG-based gait are used as a demonstrative reference for skill learning (2.3.3). For practicability and simplicity, this process is conducted in a simulated environment in Gazebo [Koenig & Howard, 2004], and entirely offline; within the paradigm of IL, beyond a form of TO (2.3.3), it can also be viewed as a form of Behavior Cloning (BC), wherein control patterns are learned without direct interaction with a teacher or supervisor agent [Ravichandar et al., 2020]. By focusing on trajectory control in joint space, it bypasses the computational overhead of Inverse Kinematics (IK) solvers, and avoids complex singularities that may occur in task-space control [Hersch et al., 2008; Ravichandar et al., 2020]. Additionally, by imitating open-loop control patterns directly, the previous practical validation of the original gait on real platforms is leveraged *i.e.* similar walking performance is expected.

2.4.1 Central Pattern Generators in Closed-loop Gait Control

On the line of bioinspired locomotion control, as reviewed in 2.2.4, a popular approach to achieve limit cycle walking is through oscillator networks exploiting the notion of Central Pattern Generators (CPGs); one such case is presented in Matos & Santos [2012, 2014], wherein open-loop rhythmic patterns are generated by *unit-burst* elements for each leg DoF of an OP humanoid robot (Figure 8), as the result of weighted contributions of RBFs with sinusoidal and bell-shaped profiles:

$$f(\phi_i) = -A\omega \sin(\phi_i + c) \quad (26)$$

$$f(\phi_i) = \frac{-A\omega(\phi_i + c)}{h^2} \exp\left(-\frac{(\phi_i + c)^2}{2h^2}\right) \quad (27)$$

where A and ω scale the amplitude, and, similarly to (5), c phase-shifts the center of the RBF, with h defining the width of the bell shape in (27); in both cases, ϕ_i denotes the motion phase in leg $i \in \{\text{left, right}\}$. In each DoF, a custom set of RBFs influence the evolution of a nonlinear dynamical system:

$$\dot{y}_{i,hPitch}^{CPG} = -\alpha(y_{i,hPitch}^{CPG} - O_{i,hPitch}) + f_{hPitch}^{flexion} + f_{hPitch}^{compass} \quad (28)$$

$$\dot{y}_{i,hRoll}^{CPG} = -\alpha(y_{i,hRoll}^{CPG} - O_{i,hRoll}) + f_{hRoll}^{balancing} \quad (29)$$

$$\dot{y}_{i,hYaw}^{CPG} = -\alpha(y_{i,hYaw}^{CPG} - O_{i,hYaw}) + f_{hYaw}^{rotation} \quad (30)$$

$$\dot{y}_{i,kPitch}^{CPG} = -\alpha(y_{i,kPitch}^{CPG} - O_{i,kPitch}) + f_{kPitch}^{flexion} + f_{kPitch}^{yield} \quad (31)$$

$$\dot{y}_{i,aPitch}^{CPG} = -\alpha(y_{i,aPitch}^{CPG} - O_{i,aPitch}) + f_{aPitch}^{flexion} + f_{aPitch}^{yield} + f_{aPitch}^{compass} \quad (32)$$

$$\dot{y}_{i,aRoll}^{CPG} = -\alpha(y_{i,aRoll}^{CPG} - O_{i,aRoll}) + f_{aRoll}^{balancing} \quad (33)$$

exhibiting first-order point-attractor dynamics towards $O_{i,k}$, with proportional gain α . Individual motions represent different aspects and/or stages of the locomotion process, requiring platform-specific parameterization, as well as reflecting morphological symmetries of the platform; per (28 - 33), these patterns are divided into 1) *balancing*, shifting the CoM over the supporting leg and unloading the swinging leg in alternating fashion by acting on roll joints:

$$f_{i,hRoll}^{balancing} = -A^{balancing} \omega \sin(\phi_i) \quad (34)$$

$$f_{i,aRoll}^{balancing} = -f_{i,hRoll}^{balancing} \quad (35)$$

2) *flexion* of pitch joints for higher foot clearance during leg swing:

$$f_{i,hPitch}^{flexion} = -\frac{A_{hip}^{flexion} \omega \phi_i}{h^2} \exp\left(-\frac{\phi_i^2}{2h^2}\right) \quad (36)$$

$$f_{i,kPitch}^{flexion} = -\frac{A_{knee}^{flexion} \omega \phi_i}{h^2} \exp\left(-\frac{\phi_i^2}{2h^2}\right) \quad (37)$$

$$f_{i,aPitch}^{flexion} = -(f_{i,hPitch}^{flexion} + f_{i,kPitch}^{flexion}) \quad (38)$$

a 3) *compass* motion that simultaneously moves the swing foot forward and the standing foot backward, displacing the robot along the walking direction:

$$f_{i,hPitch}^{compass} = -A^{compass} \omega \sin\left(\phi_i + \frac{\pi}{2}\right) \quad (39)$$

$$f_{i,aPitch}^{compass} = -f_{i,hPitch}^{compass} \quad (40)$$

knee and ankle 4) *yielding* aimed at absorbing the foot impact and stabilizing the CoM vertical trajectory:

$$f_{i,hPitch}^{yield} = -A^{yield} \omega \sin(\phi_i + \pi) \quad (41)$$

$$f_{i,kPitch}^{yield} = -2A^{yield} \omega \sin(\phi_i + \pi) \quad (42)$$

$$f_{i,aPitch}^{yield} = -f_{i,hPitch}^{yield} \quad (43)$$

and, lastly, 5) *rotation* helps to further stabilize the movement in the transverse plane and control the walking direction:

$$f_{i,\text{hYaw}}^{\text{rotation}} = -A^{\text{rotation}} \omega \sin \left(\phi_i + \frac{\pi}{2} \right) \quad (44)$$

(34 - 44) describe motions applied in both legs, with phase $\phi_i \in [-\pi, \pi]$ shared between same-leg DoFs, and with a half-cycle phase delay between left and right legs, enforced through coupled rhythm generators providing inter-leg coordination:

$$\dot{\phi}_{\text{left}} = \omega + K_C \sin \left(\dot{\phi}_{\text{left}} - \dot{\phi}_{\text{right}} + \pi \right) \quad (45)$$

$$\dot{\phi}_{\text{right}} = \omega + K_C \sin \left(\dot{\phi}_{\text{right}} - \dot{\phi}_{\text{left}} + \pi \right) \quad (46)$$

where ω , in addition to scaling amplitude in individual motions (34 - 44), denotes the walking frequency; K_C represents the coupling strength, and a phase shift of π is enforced on ϕ_i , defining an anti-phase relationship between contra-lateral motion generators on each leg. Phase coupling presents appealing properties, straightforward maintenance of relationships and entrainment, which are used to achieved inter-limb coordination among unit generators [André et al., 2016b; Matos & Santos, 2014; Matos, 2014]

Conceptually, (26 - 44) represent a motor control approach similar to DMP dynamics (2.3.1), extending stable attractors³³ with parametrizable nonlinear RBFs (26 - 27)³⁴, and employing a separate dynamical system to drive temporal evolution implicitly (45); despite being developed within the boundaries of bipedal walking control, in contrast with the general purpose inspiration behind DMP design, they still retain stable behavior through attractor dynamics, time invariance and lack of explicit time dependency, as well as flexible online modulation. Incidentally, rhythmic DMPs have been used in bioinspired CPG-based control [André et al., 2015; Nakanishi et al., 2004; Zhou et al., 2016].

A phase regulation mechanism is proposed in the original work by Matos & Santos [2014], through which gait events are used to reset the phase value in stance-swing and swing-stance transitions. Alternatively, later iterations of this approach adopted instead a perceptual feedback system that modulates $\dot{\phi}_i$ according to measured Ground Reaction Forces (GRFs) [André et al., 2016b, 2015; Owaki et al., 2012]:

$$\dot{\phi}_{\text{left}} = \omega + K_C \sin \left(\dot{\phi}_{\text{left}} - \dot{\phi}_{\text{right}} + \pi \right) + K_F \sin \left(\sum_{j=1}^{N_{\text{left}}^{\text{FS}}} x_j^{\text{FS}} \right) \quad (47)$$

$$\dot{\phi}_{\text{right}} = \omega + K_C \sin \left(\dot{\phi}_{\text{right}} - \dot{\phi}_{\text{left}} + \pi \right) + K_F \sin \left(\sum_{j=1}^{N_{\text{right}}^{\text{FS}}} x_j^{\text{FS}} \right) \quad (48)$$

where x_j^{FS} are direct measurements from each of the N_i^{FS} , $i \in \{\text{left}, \text{right}\}$ Force Sensors (FSs) installed on the foot $i \in \{\text{left}, \text{right}\}$, and K_F scales the strength of the corrective feedback; in practice, (47 - 48) speed up phase evolution in case of non-null force readings while legs are in swing, and vice-versa if no force is measured during stance, effectively synchronizing y^{CPG} (28 - 33) with ground contact. This represents a form of dynamic entrainment mirroring the behavior of biological gait [Hogan & Sternad,

³³With a direct equivalence between $O_{i,k}$ and goal state g in (3) and (11).

³⁴Which can be viewed as *primitive* elements of the resulting patterns.




Figure 10: Resulting kinematic trajectories y^{CPG} for all $K = 6$ DoFs of the right leg when employing (28 - 33) for a full period *i.e.* $[-\pi, \pi]$; Values of A^{motion} , $O_{i,j}$, ω and h presented in Table 1 (OP).

2013]. Similar closed-loop implementations also introduced corrections to ankle pitch and roll in an effort to ensure feet remain parallel to the ground at all times, increasing stability [Macedo et al., 2014; Macedo, 2014]:

$$\dot{O}_{i,\text{aPitch}} = K_0 \left(\sum_{j=1}^{N_{i,\text{front}}^{\text{FS}}} x_j^{\text{FS}} - \sum_{j=1}^{N_{i,\text{rear}}^{\text{FS}}} x_j^{\text{FS}} \right) \quad (49)$$

$$\dot{O}_{i,\text{aRoll}} = K_0 \left(\sum_{j=1}^{N_{i,\text{right}}^{\text{FS}}} x_j^{\text{FS}} - \sum_{j=1}^{N_{i,\text{left}}^{\text{FS}}} x_j^{\text{FS}} \right) \quad (50)$$

where ankle pitch offset $O_{i,\text{aPitch}}$ (32) is adjusted proportionally to K_0 ³⁵ and to the ratio of weight distribution on the robot feet *i.e.* increased whenever the frontal section outweighs the rear counterpart, and vice-versa; the same rationale is employed for the ankle roll offset $O_{i,\text{aRoll}}$ (33), using readings from sensors placed at left and right foot sections [Macedo, 2014].

2.4.2 Generic Legged Control

The CPG walking approach described has been extensively validated in previous studies [André et al., 2015; Macedo, 2014; Matos & Santos, 2014], demonstrating its ability to generate stable locomotion patterns in a real OP robot. However, the need for reparameterization - whether through hand-tuning or by resorting to optimization techniques - across different locomotion contexts limits its generalizability to other platforms and environments.

³⁵In practice, both force-based feedback gains K_F and K_O should be set to low values for stable behavior.




Figure 11: Kinematic overview of how CPG walking [Matos & Santos, 2012, 2014] is reproduced with the generic MP-based control strategy, integrating Force Sensor (FS) feedback loops into DMP dynamics; α , β and, notably, $f_k^*(\phi_i)$, $i \in \{\text{left, right}\}$ are shared between both legs; both hip yaw joints are mechanically coupled in the NAO; lastly, N_i^{FS} evenly distributed FSs measuring GRFs in both feet are assumed.

Alternatively, by encoding a combination of any arbitrary number of RBFs, the proposed generic strategy (2.3) offers greater flexibility and a broader range of applications, as well as potentially less dimensional when sharing a common primitive library $\mathbf{P}$. In addition to fully replacing (28 - 33) with DMP dynamics in (11) (2.3.1), replicating CPG walking requires adjusting three main aspects of the general strategy (2.3):

- **Preset goals states** in each leg DoF, defined as the hand-tuned offset values in (28 - 33); in all $K = 6$ leg DoFs, joint position follows (11) with $g_k = O_{i,k}$, $\forall k \in \{1, 2, \dots, K\}$ representing an initial baseline for control trajectories;
- **Coupled closed-loop canonical systems** driving the motions in each leg, by merging (6) and (47 - 48), thus integrating synchronization feedback and establishing inter-leg coordination³⁶:

$$\tau \dot{\phi}_{left} = 1 + K_C \sin \left(\dot{\phi}_{left} - \dot{\phi}_{right} + \pi \right) + K_F \sin \left(\sum_{j=1}^{N_{left}^{FS}} x_j^{FS} \right) \quad (51)$$

$$\tau \dot{\phi}_{right} = 1 + K_C \sin \left(\dot{\phi}_{right} - \dot{\phi}_{left} + \pi \right) + K_F \sin \left(\sum_{j=1}^{N_{right}^{FS}} x_j^{FS} \right) \quad (52)$$

- **Dynamic goal update** in ankle *pitch* and *roll* DoFs for stability, by repurposing the corrective feedback mechanism in (49 - 50), in line with similar goal switching methods [Ijspeert et al., 2013]:

³⁶The phase domain $\Phi = [-\pi, \pi]$ is implied.

$$\tau \dot{g}_{i,\text{aPitch}} = K_0 \left(\sum_{j=1}^{N_{i,\text{front}}^{\text{FS}}} x_j^{\text{FS}} - \sum_{j=1}^{N_{i,\text{rear}}^{\text{FS}}} x_j^{\text{FS}} \right) \quad (53)$$

$$\tau \dot{g}_{i,\text{aRoll}} = K_0 \left(\sum_{j=1}^{N_{i,\text{right}}^{\text{FS}}} x_j^{\text{FS}} - \sum_{j=1}^{N_{i,\text{left}}^{\text{FS}}} x_j^{\text{FS}} \right) \quad (54)$$

With this, beyond the appeal of MP theory highlighted in 2.3, motor skills in $\mathbf{S}$ are restricted to rhythmic, asymmetrical, bilateral leg movement, regulated in real time by GRFs measurements. By exploiting CPGs as a functional footprint, $\mathcal{S}^1$ encodes archetypal walking at a fundamental level; metaparameters $\{\tau, \mathbf{g}\}$ directly define walking velocity³⁷ and default joint offsets, which can be updated to cope with different terrain slopes [André et al., 2015]. Ultimately, this approach is extensible to other forms of legged locomotion e.g. running or hopping. Figure 11 summarizes the resulting kinematic control rules applied to leg DoFs within the control structure in Figure 9.

2.4.3 Cloning Walking Behavior

In a practical application of what has been discussed thus far, given CPG-based walking as reference control patterns (2.4.1), the motor skill learning procedure (2.3.3) was conducted to parametrize $\mathbf{P}$, $\mathbf{S}^{\text{OP}}$ and $\mathbf{S}^{\text{NAO}}$ (Figure 9) in order for the legged control formulation of the generic control strategy (Figure 11) to accurately reproduce $\mathbf{y}^{\text{CPG}}$ for a finite time duration.

To that end, a preliminary step involved empirically tuning CPG parameters $\{O_{i,k}, A^{\text{motion}}\}$ in the OP³⁸ and the NAO until similar stable forward walking on flat terrain was achieved. A single period³⁹ of the underlying open-loop right leg kinematics $\mathbf{y}_k^{\text{CPG}}$ for all $K = 6$ right leg DoFs, as well as the associated phase ϕ_{right} , was replicated with (28 - 33) resulting in fixed-length sequences $\mathbf{y}^{\text{CPG}} \in \mathbb{R}^T$ and $\phi_{\text{right}}^{\text{CPG}} \in \mathbb{R}^T$. Reference time scaling τ^{CPG} was estimated⁴⁰ directly from walking frequency ω (45) i.e. $\tau^{\text{CPG}} \simeq \omega^{-1}$.

θ^{best} was then found iteratively with BBO - in the form of CMA-ES - by minimizing $J(\theta, \mathbf{y}^{\text{CPG}})$ (23 - 24). As illustrated in Figure 12, this was conducted in two separate stages; initially in the OP through 1) expansive simultaneous parameterization of $\mathbf{P}$ and $\mathbf{S}^{\text{OP}}$; and subsequently in the NAO, given $\mathbf{P}$, through 2) intrinsic learning of $\mathbf{S}^{\text{NAO}}$. For each candidate θ , given $\{\mathbf{y}^{\text{CPG}}, \phi_{\text{right}}^{\text{CPG}}, \tau^{\text{CPG}}\}$, an experimental trajectory $\mathbf{y}^\theta$ was constructed by updating DMP state from fixed initial reference values, as described in Algorithm 1.

In its turn, $\mathbf{P}$ was populated *a priori* with P rhythmic DMPs, parametrized with a fixed set of N homogeneous kernels⁴¹ $\Psi_n \in \mathbb{R}^N$ according to Table 2, whose weighting $\mathbf{W}$ (10) was only learned during the initial expansive learning in the OP. $T = 150$ (through a fixed integration step of $\Delta t = 0.008$ seconds) and $N = 10$ were hand-tuned through trial-and-error to fit CPG reference trajectories. Considering the $K = 6$ leg DoFs of both platforms, dimensionality was $n_{\text{exp}} = 22P$ in the first stage, subsequently

³⁷Within reasonable boundaries; functional walking patterns change topologically at different velocities [Barroso, 2016].

³⁸Available in the literature for the OP [André et al., 2016a,b, 2015; André et al., 2014; Matos & Santos, 2012, 2014].

³⁹ ϕ_{right} sweeping the full phase domain Φ i.e. linearly evolving from $-\pi$ to π .

⁴⁰Eliminating the need to also parametrize τ during learning, and accommodating $\mathbf{y}^{\text{REF}}$ of any length.

⁴¹Fixed width and evenly spaced over phase domain $\Phi = [-\pi, \pi]$, according to (7) and (8).

Algorithm 1 Pseudo-code describing the computation of $J(\theta, \mathbf{y}^{\text{REF}})$ during optimization, with $\mathbf{y}^{\text{REF}} = \mathbf{y}^{\text{CPG}}$ described by CPG dynamics (28 - 33); As depicted in Figure 12, step #4 is only conducted in the initial learning stage, with $\mathbf{P}$ reused in the subsequent intrinsic learning procedure through constant $\mathbf{W}_0$. Disregarding masking operation in #9, complexity follows $\mathcal{O}(KT)$.

```

1: procedure COST ( $\theta, \mathbf{y}^{\text{REF}}, \phi^{\text{REF}}, \tau^{\text{REF}}$ )
2:    $\{\mathbf{W}^\theta, \Lambda^\theta, \Delta^\theta\} \leftarrow \theta$ 
3:   if  $\mathbf{W}^\theta \neq \mathbf{W}_0$  then ▷ Figure 12
4:      $\mathbf{P}^\theta \leftarrow \text{PARAMETRIZE}(\mathbf{W}^\theta)$  ▷  $\forall p \in \{1, 2, \dots, P\}$ 
5:    $\mathcal{S}^\theta \leftarrow \{\Lambda^\theta, \Delta^\theta\}$ 
6:   for each  $k \in \{1, 2, \dots, K\}$  do
7:      $y_{k,1}^\theta \leftarrow y_{k,1}^{\text{REF}}$ 
8:     for each  $t \in \{2, 3, \dots, T\}$  do
9:        $f(\phi_t^{\text{REF}}) \leftarrow \lambda_k^{*\theta^\top} \mathbf{m} \mathbf{P}^\theta (\phi_t^{\text{REF}} + \delta_k^{*\theta})$  ▷ (18)
10:       $\ddot{y}_{k,t}^\theta \leftarrow (\tau^{\text{REF}})^{-2} (\alpha_{y,k} (\beta_{y,k} (g_{\text{right},k} - y_{k,t-1}^\theta - \dot{y}_{k,t-1}^\theta) + f(\phi_t^{\text{REF}})))$  ▷ (11)
11:       $\{y_{k,t}^\theta, \dot{y}_{k,t}^\theta\} \leftarrow \text{UPDATE} (y_{k,t-1}^\theta, \dot{y}_{k,t-1}^\theta, \ddot{y}_{k,t}^\theta, \tau^{\text{REF}})$  ▷ e.g. Euler method
12:       $J_k(\theta, \mathbf{y}_k^{\text{REF}}) \leftarrow \sum_{t=1}^T (y_{k,t}^\theta - y_{k,t}^{\text{REF}})^2$ 
13:     $J(\theta, \mathbf{y}^{\text{REF}}) \leftarrow \sum_{k=1}^K J_k(\theta, \mathbf{y}_k^{\text{REF}})$  ▷ (23)
14:  return  $J(\theta, \mathbf{y}^{\text{REF}})$ 

```

decreasing to $n_{\text{int}} = 12P$ during the intrinsic learning stage (Figure 7 and 12). In the case of CMA-ES, the computational complexity of each update generally follows $\mathcal{O}(n^3)$ [Hansen et al., 2003]; as n grows linearly with P , in order to identify a reasonable trade-off between trajectory resolution and resulting overhead, the learning procedure in the OP was repeated 3 times for each $P \in \{1, 2, 3, 4, 5, 6\}$. In both robots, all DMP transformation systems were critically damped with $\{\alpha_{y,k}, \beta_{y,k}\} = \{25, 6.25\}$.

For practicality and simplicity⁴², CMA-ES was hand-tuned with $\lambda^{\text{BBO}} = 50$ candidates per generation and $\sigma_0^{\text{BBO}} = 100$ as the initial exploration level (Table 3), fixed for all P in both robots; these parameters were intentionally overestimated⁴³ in order foster early exploration and prevent premature convergence to local optima. Individual basis weights w_n (4) and elements in Λ were constrained to $[-1000, 1000]$, while Δ was constrained to $[-\pi, \pi]$ (15). Moreover, given the morphological similarity between both platforms, $\mathbf{S}^{\text{NAO}}$ is initialized with $\mathbf{S}^{\text{OP}}$ during the second learning stage (Figure 12) for faster convergence. Optimization results were objectively evaluated through $J(\theta^{\text{best}}, \mathbf{y}^{\text{CPG}})$ and total convergence time I^{conv} ; subsequently, learned DMP control patterns $\mathbf{y}^{\text{DMP}} = \mathbf{y}^{\theta^{\text{best}}}$ were validated in walking trials in a simulated environment, and the resulting gait profiled in terms of velocity, stability and energetic efficiency. For a detailed overview of the CMA-ES algorithm, the reader is referred to Hansen [2016] and Hansen & Ostermeier [2001]. A description of the implementation developed in this work, including convergence criteria considered during optimization, is provided in Appendix A: Practical Implementation.

⁴²Optimal learning performance was not the primary focus of this study; due to the multimodal nature of the parameter set, a more careful tuning process would be required in real-world applications, guided by technical guidelines and practical rules of thumb e.g. parameter normalization and scaling λ^{BBO} with n [Hansen, 2016; Hansen et al., 2003].

⁴³At the cost of computational overhead and convergence time.




Figure 12: Conceptual diagram describing the two-stage procedure employed to teach bipedal walking to both robots; in the first, expansive learning (2.3.3) is used to populate $\mathbf{P}$ and $\mathbf{S}^{\text{OP}}$; in the latter, intrinsic learning is used to transfer skill parameterization to $\mathbf{S}^{\text{NAO}}$, reusing MPs in $\mathbf{P}$. In both robots, $K = 6$.

2.5 Results

2.5.1 Motor Skill Learning

Table 1 presents the hand-tuned CPG parameters obtained during the preliminary stage. In practice, parameterization was sourced from previous research [André et al., 2016a,b, 2015; André et al., 2014; Matos & Santos, 2012, 2014] and further tuned through trial-and-error to generate stable forward displacement in flat terrain; the resulting CPG gait used as reference for motor skill learning was, in terms of total distance and average velocity, similar in both robots (2.5.2). In the specific case of the OP, the simulated model was found to be slightly asymmetrical, which caused a slight deviation to the left with each step during the *balancing* motion (as illustrated in previous works [André et al., 2016b; Matos & Santos, 2014]), proportional to $A^{\text{balancing}}$; a practical workaround was found in defining an offset of 2 degrees⁴⁴ on both roll DoFs (hRoll and aRoll). Generally, values of A^{motion} were higher in the NAO than in the OP, which can be explained by different physical dimensions *i.e.* leg length of both robots (Appendix B: Humanoid Platforms). Moreover, in the NAO, it was found that walking performance and stability improved when completely disabling the *rotation* motion: the described CPG control patterns (2.4.1) assume independently controlled DoFs, while the mechanical coupling between left and right hYaw in the NAO enforces symmetrical movement. As such, $A^{\text{rotation}} = 0$ in the NAO (Table 1).

Figure 13 and 14, illustrate, respectively, the CPG cloning process for different number of MPs $P \in \{1, 2, 3, 4, 5, 6\}$, as well as average final cost $J(\theta^{\text{best}}, \mathbf{y}^{\text{CPG}})$ and measured Variance Accounted For (VAF) between reference and resulting DMP trajectories, across 3 repetitions; the latter, in particular, represents an objective metric to evaluate the resulting trajectories, which - in contrast with cost - is not dependent on the length and scale of the CPG references (23). Table 4 details both metrics, and presents

⁴⁴Applied symmetrically in hRoll and aRoll in order to preserve stability *i.e.* $O_{\text{hRoll}} = -O_{\text{aRoll}} = 2$ (Table 1).

Table 1: Hand-tuned CPG parameters (28 - 33); same parameterization used in both legs;

Parameter	OP	NAO
α	10	
ω	4.9	
h	$\pi/6$	
$A^{\text{balancing}}$	11	-25
$A^{\text{flexion}}_{\text{hip}}$	20	60
$A^{\text{flexion}}_{\text{knee}}$	40	120
A^{compass}	11	22
A^{yield}	2	1
A^{rotation}	10	0
O_{hPitch}	-20	-35
O_{hRoll}	2	0
O_{hYaw}	0	0
O_{kPitch}	40	50
O_{aPitch}	-20	-25
O_{aRoll}	-2	0
K_F	$1 \cdot 10^{-2}$	$5 \cdot 10^{-2}$
K_C	$5 \cdot 10^{-3}$	
K_O	$2.5 \cdot 10^{-3}$	

Table 2: DMP parameters employed in the generic control strategy (2.4.2).

Parameter	OP	NAO
α_y	25	
β_y	6.25	
τ^{REF}	0.19	
c_n	$n \frac{2\pi}{N} - \pi$ (7)	
h_n	0.30	
T	1000	
N	10	
P	1-6	

Table 3: CMA-ES parameters employed in the motor skill learning procedure (2.4.3).

Parameter	OP	NAO
λ^{BBO}	50	
σ_0^{BBO}	100	

estimates of the error per sample $\hat{e}_y = 1/KT \times J(\theta^{\text{best}}, \mathbf{y}^{\text{CPG}})$ in each case. In both robots, the quality of the resulting solution - as evidenced in a lower cost and higher VAF - increased with P ; on the other hand, the convergence time I^{conv} also increased significantly with more MPs, a natural result of higher dimensionality n , which, as presented (2.4.3), varies linearly with P i.e. $22P$ in the OP and $12P$ in the NAO; accordingly, both convergence times and initial costs were noticeably smaller in the NAO. Despite similar trends observed in both robots, lower final costs and higher VAF values were consistently obtained in the OP, hinting at a potential overfit of the MPs learned during the initial expansive learning stage. Notwithstanding, accurate approximations were found in both cases when $P > 1$, suggesting that at least 2 MPs are required to replicate CPG patterns with the control strategy proposed in this work. In the specific case of the OP, $\hat{e}_y \leq 0.08^\circ$ for all $P > 1$ (Table 4), while this value is slightly higher in the NAO at $\hat{e}_y \leq 1^\circ$, measured VAF is above 0.80 for all $P > 1$.

In the OP, as depicted in Figure 15, $P = 3$ MPs were found to suffice to accurately replicate reference CPG patterns, with a near complete overlap between reference and learned control patterns; in the case of the NAO, however, slight deviations between both were evident (Figure 16), which decreased for higher values of P , in accordance with measured cost and VAF values (Figure 14). Notably, due to the *rotation* motion being disabled in the NAO (Table 1), the Hip Yaw DoF is fixed at neutral position.

The resulting set of $P = 3$ MPs is presented in Figure 17, along with learned robot-specific motor skill $\mathcal{S}^1 = \{\mathbf{\Lambda}, \mathbf{\Delta}\}$ composing $\mathbf{S}^{\text{OP}}$; similarly, analogous motor skills learned in the NAO with $P \in \{3, 4, 5\}$, integrating $\mathbf{S}^{\text{NAO}}$, are presented in Figure 18. While the shape of individual MPs is not particularly




Figure 13: Evolution of $J(\theta, \mathbf{y}^{\text{CPG}})$ during skill learning (2.4.3), for $P \in \{1, 2, 3, 4, 5, 6\}$, in the OP (**Top**) and the NAO (**Bottom**); final values presented in Figure 14 and Table 4.




Figure 14: Average final optimization cost $J(\theta^{\text{best}}, \mathbf{y}^{\text{CPG}})$ (**Left**) and Variance Accounted For (VAF) between DMP learned trajectories $\mathbf{y}^{\text{DMP}}$ and reference $\mathbf{y}^{\text{CPG}}$ (**Right**), with $P \in \{1, 2, 3, 4, 5, 6\}$ (Table 4), across 3 repetitions of the skill learning procedure (2.4.3); error bars illustrate measured standard deviation.

Table 4: Average $J(\theta^{\text{best}}, \mathbf{y}^{\text{CPG}})$ and Variance Accounted For (VAF) (Figure 14), as well as estimated error per sample $\hat{e}_y = 1/KT \times J(\theta^{\text{best}}, \mathbf{y}^{\text{CPG}})$, across 3 repetitions of the skill learning procedure (2.4.3).

P	OP			NAO		
	$J(\theta^{\text{best}}, \mathbf{y}^{\text{CPG}})$	$\hat{e}_y$	VAF	$J(\theta^{\text{best}}, \mathbf{y}^{\text{CPG}})$	$\hat{e}_y$	VAF
1	976.05 ± 80.66	1.08	0.88 ± 0.02	2311.84 ± 94.54	2.57	0.71 ± 0.03
2	68.71 ± 31.46	0.08	0.99 ± 0.00	900.14 ± 236.38	1.00	0.80 ± 0.01
3	22.79 ± 7.25	0.03	1.00 ± 0.00	525.37 ± 78.23	0.58	0.82 ± 0.01
4	17.00 ± 1.83	0.02	1.00 ± 0.00	424.37 ± 100.14	0.47	0.83 ± 0.00
5	5.56 ± 0.60	0.00	1.00 ± 0.00	256.24 ± 154.61	0.29	0.83 ± 0.01
6	3.26 ± 0.97	0.00	1.00 ± 0.00	189.27 ± 5.12	0.21	0.83 ± 0.00




Figure 15: Reconstructed DMP trajectories $\mathbf{y}^{\text{DMP}} = \mathbf{y}^{\theta^{\text{best}}}$ for each of the $K = 6$ DoFs in the OP, with $P = 3$, as well as reference CPG curves (Figure 10).






Figure 16: Reconstructed DMP trajectories $\mathbf{y}^{\text{DMP}} = \mathbf{y}^{\theta^{\text{best}}}$ for each of the $K = 6$ DoFs in the NAO, with $P \in \{3, 4, 5\}$, as well as reference CPG curves (Figure 10).


 Figure 17: Resulting set of $P = 3$ MPs and motor skill $\mathcal{S} = \{\Lambda, \Delta\}$ learned in the OP.

 Figure 18: Resulting motor skills $\mathcal{S} = \{\Lambda, \Delta\}$ learned in the NAO with $P = \{3, 4, 5\}$.

noteworthy - in itself an indicator of its abstract and generic nature - analysis of Λ and Δ can provide insights into their functional role in pattern generation; specifically, low $\lambda_{\text{hYaw}}^p, \forall p \in \{1, 2, \dots, P\}$ agree with the null amplitude of the rotation motion (Figure 16).

2.5.2 Dynamic Walking Performance

Both original and learned locomotion were validated in short walking trials with simulated OP and NAO robots in Gazebo [Koenig & Howard, 2004] (Figure 19); a total of 7 seconds of walking were reproduced, with experimental data from the last 5 seconds recorded for analysis. In practice, different symmetries were applied in order to align control patterns (Figure 10, 15 and 16) with DoF orientations in each robot; for consistency between both robots, results are presented according to International Society of Biomechanics (ISB) guidelines for reporting kinematic data of the lower limbs [Wu et al., 2002]. Moreover, with reproducibility in mind, noise seeds⁴⁵ are reused in CPG and DMP walking.

Figures 20-25 provide a kinematic and dynamic profile of the hand-tuned CPG gait used as demonstration, namely control commands $y_{i,k}^{\text{CPG}}$ and tracked state $q_{i,k}^{\text{CPG}}$ ⁴⁶, DoF-specific phase diagrams of $q_{i,k}^{\text{CPG}}$, leg-specific phase ϕ_i and total GRF measurements $\sum x_j^{\text{FS}}$, as well as spatial coordinates of the robot's CoM over time $\{x_{\text{CoM}}, y_{\text{CoM}}, z_{\text{CoM}}\}$; similarly, Figures 26-31 depict analogous information for the resulting

⁴⁵Influencing electrical and measurement noise during simulation.

⁴⁶Control commands constituting setpoints for low-level Proportional-Integral-Derivative (PID) feedback controllers. PID parameterization in both robots provided in Appendix B: Humanoid Platforms.




Figure 19: Snapshots of the simulated OP and NAO bipeds walking on flat ground.

DMP walking achieved with the generic legged control strategy described (2.4.2) by prescribing $y_{i,k}^{\text{DMP}}$ to each leg DoF, as parametrized with the motor skill learning procedure (2.4.3).

In all trials, prescribed trajectories were accurately tracked, leading to stable behavior in all DoFs, as is demonstrated through the closed loops in the phase diagrams in Figure 22-23 and Figure 28-29. The resulting displacement of the robot occurred mostly in the frontal direction, with alternating left/right foot strikes, visually correlated with an increase in total GRF *i.e.* each strike led to a noted peak in measured force. On the other hand, as can be seen in Figure 25 and 31, the effect of the force-based feedback mechanism (2.4.1) was more prevalent in the NAO, despite both robots recording similar GRF ranges throughout the trials. In the specific case of DMP walking in the NAO, transient instability was observed between $t = 4$ and $t = 5$ seconds (Figure 31), evidenced in higher GRF peaks and irregularity in CoM position; in the remainder of the experiment, CoM movement was shown to be regular and periodic.

Table 5 presents objective metrics comparing hand-tuned CPG and learned DMP locomotion in each robot, including distance traveled (*i.e.* total forward displacement), average and maximum walking velocity, total energy expenditure, gait duty factor and practical estimation of the Froude number:

$$\text{Fr} = \frac{\bar{v}^2}{gL} \quad (55)$$

based on the average walking velocity $\bar{v}$, as estimated at the robot's CoM, with L the robot's leg length; as well as the effective CoT or the duration of the walking trial:

$$\text{CoT} = \frac{\bar{P}}{mg\bar{v}} \quad (56)$$




Figure 20: DoF position control commands y_k^{CPG} (Figure 10) and resulting state q_k^{CPG} values in the OP during CPG walking on flat ground. Each DoF is presented according to ISB joint coordinate systems.




Figure 21: DoF position control commands y_k^{CPG} (Figure 10) and resulting state q_k^{CPG} values in the NAO during CPG walking on flat ground. Each DoF is presented according to ISB joint coordinate systems. Hip Yaw is omitted due to being fixed at 0° (Figure 16).




Figure 22: Phase diagrams of leg DoFs in the OP recorded during CPG walking on flat ground.




Figure 23: Phase diagrams of leg DoFs in the NAO recorded during CPG walking on flat ground. Hip Yaw is omitted due to being fixed at 0° (Figure 16).




Figure 24: Leg Phase ϕ_i and total GRF measured $\sum x_j^{\text{FS}}$ (**Left**), as well as CoM position $\{x_{\text{CoM}}, y_{\text{CoM}}, z_{\text{CoM}}\}$ and foot strike locations (**Right**) in the OP during CPG walking on flat ground.




Figure 25: Leg Phase ϕ_i and total GRF measured $\sum x_j^{\text{FS}}$ (**Left**), as well as CoM position $\{x_{\text{CoM}}, y_{\text{CoM}}, z_{\text{CoM}}\}$ and foot strike locations (**Right**) in the NAO during CPG walking on flat ground.




Figure 26: DoF position control commands y_k^{DMP} (Figure 15) and resulting state q_k^{DMP} in the OP values during DMP walking on flat ground. Each DoF is presented according to ISB joint coordinate systems.




Figure 27: DoF position control commands y_k^{DMP} (Figure 16) and resulting state q_k^{DMP} in the OP values during DMP walking on flat ground. Each DoF is presented according to ISB joint coordinate systems. Hip Yaw is omitted due to being fixed at 0° (Figure 16).




Figure 28: Phase diagrams of leg DoFs in the OP recorded during DMP walking on flat ground.




Figure 29: Phase diagrams of leg DoFs in the NAO recorded during DMP walking on flat ground. Hip Yaw is omitted due to being fixed at 0° (Figure 16).




Figure 30: Leg Phase ϕ_i and total GRF measured $\sum x_j^{\text{FS}}$ (**Left**), as well as CoM position $\{x_{\text{CoM}}, y_{\text{CoM}}, z_{\text{CoM}}\}$ and foot strike locations (**Right**) in the OP during DMP walking on flat ground.




Figure 31: Leg Phase ϕ_i and total GRF measured $\sum x_j^{\text{FS}}$ (**Left**), as well as CoM position $\{x_{\text{CoM}}, y_{\text{CoM}}, z_{\text{CoM}}\}$ and foot strike locations (**Right**) in the NAO during DMP walking on flat ground.

Table 5: Gait metrics obtained during CPG and DMP simulated walking with both robots.

	OP			NAO		
	CPG	DMP	+/- (%)	CPG	DMP	+/- (%)
Distance (m)	0.675	0.655	-2.97	0.658	0.647	-1.58
Avg. Velocity (m/s)	0.115	0.115	+0.20	0.101	0.105	+3.49
Total Energy Expenditure (J)	102.234	94.002	-8.05	223.168	272.366	+22.05
Average Power (P)	20.537	18.910	-7.92	44.678	54.648	+22.31
Cost of Transport	5.327	5.048	-5.24	6.310	7.825	+24.01
Froude Number	0.006	0.006	+0.41	0.004	0.005	+7.10
Double Support (%)	25	21	-16.00	13	15	+15.38
Single Support (%)	75	79	+5.33	87	85	-2.30
Duty Factor	0.613	0.605	-1.30	0.558	0.565	+1.30
Average (Abs.)			4.76			9.97

where $\bar{P}$ is the average power over the walking trial, and m the robot's total weight. Per convention, g denotes the gravitational acceleration in both (55) and (56). Documented values of m and L for the OP and the NAO are provided in Appendix B: Humanoid Platforms.

Overall, both robots traveled roughly 0.65-0.67 meters at velocities of approximately 0.10-0.12 meters per second, exhibiting duty factors between 0.56 and 0.61, in line with the characteristic values in human locomotion [Matos & Santos, 2014]; however, walking in the NAO had a noticeably higher energetic footprint than in the OP. On one hand, this was to be expected as the former is a larger and heavier robot, and higher amplitude in control trajectories (Figure 16) imply higher torques; on the other, CoT values were relatively high in both robots, particularly when compared to reported values in the literature for bipedal platforms [Ijspeert, 2014; Mikolajczyk et al., 2022; Seok et al., 2015], highlighting the significant energetic cost of the resulting gait. In all trials, Froude number values were very low, which is in agreement with the slow locomotion velocities exhibited by both robots.

Regarding the quality of replicating the original CPG patterns with the generic control strategy, on average, locomotion in the NAO exhibited higher variation between CPG and DMP locomotion, notably a 22% increase in energetic consumption and average power and 24% in CoT, along with a decrease in distance traveled and maximum velocity; this was, to an extent, expected when taking into the account the lower-quality approximations obtained during the skill learning procedure, and transient instability depicted in Figure 31. By contrast, DMP walking in the OP was closer to the original hand-tuned CPG approach, registering the same average velocity (0.115 m/s) and very similar displacement (0.675 vs 0.655 m), with a decrease of 7-8 % in energetic consumption. Globally, gait metrics varied an average of 4.76% in the OP and 9.97% in the NAO.

2.6 Discussion

The practical application of the generic control strategy (2.3) to the legged control of multiple humanoid robots, specifically within the context of bipedal walking on flat terrain (2.4), demonstrated how a shared task- and platform-agnostic MP library can encode a sufficiently rich motor repertoire to replicate, with reduced complexity, dynamic and high-dimensional motor skills with considerable success and accuracy; the results achieved, both during skill learning and subsequent reproduction of learned behaviors in a simulated environment, attest to the flexibility and general appeal of MP-based control while highlighting the valuable role of MPs can have in the simplification and dimensionality reduction of motor generation and control in robotics.

The hand-tuned (Table 1) closed-loop CPG control patterns used as reference in motor skill learning (2.4.3), illustrated in Figures 20-25, were found to be regular, periodic, and efficiently tracked by low-level linear feedback (PID) control in both simulated robot models; however, as this process prioritized stability over performance, the resulting walking velocity was low, directly contributing to low Fr and high⁴⁷ CoT values *i.e.* slow, quasi-static and inefficient locomotion (Table 5); moreover, CoM path is tied to robot

⁴⁷CoT around 5-5.3 for the OP and 6.3-7.8 in the NAO (Table 5) imply more than 5 times the gravitational energy required to move an object of similar mass over the same distance.

morphology and dimensions; in the case of the NAO, it can be seen that a more angular trajectory was obtained, albeit still periodic, with a slight deviation to the left (Figure 25), which may suggest excessive acceleration during stepping motions. Notwithstanding, it should be noted the primary goal was obtaining similar functional control trajectories in both robots, and not in achieving highly-performant gait; to that end, a more comprehensive preliminary tuning process is warranted, potentially leveraging automated ML techniques that integrate walking trials and optimize velocity and dynamic stability metrics, and achieving Fr values characteristic of dynamic walking *i.e.* $0.2 \gtrsim Fr \lesssim 0.5$ [Alexander, 1991; Mikolajczyk et al., 2022]. Conversely, this discussion focuses on how well reference walking patterns are mimicked by the implemented MP control, rather than objective walking performance. As noted in 2.4.1, the CPG walking approach has been validated in previous works [André et al., 2015; Matos & Santos, 2014]. Ultimately, learned gait can provide a marginally stable behavior baseline for subsequent real-time tuning strategies.

The results of the motor skill learning procedure (2.5.1) show that the MP-based control strategy (2.3) was able to accurately replicate the reference CPG walking patterns under different configurations, namely with different sizes of primitive library $\mathbf{P}$, in both the OP and the NAO. Moreover, as anticipated, in the second learning stage - involving intrinsic parameterization of $\mathbf{S}^{\text{NAO}}$ - solutions were found significantly faster (Figure 13), below 300 k generations for all P , in contrast to roughly 100 k to 2 M in the initial learning stage of $\mathbf{P}$ and $\mathbf{S}^{\text{OP}}$. Generally, convergence time increased with P , which was, to an extent, expected: in the specific case of CMA-ES, convergence has been empirically estimated to be reached in approximately $O(n \log n)$ generations [Hansen et al., 2003; Hansen & Ostermeier, 2001]. This does not, however, accurately describe the actual increase in computational effort with P , which follows $\mathcal{O}(n^3)^{48}$ and therefore has a greater impact than illustrated. In addition, initial cost values were also observed to be greatly reduced in the NAO, explained by both the decreased dimensionality of the optimization problem ($n = 12P = 36$) and by the initialization the motor skill $\mathbf{S}^{\text{NAO}}$ to previously learned $\mathbf{S}^{\text{OP}}$. This highlights the appeal of acquiring new motor skills through intrinsic learning, which is facilitated the larger and more diverse $\mathbf{P}$ is; MP-based control thus gains additional relevance under a Lifelong Learning paradigm, as exploring new behaviors becomes incrementally more efficient.

Similarly, the quality of the resulting approximation, as objectively quantified by a lower final cost $J(\boldsymbol{\theta}^{\text{best}}, \mathbf{y}^{\text{CPG}})$, also increased with P (Figure 14); while the optimization cost is dependent on reference scale and length T , and thus may not convey what exactly represents an *accurate* reproduction of $\mathbf{y}^{\text{CPG}}$, a similar trend was observed in the VAF between CPG and DMP trajectories *i.e.* higher values implying a better approximation. This aligns with the underlying formulation (2.3.2), illustrating that a larger $\mathbf{P}$ encodes greater motor variability and therefore provides greater flexibility when combining multiple MPs to match reference patterns. In the OP, a total of $P = 3$ MPs were sufficient to accurately replicate the reference motions, as illustrated in Figure 15. On the other hand, for all values of P , final cost $J(\boldsymbol{\theta}^{\text{best}}, \mathbf{y}^{\text{CPG}})$ and trajectory VAF values were noticeably worse in the NAO than in the OP (Figure 14). Beyond unavoidable platform specificities that may cause the *a priori* solution $\mathbf{S}^{\text{OP}}$ to guide optimization towards suboptimal parameterization and/or premature convergence - unlikely when considering the morphological similarity between both platforms, but which could be escaped by increasing CMA-ES parameters λ^{BBO} and σ_0^{BBO}

⁴⁸Globally, optimization is expected to follow $\mathcal{O}(n^4 \log n)$.

(Table 2) - this trend suggests that $\mathbf{P}$ may have been overfit to the OP robot and consequently was not comprised of completely platform-invariant elemental MPs; in order to circumvent this phenomenon and ensure the general nature of the learned MPs, an alternative skill learning approach would be to perform compound *one-shot* optimization of $\mathbf{P}$, $\mathbf{S}^{\text{OP}}$ and $\mathbf{S}^{\text{NAO}}$:

$$J^*(\boldsymbol{\theta}, \mathbf{y}_{\text{OP}}^{\text{CPG}}, \mathbf{y}_{\text{NAO}}^{\text{CPG}}) = J(\boldsymbol{\theta}, \mathbf{y}_{\text{OP}}^{\text{CPG}}) + J(\boldsymbol{\theta}, \mathbf{y}_{\text{NAO}}^{\text{CPG}}) \quad (57)$$

where $\mathbf{y}_{\text{OP}}^{\text{CPG}}$ and $\mathbf{y}_{\text{NAO}}^{\text{CPG}}$ denote CPG reference trajectories for the the OP and the NAO respectively; by combining both previously learning steps into a single optimization problem, dimensionality increases to $n = 22P + 12P = 32P$. A similar alternative would be, during the initial expansive learning of $\mathbf{P}$, generating or recording a distinct set of D reference trajectories:

$$J^*(\boldsymbol{\theta}, \mathbf{y}_1^{\text{CPG}}, \mathbf{y}_2^{\text{CPG}}, \dots, \mathbf{y}_D^{\text{CPG}}) = \sum_{d=1}^D J(\boldsymbol{\theta}, \mathbf{y}_d^{\text{CPG}}) \quad (58)$$

which would preserve the lower dimensionality of intrinsic learning in the NAO, albeit, much like the previous solution, at the cost of a less straightforward initial learning stage; one other approach entails the regularization of f^p through extrinsic parameters $\mathbf{W}$ (2.3.2), in an effort to better isolate the platform-invariant MPs from platform-specific parameterization:

$$J^*(\boldsymbol{\theta}, \mathbf{y}_{\text{OP}}^{\text{CPG}}) = \alpha_1 J(\boldsymbol{\theta}, \mathbf{y}_{\text{OP}}^{\text{CPG}}) + \alpha_2 (\|\mathbf{W}\|_2)^2 \quad (59)$$

which would require, as previously discussed (2.3.3), careful tuning of α_1 and α_2 . Lastly, constraining $\mathbf{W}$ to a normalized range e.g. $[-1, 1]$ would also ensure it remains scale-invariant across different platforms.

Notwithstanding, despite localized deviations, the resulting DMP trajectories in the NAO still constituted reasonably accurate reproductions of the CPG demonstrations for $P > 2$, as can be seen in Figure 16 in the case of $P \in \{3, 4, 5\}$. Conversely, while in the case of $P = 2$, a value of $\hat{e}_y = 1^\circ$ per sample (Table 4) may be acceptable under specific contexts or tasks requirements, establishing a functional threshold solely from the results of the the learning procedure is not straightforward, as it is not clear to what degree the stability and performance of the resulting gait will be impacted.

Generally, from what is presented in Figure 14 and Table 4, it can be concluded that a single MP was unable to accurately describe the reference patterns; this is in line with the CPG formulation (2.4.1), which employs a pair of distinct RBFs (26 - 27) as basic building blocks underneath leg control patterns. On the other hand, during the initial expansive learning stage, effective dimensionality reduction occurs if $P \leq 2$ (2.3.2). While trajectories combining only $P = 2$ MPs achieved both high $\text{VAF} = 0.99$ and negligible $\hat{e}_y = 0.08^\circ$ in the OP, the higher costs and lower VAF obtained in the NAO advocate for the use of more MPs, even if at the cost of motor redundancy; a practical compromise between dimensionality, generality and accuracy in both robots was found at $P = 3$. In fact, as evident in the motor skills learned in the NAO (Figure 18), for $P = 4$, f^1 had generally lower gains than the remaining MPs; in a similar fashion, for $P = 5$, this was observed in f^1 and f^5 , reinforcing the notion that $P = 3$ MPs suffice to fully

represent $\mathbf{y}^{\text{CPG}}$, even moreso when considering there is no rotation motion in the NAO (Table 1) and only 5 DoF require non-null control commands. Overall, with $P = 3$, not only is the difference in dimensionality marginal when compared to the conventional strategy of DoF-specific MPs ($22P = 66$ in the OP and $12P = 36$ in the NAO in contrast with $KN = 60$), it is also understood that expansive learning is not frequently required if $\mathbf{P}$ provides a sufficiently diverse behavior range, in line with the previous remarks about continually growing motor libraries in L2 systems. Lastly, $P = 3$ aligns with the general remarks from Hogan & Sternad [2013] regarding the primitive structure of biological walking.

Regarding the performance of the resulting gait, the MP control strategy was able to reproduce, to a considerable extent, the reference CPG walking in both robots. Globally, a very high degree of similarity was achieved in the OP, with selected gait metrics deviating less than 5% in magnitude between CPG and DMP walking (Table 5). In contrast, the DMP reproduction in the NAO exhibited some degree of instability during locomotion, characterized by a slightly erratic CoM path with increased vertical displacement⁴⁹, higher GRF and in-place stepping (Figure 31); as prescribed motions y_k^{DMP} are still consistently tracked throughout the walking trial (Figure 27), this is possibly a result of the suboptimal approximation obtained in the learning stage, exacerbated by inertial and gravitational forces acting upon the NAO. The lack of the rotation motion may have also played a role in this behavior, as its main purpose is shifting the swinging foot forward, stabilizing CoM vertical movement and contributing to an increase in step length [Matos & Santos, 2014]. Additionally, considering the significant difference in robot weight (2.9 vs 5.3 Kg), similar ranges of $\sum x_j^{\text{FS}}$ in both robots (Figures 24-25 and 30-31) were not expected, but can be explained by model inaccuracies *i.e.* force sensors may be positioned or distributed differently in each robot⁵⁰, implying a practical need for retuning K_C , K_F and K_O (2.4.2) in real world applications; in practice, the higher K_C value used in the NAO (Table 1) leads to a visible impact of ground contacts in phase ϕ_i .

DMP walking in the NAO reported an above 20% increase in energetic consumption and estimated CoT when compared to reference CPG walking (Table 5). Nevertheless, regular forward displacement was still evident, with only a small decrease in total distance traveled and similar increases in average velocity and duty factor values over the whole duration of walking trial; on average, gait metrics exhibited 9.97% variation from reference CPG walking, objectively representing over 90% similarity with reference gait. From all that was discussed in this section, this increase is seen not as a result of the proposed control approach but rather on the conservative approach followed during the preliminary tuning stage; instead, it can be concluded that, using measured variations of $<10\%$ in both robots as a KPI, the generic MP strategy was able to accurately imitate CPG locomotion in multiple humanoid platforms. Conceptually, the results obtained allow formulating a concrete answer to **RQ1**: structuring leg control on top of a general task- and platform-agnostic primitive library and employing IL techniques to imitate bipedal walking demonstrations allows the accomplishment of functional locomotion in different humanoid platforms.

⁴⁹Also influenced by the increased dimensions of the NAO and A^{motion} (Table 1), when compared to the OP.

⁵⁰Force-Sensing Resistors (FSR) were added the feet of the sensorless simulation model of the NAO used in this work.

2.7 Conclusion

This chapter describes a generic open-loop control strategy exploiting the notion of Motor Primitives (MPs) to simplify and generalize motor generation and skill acquisition. Extending previous research, DMP systems are used to encode stable attractor dynamics that prescribe open-loop commands to multiple DoFs, modulated by a combination of generic and abstract MPs providing flexibility in motor generation, and taking advantage of the several appealing features of the DMP framework. Conceptually, a general task- and platform-agnostic behavior library is defined containing all general-purpose primitive elements, with platform-specific motor skill libraries encoding movement policies by scaling and phase-shifting individual MPs in each available DoF. Beyond the 1) general and abstract nature of the MP library, making it implicitly generalizable, this structure/approach 1) simplifies skill acquisition and refinement by reducing the dimensionality of motor skills; 2) does not impose any control modality *i.e.* position or torque control, task or joint-space; 3) is not tied down to the actual dynamics *i.e.* fits most DMP extensions and discrete and rhythmic variants; 4) discards the need for detailed whole-body kinematic and dynamic models, allowing modular usage *e.g.* for each limb in humanoids; and 5) accommodates for a number of closed-loop control mechanisms, as well contextual skill definition in behavior generalization.

As a practical demonstration of the proposed approach, this control structure was applied to the legged control of two humanoid robots, the Robotis OP 1 and the SoftBank Robotics NAO V5, within the context of bipedal locomotion, whereby functional walking patterns generated with a previously validated bioinspired approach based on Central Pattern Generators (CPGs) are taught to both robots in a multi-level IL/Behavior Cloning (BC) procedure, illustrating the flexibility of the control structure in a learning scenario, and the simplicity of acquiring and refining motor skills. Not only was skill learning straightforward in both robots, walking motions were successfully reproduced by both in a simulation environment, with selected gait metrics exhibiting over 90% similarity when compared to the original gait.

The results achieved demonstrate how *generic bioinspired MP-based control, exploiting non-linear dynamical systems for motor generation, can effectively be used to imitate complex, highly-dynamic and high-dimensional behaviors such as walking in multiple platforms*, providing an objective answer to **RQ1**; moreover, its generic nature - without any task-specific information integrated *a priori* into the control structure or planning stage - allows application in a myriad of robotic problems and aligns with the remarks by Sutton [2019] on leveraging learning over task-oriented design of artificial systems. Overall, it reinforces the appeal of primitive theory in the simplification of complex phenomena, namely articulated robotic control, and exhibits multiple traits of Lifelong Learning (L2) systems, such as transfer and adaptation of behaviors, exploiting task similarity, noise tolerance and resource efficiency (1.1) [Kudithipudi et al., 2022], in addition of the open-ended nature of both primitive and skill libraries, understood to grow incrementally as the agent explores new behaviors.

Within the context of this thesis, Objectives **01**, relating to learning MP combinations from demonstrations, and **03**, about generalizing locomotion to different robots with ML, were accomplished; in addition, achieving **02** was advanced with the introduction of a task- and platform-agnostic MP library in **P**. Extending existing research, the main contributions of this study consist of a MP-based controller employing

a generic motor library, validated in a high-dimensional and dynamic task such as 3D walking, and a practical method for learning functional bipedal walking by imitation.

On the other hand, this work did not address the generalization of goal states for each MP, which would be useful, for instance, when sequencing different primitives [Kulvicius et al., 2012; Pastor et al., 2009; Stulp et al., 2012]; along the same lines, reusing the same MP with different scales and phase shifts is not accounted for by the current strategy⁵¹. Moreover, throughout this work the number of primitives P' employed in a motor skill is treated as an input parameter of the learning procedure; an automated way to find this value can even further streamline skill learning and acquisition.

In the short-term, 1) a more comprehensive cross-validation study is needed, involving a larger number of - ideally real - platforms and different locomotion modalities e.g. terrains and velocities; there is also practical appeal in 2) employing RL to autonomously update motor skills and MPs in real time, by exploring the available action space, while potentially informed by perceptual input; 3) selecting and/or formulating motor skills as a function of a general contextual descriptor $\mathcal{T}$ can provide greater flexibility in behavior generalization *i.e.* empirically estimating a generalization function $\mathcal{G} : \mathcal{T} \rightarrow [\mathcal{S}, \tau, \mathbf{y}_0, \mathbf{g}]$ with e.g. contextual RL or GPR. 4) extending motor skills to $\mathcal{S} = \{\mathbf{\Lambda}, \mathbf{\Delta}, \mathbf{T}, \mathbf{m}\}$, where $\mathbf{T}$ defines e.g. individual time scales for each MP in the combination, analogous to τ ; lastly, at a higher level, the design of 5) a hierarchical decision tree to objectively identify in real-time, and when faced with suboptimal performance or perceptual stimuli, when to employ different learning modalities *i.e.* expansive, intrinsic, extrinsic or sequential (2.4.3), would greatly complement the proposed approach.

⁵¹In practice, this can be accomplished by adding duplicate entries within $\mathbf{P}$ which, from an engineering perspective can be simplified with *symbolic* representations.

Dynamical Perceptual Memory Model

In this chapter, the role of perceptual information in robotic control beyond low-level dynamical feedback mechanisms and complex vision-based solutions is discussed, and its relevance in contextual identification is presented as an attractive feature on the path to autonomous Lifelong Learning (L2) robotic agents. A novel data-driven perceptual memory model based on dynamical systems is developed by repurposing some of the key concepts discussed in Chapter 2, framed within the general paradigm of high-level cognitive modeling. The proposed model is validated in an experimental dataset constructed from walking trials with healthy human subjects, through the identification of gait context (velocity and terrain slope) and the detection of slip-like disturbances; additionally, the presence of Perceptual Primitives (PPs) in recorded data is explored as a dimensionality reduction strategy, and incorporated into the model.

The accomplishment of objective **02** is completed, and objectives **01**, **04** and **05** (1.1) are achieved. An accuracy of $\geq 80\%$ in context identification, as well as $\geq 20\%$ dimensionality reduction with PPs, are established as Key Performance Indicators (KPIs) of the proposed methods in a gait monitoring scenario.

3.1 Introduction

As presented in previous chapters, robotic control in the real world presents a challenging task, no less complex when considering the great number of variables at play in the interaction between a robot and its surroundings; the complexity inherent to interacting with the real world presents a sizeable hurdle to overcome on the path towards true robotic adaptation and, in practice, the robustness of any autonomous system is greatly influenced not only by its ability to handle unforeseen perturbations and disturbances to its function - regardless of cause or origin, particularly in dynamic environments - but also by how effectively it can leverage high-level cognitive and reasoning processes to enhance its performance over time [Pastor et al., 2012; Rückert et al., 2013; Tenorth & Beetz, 2009]

Contrary to humans, which are particularly skilled at learning from observation, fine-tuning learned behaviors with further interactions, anticipating how each action impacts their surroundings and applying corrections in real-time, for robots to actively employ perceptual information in such ways parameterization can prove to be complex and time-consuming [Pastor et al., 2011a; Rothfuss et al., 2018; Rückert et al.,

2013]. In nature, perception plays a crucial role in Imitation Learning (IL) [Schaal, 1999], not only as it provides the bases through which organisms observe and model their environment, but also by guiding exploration of new behaviors through priming sensorial inputs. Nevertheless, one crucial challenge in cognitive robotics is formulating general and straightforward mechanisms for encoding, storing and repurposing spatiotemporal features from perception [Rothfuss et al., 2018]. Despite the longstanding academic interest in bioinspired robotic control strategies [Ijspeert, 2014], high-level sensory architectures are rarely employed by robotic platforms and automated systems for robust adaptability and generality; instead, perceptual information is often restricted to localized low-level feedback mechanisms with varying relevance in the global planning stage, such as phase regulation mechanisms in rhythmic controllers [Matos & Santos, 2014; Žabkar & Leonardis, 2016], motion capture and vision systems [Ijspeert et al., 2003; Kragic & Christensen, 2003; Rozo et al., 2016] or force sensors [Hsiao et al., 2010; Li et al., 2018], among others, that use sensor information in a transient fashion, and thus overlook valuable insights about the *stereotypical* nature of perceptual patterns. Dated methods relying on particle filters [Plagemann et al., 2006] and Neural Networks (NNs) [Christensen et al., 2008; Dev Anand et al., 2012] are focused first and foremost in failure detection; to add to their inherent complexity, these lack generality and are heavily oriented towards industrial robots and environments. In the case of robotic manipulation, integrating high-dimensional sensor data (e.g. vision) into control is often complex and challenging [Levine et al., 2018]; modern solutions employing Deep Neural Networks (DNNs) to achieve *end-to-end* training by building complex mappings from raw sensor data to motor commands [Bahl et al., 2020; Levine et al., 2016, 2018], not only disregard or camouflage higher-level cognitive structures, but are still computationally greedy or heavily dependent on simulation, which limits their applicability to real robotic platforms [Chatzilygeroudis et al., 2020; Guan et al., 2021; Sünderhauf et al., 2018; Vernon, 2022], particularly when taking into account demanding performance requirements of real-time control loops [Lenz et al., 2015].

At a higher level, while general semantic knowledge frameworks have been a longstanding goal pursued by AI researchers, illustrated in projects such as Cyc [Lenat, 1995; Matuszek et al., 2006], applying such tools to robotic problems and platforms was, for a long time, not a trivial task, due to a notable lack of symbolic grounding *i.e.* purely ontology¹-based frameworks were generally too high-level and conceptually abstract to be directly and/or efficiently applied in the physical, task-specific contexts in which robots operate [Tenorth & Beetz, 2009]. Despite advances on high-level cognitive architectures over the last decades [Kotseruba & Tsotsos, 2018; Vernon, 2022], particularly in the form of data-driven *emergent* frameworks [Beetz et al., 2010; Byrne, 2015; Crawford et al., 2016; Rothfuss et al., 2018; Sun, 2007; Verschure & Althaus, 2003; Yang et al., 2022], modern solutions are still oriented towards high-level semantics, tailored to specific tasks and environments, or otherwise primed for specific sensor modalities (e.g. vision) [Qiao et al., 2022]; moreover, episodic recall of past experiences, lifelong incremental learning of existing knowledge databases, and simplification strategies such as Perceptual Primitives (PPs), are generally overlooked [Kotseruba & Tsotsos, 2018; Vernon, 2022]. Conversely, when pursuing reactive and versatile robot movement, robust learning and generalization can benefit from grounded associations of motor skills with perceptual patterns, mirroring human cognition [Guan et al., 2021; Pastor et al., 2012].

¹Structured knowledge database, emphasizing relationships between concepts [Guan et al., 2021; Il Hong Suh et al., 2007].

To that end, as suggested by Rothfuss et al. [2018] and Guan et al. [2021], a compact, straightforward and generalizable computational model based on active (data-driven) perception, able to archive and accurately recall experienced patterns, is needed in the context of adaptive robot motor control. Notwithstanding, from a bioinspired perspective, while neurological research and cognitive science have shown that hierarchical models are better suited to represent human memory models, emergent cognitive architectures and robotists mostly focus on dynamical systems [Vernon, 2022; Žabkar & Leonardis, 2016], which prompts the following Research Question (RQ):

RQ2: How can dynamical systems encode stereotypical episodic information associated with different perceptual contexts, while accounting for primitive elements?

Accordingly, herein an alternative memory model is proposed that 1) exploits dynamical systems as an underlying time-invariant representation of stereotypical and amodal (*i.e.* sensor-agnostic) data recordings; 2) explores the existence of *primitive* structures in perceptual patterns *i.e.* Perceptual Primitives (PPs); and 3) proposes a set of practical algorithms for real-time classification of transient sensor input. Focusing on stable and perturbed human locomotion, and using motion capture as perceptual input, episodic perceptual memories of different walking modalities (velocities and terrain slopes) are constructed with the learning strategy described in Chapter 2 (2.3.3), and practical algorithms are designed to exploit the perceptual database for context identification and perturbation detection; performance is evaluated with measured detection *sensitivity*, *specificity* and *accuracy*, as well as *detection time* of non-stereotypical or perturbed gait. Objectives **01**, **04** and **05** are effectively achieved, and by repurposing the control structure introduced in Chapter 2, **02** is completed. From a scientific perspective, this work contributes with a 1) *generic episodic perceptual model* and practical 2) *context identification and tracking algorithms* validated in human walking.

This chapter is organized as follows: Section 3.2 provides an overview of the relevant concepts discussed and built upon in this work, focusing on biological insights in human cognition and practical, bioinspired engineering applications. The perceptual memory model is formalized in Section 3.3, informed by probability theory and derived from the popular DMP framework; the parallels with the control strategy introduced in the previous chapter are discussed, both regarding reduced dimensionality, shared information through PPs and memory learning procedure. Section 3.4 describes the methodology used to validate the proposed perceptual model, in the particular case of context identification and slip detection in human healthy gait, including the construction of the experimental dataset and preliminary signal processing steps. Experimental results are presented in Section 3.5, followed by an in-depth discussion about their significance in Section 3.6. Lastly, Section 3.7 summarily reviews the theoretical and practical approaches followed in this work; general conclusions, in conjunction with experimental guidelines for practical implementation are extrapolated, identifying appealing research to pursue in the short term.

3.2 Background

3.2.1 Paradigms of Human Memory

Memory plays a vital role in many cognitive processes e.g. learning, decision making and problem solving, allowing humans to archive, catalog and recover information [Squire, 1994; Sridhar et al., 2023]. Neuroscience generally classifies memory as either *long-term*, relating to structured knowledge stored for extended periods of time; and *short-term*², describing immediate or transient information gathered during an ongoing cognitive process [Kotseruba & Tsotsos, 2018; Rothfuss et al., 2018]; beyond this categorization, from a multi-level perspective of human cognition, long-term memory is also viewed as either *declarative* or *non-declarative*: the first is inherently associative and requires conscious recall, being related to higher-level cognitive function (e.g. hippocampus or prefrontal cortexes) and generally divided into *semantic* (abstract and generalized knowledge, not necessarily grounded in perception) and *episodic* (past experiences associated with specific spatiotemporal contexts) [Hasselmo, 2011; Kumar, 2021; Ranganath, 2024; Tulving, 1984]³; in contrast, non-declarative or *implicit* memory deals with lower-level instinctual behavior, including procedural memory⁴, perceptual priming, motor reflexes - such as observed in locomotion (2.2.1) - and conditioned responses, being usually associated with activity in motor centers (e.g. basal ganglia and cerebellum) and relying heavily on proprioceptive elements [Knowlton & Schorn, 2024; Squire, 1994]. These elements constitute low-level schemas encoding how specific actions and/or events *feel* to an individual, providing sensory feedback and spatial awareness. Humans, as well as most vertebrates, continuously construct and update proprioceptive memories from each interaction with the environment, which play a crucial role in movement and behavior generalization at all levels; these structures retain not only information about sensor patterns (e.g. receptors in skin, muscle, joints), but also their motor counterparts (e.g. limb and joint positions), informing motor control by anticipating how each action impacts their surroundings and applying corrections in real time [Pastor et al., 2011a]. Incidentally, scientific research has shown that specific areas of the primate brain are simultaneously active during perception *and* action, suggesting a deep connection between both [Schaal, 1999; Sridhar et al., 2023]. This is also evidenced in heightened activity in both the motor and, most notably, sensory cortexes while reading movement-related words [Kumar, 2021]. Moreover, this phenomenon has been observed whether the underlying motion is simple (e.g. reaching and grabbing) or highly complex (e.g. walking, riding a bicycle, juggling objects) [Žabkar & Leonardis, 2016].

In practice, proprioceptive information facilitates *grounding* of high-level declarative memories, particularly movement-related episodic experiences e.g. dancing, sports, or playing an instrument, integrating contextual sensorimotor patterns into higher-level cognitive structures [Hasselmo, 2011; Tenenbaum & Beetz, 2009]; such associations lead to richer and multidimensional episodic memories, through which individuals are able to 1) predict outcome of current experiences and actions; 2) extrapolate general knowledge

²Also referred to as *working* memory [Rothfuss et al., 2018].

³Colloquially, semantic memory encodes *what you know* and episodic memory *what you remember* [Nuxoll & Laird, 2007].

⁴Action-centric representations of learned motor and cognitive skills e.g. sequence of MPs (Chapter 2, 2.2.4).

from past experiences; and exploit it to 3) modulate and enhance current behaviors [Rothfuss et al., 2018].

Furthermore, memory is deeply tied to learning, and thus an essential component of L2 cognitive architectures [Kotseruba & Tsotsos, 2018; Kudithipudi et al., 2022]. Memory formation is a complex process from a neurological, physiological and psychological lens, involving synaptic plasticity through Hebbian mechanisms [Kumar, 2021; Sridhar et al., 2023]; while there are still unknowns as to what exactly is happening behind the scenes, research has identified two particular stages to memory creation: 1) fast, real-time construction of short-term working memories, often grounded on sensor input, and 2) posterior *consolidation*, carried out during rest periods e.g. sleep, where these memories are incorporated and/or translated into higher-level long-term representations [Sridhar et al., 2023; Stickgold, 2005; Wolpert et al., 2011; Žabkar & Leonardis, 2016].

3.2.2 Skill Memories and Perceptual Primitives

In the particular case of learned motor skills, an overlap between episodic and implicit - procedural and proprioceptive - memory has been hypothesized [Nadel, 1994; Willingham, 1998; Žabkar & Leonardis, 2016], to the point that such skill memories are intrinsically linked to motor-centric representations simultaneously grounded on proprioceptive information [Knowlton & Schorn, 2024; Willingham, 1998]. This interplay between conscious empirical recollection and unconscious muscle memory is built upon *stereotypical* - i.e. recurring in a consistent pattern - sensorimotor patterns in repeated execution of motor skills. Moreover, although repeated execution - i.e. training - tends to reduce variability in movement [Schmidt et al., 2018], these patterns remain inherently stochastic, with human movement exhibiting natural variability — a characteristic that has been linked to more efficient motion generation and skill learning mechanisms [Rückert et al., 2013]. Notably, stereotypical behavior represents one of the hallmark features of MP-based control [Pastor et al., 2013, 2012, 2011b].

The pervasiveness of proprioceptive information across different levels of memory formation and skill learning illustrates not only the intertwined nature of perception and action, but also the significance of perception in high-level cognitive processes. Along these lines, just as motor skills have been shown to rely on foundational Motor Primitives (Chapter 2, 2.2.4), perception has been suggested to be structured around Perceptual Primitives (PPs) i.e. low-level elemental units of sensory information that, like MPs, compose low-dimensional components of intricate, complex perceptual patterns [Aslin & Smith, 1988; Frégnac & Bathellier, 2015; Mausfeld, 2003; Palmer, 2015]. While research in this particular subject is, by and large, focused on visual perception [Chen, 1982, 2005; Mayeli, 2019] in the scope of object and shape recognition⁵ [Forsyth et al., 1997; Hock, 2015; Kaiser et al., 2019], the underlying idea has been extended to auditory and haptic/somatosensory primitives [Klatzky & Reed, 2016; Poeppel & Idsardi, 2011; Stiles & Shimojo, 2015a], as well as linguistics [Endress et al., 2007, 2009]. Conversely, studies on facial recognition [Curby & Teichmann, 2022; Gauthier et al., 2003] and speech comprehension [Chakrabarty & Elhilali, 2019] support an alternative holistic view of human perception, built on the Gestalt premise of *the whole is more than the sum of its parts*, that fundamentally challenges a modular and compositional

⁵Often through the use of topological or geometric primitives e.g. spherical and cubic shapes.

organization [de-Wit & Wagemans, 2015] and questions *how* and *where* - if at all - PPs emerge, suggesting instead both paradigms to be relevant in the big picture of perceptual cognition, possibility on a task-specific basis. Notwithstanding, PPs have been exploited by engineers in computer vision and image processing [Higgins et al., 2018; Huang et al., 2009, 2010; Irani & Ware, 2003; Kaiser et al., 2019; Shams & von der Malsburg, 2002; Zhang et al., 2013; Zlatoff et al., 2008], speech recognition [Deverin et al., 2024; Varnet et al., 2015] and robotics⁶ [Coelho et al., 2001; Grotz et al., 2017; Pugeault et al., 2010]. In addition, the idea of *crossmodal plasticity*, supported by scientific evidence of auditory and tactile information being processed by the visual cortex in vision-impaired individuals, suggests that the brain blends available sensory inputs into a single representation [Spence, 2015; Stiles & Shimojo, 2015a] and highlights the appeal of sensor-agnostic models of perception, able to deal with heterogeneous or multimodal perceptual systems coherently, as explored, for instance, in Ruesch et al. [2008].

Despite the lack of - to date - a comprehensive theory on human perception [Cahen & Tacca, 2013], from the perspective of artificial cognitive architectures mimicking human memory [Kotseruba & Tsotsos, 2018], these insights can be exploited for the development of bioinspired memory models that archive and process information in a compact and, most notably, associative way *i.e.* coupling perception and action at different levels, from abstract semantic relationships to low-level raw measurements and motor commands. This notion of Associative Skill Memories (ASMs) [Pastor et al., 2013] relies on the stochastic and stereotypical nature of episodic repetitive experiences, building on the idea of *affordances* [Gibson, 2014; Guan et al., 2021; Pastor et al., 2013] *i.e.* *similar tasks afford similar movement plans* as a way to effectively ground motor policies with context-specific perceptual information, while looking at perception in an abstract and amodal way. In a practical sense, methods for behavior generalization and contextual learning, such as those referenced in Chapter 2 (2.3.3), account for generic context descriptors which usually rely, to variable extents, in perceptual systems *e.g.* ball tracking in table tennis [Kober et al., 2012; Mülling et al., 2013, 2010] and force sensors in grasping [Kalakrishnan et al., 2011; Li et al., 2018]; in such systems, motor policies are generally anchored on contextual perceptual cues, thus constituting a sort of grounded procedural or skill memory.

3.2.3 Perception in Artificial Cognition

Long-running bio-inspired projects such as *Cyc* [Lenat, 1995] and *ConceptNet* [Liu & Singh, 2004] aim to archive human knowledge through structured semantic relationships in an ontological form, enabling reasoning and high-level cognitive mechanisms *e.g.* logical inference and language parsing. However, while foundational to knowledge representation in modern AI, their aim revolves around integrating large-scale explicit information across diverse domains, without direct grounding in sensory or real-world data, functioning more as static - manually curated - repositories than dynamic and/or adaptive architectures.

By contrast, mirroring human cognition in artificial systems has been pursued by researchers and engineers at varying degrees of abstraction [Kotseruba & Tsotsos, 2018; Vernon, 2022], from high-level *symbolic* or *cognitivist* approaches such as *Soar* [Laird et al., 1987], *ACT-R* [Anderson, 1993] and *ICARUS* [Lan-

⁶With a logical overlap between kinematic PPs and MPs.




Figure 32: General overview of different *symbolic*, *emergent* and *hybrid* cognitive architectures (**Bottom**) and associated sensor modalities (**Top**); adapted from Kotseruba et al. [2016].

gley & Choi, 2006] focusing on knowledge-based reasoning and decision making, to data-driven *emergent* or *embodied* frameworks like *DAC* [Verschure & Althaus, 2003], *SPA* [Crawford et al., 2016] and *HTM* [Byrne, 2015], where cognitive structures arise from direct interaction with the world [Vernon, 2022], and *hybrid* solutions that combine both paradigms to varying extents, such as *CLARION* [Sun, 2007] and *CRAM* [Beetz et al., 2010]. Memory is present in these frameworks at different levels of biological plausibility (3.2.1), albeit, generally, in limited capacity: reflecting the design principle, symbolic architectures focus on high-level semantics, while their emergent and hybrid counterparts build procedural memories as perceptually-coupled motor policies or control rules; in addition, short-term, task-specific data buffers constitute functional working memories in most frameworks [Kotseruba & Tsotsos, 2018]. Some bioinspired models cache task- or context-relevant knowledge within working memory as an *activation* of long-term structures, which increases with repeated recruitment; within a motor control context, a similar mechanism is implemented in MoMPs (Chapter 2) [Mülling et al., 2013, 2010]. On the other hand, with few exceptions [Beetz et al., 2010; Mohan et al., 2014], the role of episodic memories is generally overlooked [Bölöni, 2011; Menager & Choi, 2016], even in emergent architectures, despite its well-documented importance in anticipatory processes and lifelong learning [Tulving, 1984; Vernon, 2022].

Regardless of the underlying methodology, modeling and integrating perception - as the basis of all real-world interaction - remains a crucial aspect of any cognitive architecture (Figure 32), the extent of which generally correlates with the level of grounding embedded into the structure. Symbolic models

tend to reduce perception to predefined features and/or abstract representations, making them effective in structured environments but limiting their adaptability to novel or unforeseen scenarios [Harnad, 1990; Matheson & Barsalou, 2018]; conversely, their emergent and hybrid counterparts place a stronger emphasis on processing raw sensor input, achieving richer, context-specific representations - albeit in a distributed and often opaque structure - better suited to dynamic environments [Eliasmith, 2013]. Of note is the prevalence of vision among sensor modalities (Figure 32), by and large the most explored perceptual system available in artificial systems, often relying on bioinspired structures, conventional image processing and cutting-edge techniques to perform object detection, shape recognition and motion tracking [Kotseruba & Tsotsos, 2018]; some architectures complement vision with tools such as Support Vector Machines (SVMs) to extract high-level abstract features (e.g. age, gender) [Martínez-Gómez et al., 2014], or NNs and DL in spatial navigation [Madl et al., 2016; Martens et al., 1998] and other robotic tasks [Lenz et al., 2015; Levine et al., 2016; Rothfuss et al., 2018; Yang et al., 2022]. Beyond vision, audition is also a common perceptual modality mainly focusing on functional communication through speech recognition [Epstein et al., 2011] and natural language processing (e.g. speech-to-text mechanisms) [Grossberg & Myers, 2000; Romero-Garcés et al., 2015]. Lastly, tactile and proprioceptive inputs - e.g. limb position and joint effort - have a prominent role in robotic control, and are usually implemented with force sensors, joint feedback or inertial measurements [Ivaldi et al., 2014; Kotseruba & Tsotsos, 2018; Mohan et al., 2014].

Perception is therefore not a monolithic concept, encompassing a rather diverse range of sensor modalities. As a result, cognitive architectures grapple with the challenge of integrating and reconciling disparate sensory streams into a coherent and contextually meaningful representation [Zmigrod & Hommel, 2013]. This multimodal integration, often accomplished with feature integration or sensor fusion [Fung et al., 2017; Marsh et al., 2022; Zmigrod & Hommel, 2013], is particularly critical in domains like robotics and human-robot interaction, where interpreting nuanced, multimodal inputs (e.g. gestures, speech, environmental changes, and other sensory cues) is essential for effective collaboration and robust task execution [Akkaladevi et al., 2021; Ivaldi et al., 2014; Kotseruba & Tsotsos, 2018; Ruesch et al., 2008]; in addition, it can feed bottom-up attention mechanisms that guide exploration and reactive behavior, as addressed in Ruesch et al. [2008]. While some cognitive models are able to learn associations between different inputs during perceptual learning and/or memory formation [Martens et al., 1998; Verschure & Althaus, 2003], most ignore crossmodal interaction or shared information - e.g. PPs - across different sensors. Overall, cognitive frameworks are generally either vision-centric or tailored for a reduced number of sensors, treat different inputs in isolation and with *off-the-shelf* solutions, or derive task-specific fusion methodologies that lack generality and scalability [Kotseruba & Tsotsos, 2018; Vernon, 2022]; as a result, treating active perception in a sensor-agnostic (amodal) approach remains mostly unexplored.

For an extensive practical and taxonomical survey on cognitive architectures, including the role of perception in artificial cognition, the reader is referred to Kotseruba & Tsotsos [2018].

3.2.4 Related Work

In the context of this study, data-driven research where episodic memories are created from perceptual input, particularly from multimodal sensor sets or, alternatively, in sensor-agnostic (amodal) fashion, is considered relevant literature, with special attention given to works in robotic control, human-robot interaction and behavioral monitoring.

Commonly, emergent architectures employ dynamical systems to model grounded *subsymbolic* processes, such as motor control and perceptual processing [Vernon, 2022]. In fact, in line with the versatility of parametrized representations for motor policies (2.3.1), DMPs have been used to codify heterogeneous or multimodal quantities, commonly kinematic and dynamic control patterns e.g. compliant MPs [Deniša et al., 2016, 2013; Petrič et al., 2018, 2014] as well as position and orientation simultaneously [Magyar et al., 2019; Pastor et al., 2009, 2011b; Ude et al., 2014] in grasping tasks⁷. Beyond DMPs, Kalakrishnan et al. [2011] and used simpler dynamical systems to mix trajectories in task space and torque profiles; Rozo et al. [2016] used GMMs with similar kinematic and dynamic trajectories. However, in these works the focus is largely on motor control and policy representations, and not in multimodal perceptual input and/or data-efficient representation of dynamic patterns; conversely, describing multimodal sensor patterns with dynamical systems has been explored by the author [André et al., 2016b; André et al., 2014] in simulated robotic locomotion.

In dated examples of symbolic architectures adapted to robotic motor control, Il Hong Suh et al. [2007] and Tenorth & Beetz [2009] developed general knowledge management and symbolic reasoning frameworks in Prolog⁸, focusing on human-robot interaction in household environments with mobile platforms. Both works utilize high-level primitive semantics by decomposing the action space into *subactions* and create associative relations by coupling action planning with extrapolated perceptual cues. However, in both cases an extensive set of internal and environmental sensor devices was required. Moreover, continuous improvement of learned behaviors, as well as expanding the knowledge database when facing unknown environments, is not explored. The latter would eventually evolve into a general cognitive architecture in *CRAM* [Beetz et al., 2010], notably modeling episodic memories, and whose capabilities have been demonstrated in varied object manipulation tasks with a PR2 anthropomorphic robot [Vernon, 2022].

With a particular focus on robotic DMP control, Pastor et al. [2011a, 2013, 2012] formalized the notion of Associative Skill Memories (ASMs), framing motor skill libraries as collections of *perceptuocoupled* motor policies; to that end, the authors suggested archiving both expected sensor readings and measured variance in conjunction with motor policies (*i.e.* parametrization), and describe simple statistical failure prediction algorithms. This particular aspect would later be implemented in André et al. [2014] and André et al. [2016b] in simulated walking with an OP humanoid robot, while extending it by also using DMP dynamics to store multimodal (joint encoders and force sensors) sensor and variance patterns in a compact representation, an approach later adopted by Deniša et al. [2016] for motor behavior generalization. In practice, this work demonstrated a capability to predict unstable gait up to 3.6 s in advance, when faced

⁷Coupled DMP systems codifying end-effector position and orientation (*i.e.* quaternions), coupled with finger configuration.

⁸Declarative programming language used in semantic reasoning and computational linguistics [Clocksin & Mellish, 2003].

with irregular of unknown terrains, and with approximate sensitivity and specificity of 87% and 95%, respectively. On the other hand, the dynamical aspect of the representation is not explored, with archived patterns fully reconstructed during recall, decreasing its overall flexibility.

From an autonomous L2 perspective, Mohan et al. [2014] proposed a hybrid cognitive architecture for the IIT iCub humanoid robot, exploiting episodic and procedural memories to enhance cumulative learning, reasoning, and behavioral generalization. The framework was applied to object manipulation tasks e.g. identifying and selecting toys from a table, employing visual and tactile perception. In a modular and open-ended structure, NNs were used to compress temporal sequences of perceptual data and task rewards into fixed-sized feature vectors through Hebbian learning, enabling fast and efficient episodic recall. An initial exploratory stage allowed the robot to gather contextual information about the task (e.g. visual, tactile and spatiotemporal properties of different toys), from which episodic memories were formed; in novel contexts *i.e.* rearranging and/or introducing new toys on the table, new behaviors were generated from the probabilistic combination of past successful episodes - bypassing the need for active exploration. In particular, the *competing* nature behind action selection is noted, whereby episodic memories with reduced relevance are gradually forgotten, while the remainder are further fine-tuned. Despite its generality, the architecture is notably vision-centric, with tactile feedback in a support role, mainly relevant during grasping control. Moreover, only a single task and environment is addressed; validation in other dynamic problems is needed. Conceptually, it also does not avoid *catastrophic* forgetting *i.e.* past episodes are eventually forgotten if not revisited, a documented key challenge in L2 [Kudithipudi et al., 2022].

Exploring ASMs at a higher degree of abstraction, Žabkar & Leonardis [2016] introduced *CHARMS*, a computational framework for motor memory in robotics, with focus on proprioception (including motor commands, joint kinematic state and end-effector position), and based on compositional hierarchies of *AND-OR* graphs to represent motor primitives and abstract them into higher-level cognition. Both online memory formation and offline memory consolidation are addressed, with the latter restructuring acquired memories for efficiency in data storage and generalization across different actions and tasks. Experiments with a real NAO robot, including bipedal walking, different upper limb motions and object manipulation tasks, validated the learning capabilities of proposed approach. However, the methodology is heavily oriented towards abstract semantics and high-level reasoning, with little attention given to the design of lower-level motor patterns. Additionally, learning from past experiences, stochasticity and episodic memory are not accounted for in the memory framework.

Alternatively, in another practical implementation of ASMs, Rothfuss et al. [2018] utilized Deep Neural Networks (DNNs) - specifically convolutional LSTM cells - in the design of an episodic memory model that encoded vision data (frame sequences) into a low-dimensional latent vector space, which was then used to compare different episodes directly in latent space through cosine similarity, and decoded to predict future frames; the proposed methodology was practically validated in robotic object manipulation tasks in unstructured environments, employing DMPs for motor control. Notably, adopting latent *subsymbolic* representations is conceptually related to the notion of PPs. On the other hand, while this work introduces a generic episodic memory model beyond any specific environment and/or task, it requires large-scale supervised training on labelled videos datasets, in line with vision-based DL studies in robotics [Levine

et al., 2016, 2018; Qiao et al., 2022; Sünderhauf et al., 2018].

Recently, Yang et al. [2022] proposed another episodic memory model based on visual and auditory perception, as a part of a hybrid cognitive architecture developed within the context of human-robot memory assistance with the Softbank Robotics Pepper humanoid. Similar to *CHARMS* [Žabkar & Leonardis, 2016], it adopts a graph-based representation, and exploits semantic pointers in the implementation of working and long-term episodic memory structures. Notably, while the proposed model does not extend to generalized high-level cognitive reasoning or planning, it addresses incremental, lifelong growth of the memory database with RL. Practical experiments demonstrated that the robot was able to improve in $\sim 20\%$ memory recall in human participants. Nonetheless, despite the hybrid *end-to-end* nature of the perceptual model, memory representation is built upon high-level semantics. In addition, it delegates perceptual processing to third-party DL-based solutions, such as *YOLOv3* for object detection from video streams, *Temporal Relation Networks* for activity recognition, and Google’s *Speech-to-text* and *Word2Vec* to parse auditory cues and build semantic relationships, respectively; this reliance on DL models (albeit pre-trained) increases the computational footprint and real-time latency of the model.

Alternatively, this work proposes a lower-level generalized model for episodic perceptual memories of stereotypical experiences, exploiting dynamical systems to achieve stability, temporal invariance, and stochastic representation. The focus is on compact, modular and data-efficient subsymbolic representation while retaining the modularity, versatility and generality of higher-level strategies; most notably, no assumption is made regarding sensor modality and level of abstraction, thus applicable to unimodal and multimodal perception, as well as both raw input and extrapolated perceptual features, and crossmodal relationships are incorporated into the model through shared PPs. In practice, despite not limited to proprioceptive input and associative learning not being explored in this work, the proposed approach follows Pastor’s formulation of skill memories by encoding both mean and variance, and accommodates for two-stage learning such as observed in *CHARMS*, while modeling perception in a cumulative, open-ended structure, under a L2 paradigm. While not addressing artificial cognition in its entirety, it can complement existing architectures by providing a straightforward data-driven model that abstracts higher-level cognition from low-level perception, fitting current knowledge about biological cognitive processes, as well as requirements for modern L2 machines.

3.3 Perceptual Memory Model

From the view of skilled robotic movement, stereotypical behavior can be leveraged to model perceptual input in a generalized amodal approach, ranging from raw measurements to extrapolated or processed high-level features. Under this paradigm, a perceptual *context* (further explored in 3.3.2) denotes a broad set of platform-specific perceptual traits - transduced stimuli - that characterize a particular environmental state at a specific moment, often in task-specific fashion. In an associative memory structure, motor tasks can be associated with multiple perceptual contexts, representing optimal motor behavior under varying conditions or reactionary behaviors to specific sensor triggers (e.g. reflexive motions). Perceptual contexts are thus largely dependent on how an agent or system structures its own perception, while motor tasks

are directly grounded in the external environment and specific objective demands.

Under the general assumption that stereotypical experiences, characterized by repetitive noisy sensor patterns with context-specific variability, comprise *cyclostationary*⁹ stochastic processes [Napolitano, 2020] exhibiting some degree of time and scale invariance, when considering K sensor sources synchronously sampled over a finite index set $\Gamma = \{1, 2, \dots, T\}$ ¹⁰, each T^{th} measurement can be seen as a Gaussian process¹¹ $\{\mathbf{x}_t(i) : t \in \Gamma, i \in \mathbb{N}, \mathbf{x}_t \in \mathbb{R}^K\}$ where $\mathbf{x}_t$ follows a K -variate normal distribution:

$$\mathbf{x}_t \sim \mathcal{N}_K(\boldsymbol{\mu}_t, \boldsymbol{\Sigma}_t) \quad (60)$$

whose parameters $\boldsymbol{\mu}_t \in \mathbb{R}^K$ and $\boldsymbol{\Sigma}_t \in \mathbb{R}^{K \times K}$ can be estimated from past iterations of the experience as, respectively, sensor mean and variance observed at the t^{th} sample:

$$\boldsymbol{\mu}_t = E[\mathbf{x}_t] \simeq \bar{\mathbf{x}}_t \quad (61)$$

$$\boldsymbol{\Sigma}_t = E[(\mathbf{x}_t - E[\mathbf{x}_t])(\mathbf{x}_t - E[\mathbf{x}_t])^\top] \simeq \text{var}(\mathbf{x}_t) \quad (62)$$

with the latter, if all K sensor sources are assumed to be independent (thus uncorrelated), defined as:

$$\boldsymbol{\Sigma}_t \simeq \text{diag}(\mathbf{s}_t^2) = \begin{bmatrix} s_{t,1}^2 & & & \\ & s_{t,2}^2 & & \\ & & \ddots & \\ & & & s_{t,K}^2 \end{bmatrix} \quad (63)$$

where $\mathbf{s}_t^2 \in \mathbb{R}^K$ aggregates individual variances $s_{t,k}^2$ registered in previous readings $x_{t,k}$ of the k^{th} sensor, $\forall k \in \{1, 2, \dots, K\}$; from a practical perspective, the sequence of vectors $\bar{\mathbf{x}}_t$ and $\mathbf{s}_t^2$, $\forall t \in \Gamma$, constitute K -dimensional time series describing *nominal* sensor value distribution over the duration of a stereotypical experience.

As proposed in André et al. [2016b], this notion can be further generalized by employing DMP dynamics (2.3.1) to describe observed sensor patterns $\{\bar{\mathbf{x}}_t, \mathbf{s}_t\}$, $\forall t \in \Gamma$ over a continuous and bounded phase domain Φ i.e. without explicit time dependency; a *perceptual memory* or *footprint* is then defined as:

$$\mathcal{M} = \{\mathbf{x}, \mathbf{s}\} \quad (64)$$

where $\mathbf{x} \in \mathbb{R}^K$ and $\mathbf{s} \in \mathbb{R}^K$ represent coupled dynamical variables describing, respectively, expected measurements and their associated standard deviation, as K -dimensional¹² DMP systems. Under this formulation, given fixed initial conditions $\{\mathbf{x}_0, \mathbf{s}_0\}$ and metaparameterization i.e. damping coefficients $\{\alpha_k, \beta_k\}$ and goal states $\{\mathbf{g}^{\mathbf{x}}, \mathbf{g}^{\mathbf{s}}\}$, expected sensor values at any given phase $\phi \in \Phi$ can be estimated or *recalled* by integrating system dynamics at a fixed step Δt and temporal scale τ ¹³, from an initial phase

⁹Exhibiting periodic autocorrelation or *seasonality* i.e. $E[x(t)] = E[x(t + T_0)]$ for all t .

¹⁰Where $\Gamma \in \mathbb{N}$.

¹¹Sequence of *i.i.d.* random variables.

¹² K transformation systems sharing a single canonical system (2.3.1).

¹³Strictly positives values i.e. $\Delta t, \tau \in \mathbb{R}^+$; larger τ and smaller Δt increase the complexity of the procedure.

value $\phi_{\min}$ until ϕ , such that $\mathbf{x}, \mathbf{s} : \Phi \rightarrow \mathbb{R}^K$. By extension, under the same conditions, for any given $\phi \in \Phi$, $\mathcal{M} : \Phi \rightarrow (\mathbb{R}^K, \mathbb{R}^K)$ provides an estimate of distribution parameters in (60), as detailed in Algorithm 2¹⁴. While this method requires $\{\mathbf{x}_0, \mathbf{s}_0\}$ and $\{\mathbf{g}^{\mathbf{x}}, \mathbf{g}^{\mathbf{s}}\}$ to be archived in $\mathcal{M}$, heuristics can simplify parameterization [Lioutikov et al., 2017] e.g. enforcing $\{\mathbf{x}_0, \mathbf{s}_0\} = \{\mathbf{g}^{\mathbf{x}}, \mathbf{g}^{\mathbf{s}}\}$.

Given (64), a top-level memory library $\mathbf{M}$ aggregates M time-invariant episodic perceptual memories:

$$\mathbf{M} = [\mathcal{M}^1, \mathcal{M}^2, \dots, \mathcal{M}^M]^\top \quad (65)$$

with $\mathbf{M} : \Phi \rightarrow (\mathbb{R}^{M \times K}, \mathbb{R}^{M \times K})$ ¹⁵ providing a bulk synchronized estimate of sensor measurement distribution *within the collection of past experiences known to an agent or system*. Under the L2 paradigm [Kudithipudi et al., 2022], mirroring declarative memory in human cognition [Sridhar et al., 2023], and in line with motor libraries $\mathbf{P}$ and $\mathbf{S}$ (Chapter 2), $\mathbf{M}$ can be continuously updated and incremented through new interactions with the environment; ultimately, it can be used to match, in real time, transient measurements with specific memories, which not only supplies contextual information to motor control, but also can be used to identify deviations from nominal behavior and predict how sensor input is likely to evolve.

The proposed model harnesses several appealing aspects of the DMP framework, such as time invariance and independence - relevant when sensor patterns are topologically similar at different time scales - as well as the generic nature with discrete¹⁶ and rhythmic representations - fitting a wide range of input patterns; and its inherent dynamical stability, which circumvents noisy measurements, leading to a more robust and symbolic representation of perceptual data. Additionally, as discussed, employing a primitive-based representation allows exploration into the existence of Perceptual Primitives (PPs) shared across $\mathbf{M}$, which can decrease dimensionality and thus increase data efficiency, simplify acquisition of new memories by focusing on adjusting existing primitives over learning from scratch, and, ultimately, facilitate memory transfer between different platforms. Repurposing the strategy described in Chapter 2 (2.3.2), as an alternative to (64), memories can be defined instead as combinations of simpler PPs:

$$\mathcal{M} = \{\mathcal{S}^{\mathbf{x}}, \mathcal{S}^{\mathbf{s}}\} \quad (66)$$

where $\mathcal{S}^{\mathbf{x}}$ and $\mathcal{S}^{\mathbf{s}}$ follow the motor skill definition in (16) and (20) to encode PP mixtures describing $\mathbf{x}$ and $\mathbf{s}$, respectively. Adopting this revised approach, the previous formulation disregarding any PP sharing can still be obtained with conventional multidimensional DMP systems *i.e.* $P = K$ and $\mathcal{S}^{\mathbf{x}} = \mathcal{S}^{\mathbf{s}} = \{\mathbf{1}_K, \mathbf{0}_K\}$ (Chapter 2, 2.3.3). In practice, PPs can be integrated into $\mathbf{P}$ (9) in a unified database, or, separating perceptual and motor cognition, stored in a separate library $\mathbf{L}$, as illustrated in Figure 33. PP sharing between $\mathcal{S}^{\mathbf{x}}$ and $\mathcal{S}^{\mathbf{s}}$ is not explored in this work, with $\mathbf{L}$ aggregating all PPs instead; notwithstanding, the masking mechanism¹⁷ in (18 - 19) can be used to select relevant PPs as $\mathbf{L}$ grows.

Through (66), similar to motor skill definition, $\mathcal{M}$ represents a latent higher-level representation of a

¹⁴Henceforth, $\mathbf{x}, \mathbf{s}$ and $\mathcal{M}$ are treated both as dynamical variables and, when applicable, functions of ϕ through Algorithm 2.

¹⁵Assuming synchronicity across all $\mathcal{M}^m$ in $\mathbf{M}$.

¹⁶Stereotypical sensor patterns of discrete experiences still imply periodicity; rhythmic DMP formulation can be used.

¹⁷As a sparse binary matrix, $\mathbf{m}$ (19) can be efficiently compressed and archived.




Figure 33: Conceptual flowchart illustrating the dynamical perceptual model proposed in this chapter; a library **M** containing perceptual memories $\mathcal{M}^m$ provides sensor-specific estimates $\hat{x}_k$ and $\hat{s}_k$ which are compared, in real time, with transient context-specific measurements x_k^* through d^m (69); phase ϕ_t and time scale τ is assumed to be inferred with sensor-based heuristic methods.

stereotypical experience, more symbolic in nature, but inherently grounded in perceptual measurements. The generality of this approach is highlighted: few assumptions are made on the underlying perception system (sensor hardware or architecture) *i.e.* both raw measurements and high-level features can be encoded, as long as acquired as stereotypical sequential patterns; multidimensional sensor inputs (e.g. vision) can be decomposed and encoded as unidimensional segments; and by not strictly enforcing PPs in its formulation, it fits both modular and holistic views of perception (3.2.2).

While, analogously to their motor counterparts, PPs can be hypothesized to represent abstract generic patterns - with $\mathcal{M}$ capturing all hardware specificities within sensor data - in this work **L** is considered to be platform-specific; this assumption stems from the observation that sensor readings are dependent on robot hardware and morphology, as was illustrated in the kinematic and dynamic profiles presented in Chapter 2 (2.5.2) and is in line the general insight that perception is highly subjective. Instead, sharing PPs across different perceptual inputs is viewed as a bioinspired tool to decrease dimensionality and facilitate memory acquisition within a single system, when compared to conventional perceptual data storage.

Forming new perceptual memories entails employing a learning procedure similar to the one described in Chapter 2 (2.3.3). While it may not be straightforward to continuously apply the described BBO methodology in real time, in practice, sensor data can be buffered in working memory for posterior offline optimization, mirroring the *consolidation* stage observed in biological organisms [Sridhar et al., 2023; Žabkar & Leonardis, 2016]. Additionally, by establishing connections between individual memories in **M** and motor policies e.g. motor skills in **S**, Associative Skill Memories (ASMs) are created which can be used to achieve robust perceptuomotor coupling, or enhance planning in higher-level cognitive strategies [Pastor et al., 2012]. Extending these associations to the underlying perceptual and motor primitive elements e.g. from **P** to **L**, opens up the door to even more powerful sensor-driven control strategies.

Algorithm 2 Pseudo-code describing the procedure to numerically estimate *i.e. recall* expected value $\mathbf{x}$ and standard deviation $\mathbf{s}$ at ϕ_t (within tolerance e_ϕ) in perceptual memory $\mathcal{M}$; $[\lambda_k^*]^\mathbf{x}$ and $[\delta_k^*]^\mathbf{x}$ denote intrinsic parameters in $\mathcal{S}^\mathbf{x}$ (15) and, similarly, $[\lambda_k^*]^\mathbf{s}$ and $[\delta_k^*]^\mathbf{s}$ intrinsic parameters in $\mathcal{S}^\mathbf{s}$ (15), per (66). Disregarding masking operations in #9 and #12, complexity follows $\mathcal{O}(KN_r)$, $N_r = \lceil \tau(\phi - \phi_{\min}) / \Delta t \rceil$.

```

1: procedure RECALL ( $\mathcal{M}, \phi_t, \tau, \Delta t, \phi_{\min}, e_\phi, \mathbf{x}_0, \mathbf{s}_0, \mathbf{g}^\mathbf{x}, \mathbf{g}^\mathbf{s}$ )
2:    $\{\mathcal{S}^\mathbf{x}, \mathcal{S}^\mathbf{s}\} \leftarrow \mathcal{M}$ 
3:    $\phi \leftarrow \phi_{\min}$ 
4:    $\dot{\phi} \leftarrow 1/\tau$  ▷ (6)
5:    $x_k \leftarrow x_{0k} \forall k \in \{1, 2, \dots, K\}$ 
6:   while  $|\phi_t - \phi| > e_\phi$  do
7:      $\phi \leftarrow \text{UPDATE}(\dot{\phi}, \tau)$  ▷ e.g. Euler method
8:     for each  $k \in \{1, 2, \dots, K\}$  do
9:        $f^\mathbf{x}(\phi) \leftarrow [\lambda_k^*]^\mathbf{x}^\top \mathbf{mP}(\phi_t + [\delta_k^*]^\mathbf{x})$  ▷ (18)
10:       $\ddot{x}_k \leftarrow \tau^{-2} (\alpha_k(\beta_k(g_k - x_k) - \dot{x}_k) + f^\mathbf{x}(\phi))$  ▷ (11)
11:       $\{x_k, \dot{x}_k\} \leftarrow \text{UPDATE}(x_k, \dot{x}_k, \ddot{x}_k, \tau)$ 
12:       $f^\mathbf{s}(\phi) \leftarrow [\lambda_k^*]^\mathbf{s}^\top \mathbf{mP}(\phi_t + [\delta_k^*]^\mathbf{s})$ 
13:       $\ddot{s}_k \leftarrow \tau^{-2} (\alpha_k(\beta_k(g_k - s_k) - \dot{s}_k) + f^\mathbf{s}(\phi))$ 
14:       $\{s_k, \dot{s}_k\} \leftarrow \text{UPDATE}(s_k, \dot{s}_k, \ddot{s}_k, \tau)$ 
15:   return  $\mathbf{x} = [x_1, x_2, \dots, x_K]^\top$  and  $\mathbf{s} = [s_1, s_2, \dots, s_K]^\top$ 

```

3.3.1 Memory Matching

Borrowing from probability theory, given the m^{th} perceptual memory $\mathcal{M}^m$ in $\mathbf{M}$ (65) and an estimate of phase ϕ_t such that $\mathcal{M}_t^m := \mathcal{M}^m(\phi_t)$ denotes the dynamic state of $\mathcal{M}^m$ at ϕ_t , the probability $P(\mathbf{x}_t^* | \mathcal{M}_t^m)$ of transient sensor vector $\mathbf{x}_t^* \in \mathbb{R}^K$ being described¹⁸ by $\mathcal{M}_t^m$, under the normality and independence assumptions, is defined as the product of K sensor-specific probability measures $P(x_{t,k}^* | \mathcal{M}_t^m)$, derived from univariate normal distributions:

$$P(\mathbf{x}_t^* | \mathcal{M}_t^m) = \prod_{k=1}^K P(x_{t,k}^* | \mathcal{M}_t^m), \quad x_{t,k}^* \sim \mathcal{N}_1(\hat{x}_{t,k}^m, \hat{s}_{t,k}^{m\,2}) \quad (67)$$

where $\hat{x}_{t,k}^m := [\mathbf{x}^m(\phi_t)]_k$ and $\hat{s}_{t,k}^m := [\mathbf{s}^m(\phi_t)]_k$ represent the state of $\mathbf{x}$ and $\mathbf{s}$ at ϕ_t , respectively, relative to the k^{th} sensor, as stored in $\mathcal{M}_t^m$; notwithstanding, as the probability associated with a single observation $x_{t,k}^*$ is, by definition, infinitesimally small [Smith, 1999], for practical purposes, an alternative sensor-specific *similarity* metric $d_k^m : \mathbb{R} \rightarrow]0, 1]$ is defined through the normalized density function of a normal distribution *i.e.* a Gaussian Radial Basis Function (RBF)¹⁹ centered at $\hat{x}_{t,k}^m$ and scaled with $\hat{s}_{t,k}^m$:

$$d_k^m(\mathbf{x}_t^*) := d(x_{t,k}^*, \mathcal{M}_t^m) = \exp\left(-\frac{(x_{t,k}^* - \hat{x}_{t,k}^m)^2}{2z_k \hat{s}_{t,k}^{m\,2}}\right) \quad (68)$$

with a global maximum at $d(\hat{x}_{t,k}, \mathcal{M}_t^m) = 1$, and where z_k linearly scales the kernel width, as a way to

¹⁸*i.e.* representing an observation or realization from the K -variate normal distribution (60) parametrized by $\mathcal{M}_t^m$ (64).

¹⁹For reference, a Gaussian RBF is a non-linear function of Euclidean distance.

parameterize a confidence interval around nominal values, contained in $\mathbf{z} \in \mathbb{R}^K$ for all K sensors²⁰. In practice, beyond the data-driven variance encoded in s , z_k can model *priming* behavior in specific sensor inputs [Squire, 1994; Sridhar et al., 2023] or, alternatively, reflect the quality of the underlying memory e.g. decreasing as specific contexts are repeatedly experienced, in line with how stochasticity in skill execution in humans decreases with training [Willingham, 1998]. Consequently, (67) can be substituted by a global similarity measure $d^m : \mathbb{R}^K \rightarrow]0, 1]$, scaling the product to a practical and intuitive range²¹:

$$d^m(\mathbf{x}_t^*) = \left(\prod_{k=1}^K d_k^m(x_{t,k}^*) \right)^{\frac{1}{K}} \quad (69)$$

Objectively, $d^m(\mathbf{x}_t^*)$ is a form of kernel distance proportional to the likelihood of sensor input $\mathbf{x}_t^*$ being described by the m^{th} memory in $\mathbf{M}$; it can be parametrized for a more lax or rigid approach to memory selection i.e. whether sensitivity is preferred over specificity, by means of parameter z_k , akin to the *standard score* is classical statistical theory. Moreover, it allows objective quantitative comparison between different sensor readings and archived memories, which has been shown to be a valuable feature of episodic memory models and high-level cognitive strategies [Rothfuss et al., 2018]. It extends previous works [André et al., 2016b], wherein memory identification is simplified to a binary classification problem solved with a classical two-tailed z -test and, consequently, it is able to deal with multiple overlaps between similar memories and provide a normalized measure for comparison. A similar methodology had been followed in the works of Mülling et al. [2013, 2010] (2.2.5) within the perspective of a sensor-driven gating network combining individual MPs; herein, it is decoupled from motor control, forming the basis of a general perceptual similarity metric that can be harnessed in different scenarios e.g. contextual RL.

Globally, $\mathbf{d}(\mathbf{x}_t^*) \in \mathbb{R}^M$ represents the aggregate measure of d^m for all the M memories in $\mathbf{M}$; from (69), its computation follows $\mathcal{O}(MK)$. In practice, given Δt - implementation-specific - and τ and $\phi_{\min}$ - estimated on a context-specific basis e.g. using heuristics - matching perceptual input $\mathbf{x}_t^*$ with $\mathbf{M}$ has complexity $\mathcal{O}(MK^2N_r)$, with $N_r = \lceil \tau(\phi - \phi_{\min})/\Delta t \rceil$ (Algorithm 2).

3.3.2 Context Identification

Building upon the perceptual model described so far, given perceptual library $\mathbf{M}$, from a dynamic context detection perspective, a perceptual *context* is formalized as the heterogeneous set:

$$\mathcal{C} = \{m, \phi, \tau\} \quad (70)$$

with $m \in \{1, 2, \dots, M\}$, $\phi \in \Phi$ and $\tau \in \mathbb{R}^+$, such that, informally, $\mathbf{M}(\mathcal{C}) := \mathcal{M}^m(\phi, \tau)$, identifying not only which memory $\mathcal{M}^m$ in $\mathbf{M}$ best describes sensor input, but also at which phase ϕ and rate τ it is expected to evolve. Correspondingly, any given $\mathcal{M}^m$ describes dynamic sequences of in the form $\{\mathcal{C}_i : \phi_{i+1} > \phi_i, m_i = m\} \forall i \in \mathbb{N}^+$. Under stable conditions, and assuming the agent is aware of

²⁰For a fixed width z across all K sensors, $\mathbf{z} = z\mathbf{1}_K$.

²¹Without the exponent $1/K$, small discrepancies between the measured and expected values would lead to very low $d^m(x_t)$.

the current perceptual context $\mathcal{C}_t$ associated with transient measurement $\mathbf{x}_t^* \in \mathbb{R}^K$ (Figure 33), a single step update²² of the dynamics in $\mathcal{M}_t^m$ provides a *dead reckoning* estimate $\{\mathbf{x}_{t+1}, \mathbf{s}_{t+1}\}$ of future input²³ $\mathbf{x}_{t+1}^*$, with the expectation that $d^m(\mathbf{x}_{t+1}^*) \geq d^m(\mathbf{x}_t^*)$ (69); conversely, a decrease in d^m implies contextual alterations in the form of 1) acceleration or deceleration in perceptual input patterns i.e. $\tau_{t+1} \neq \tau_t$ leading to inaccuracy in ϕ_{t+1} (Algorithm 2, #7); 2) transition between different experiences i.e. $m_{t+1} \neq m_t$; or 3) deviations from nominal behavior e.g. external disturbances, instability in legged locomotion, etc.

By bootstrapping this procedure on $\mathcal{C}_t$, an accurate initial estimate is strictly required, which may not be always straightforward to obtain under dynamic conditions e.g. real-time interactions; moreover, while τ is not critical to reconstruct perceptual patterns in $\mathcal{M}^m$ by virtue of the time invariance of DMP systems (2.3.1), it influences the dynamic update of ϕ_t , $\mathbf{x}_t$ and $\mathbf{s}_t$ (Algorithm 2), and is thus required for synchronicity with sensor input. Given the cumulative error in dead reckoning, poor estimates increase error in d^m over time, even when sensor input evolves within expected boundaries; this poses an engineering challenge, which can be solved - beyond directly resetting the system with supervisory input - with task-oriented heuristics e.g. temporal markers such as foot strike in bipedal walking.

In essence, when equipped with an overdescriptive $\mathbf{M}$, similarity d^m (69) drives incremental updates to the perceptual awareness of a system: when a prior context $\mathcal{C}_t$ is available, a simple and straightforward strategy to *track* the evolution of $\mathbf{x}_t^*$ - assuming τ does not change - is to update²⁴ all $\mathcal{M}^m$ in $\mathbf{M}$ and identify index m_{t+1} exhibiting highest similarity with a new measurement $\mathbf{x}_{t+1}^*$, according to Algorithm 3; alternatively, in the absence of $\mathcal{C}_t$, all sensor patterns stored in $\mathbf{M}$ are reconstructed in the *search* procedure described in Algorithm 4 (Figure 34), which virtually executes Algorithm 3 over the phase domain Φ , with fixed Δt and search scale τ_s , assigning ϕ_{t+1} and m_{t+1} associated with maximum $d^m(\mathbf{x}_{t+1}^*)$ - over $\mathbf{M}$ and over Φ - to new context $\mathcal{C}_{t+1}$. In addition, repeating this process for sequential measurements $\{\mathbf{x}_{t+1}^*, \mathbf{x}_t^*, \dots\}$ allows *bootstrapping* or *resetting* the perceptual context by estimating τ_{t+1} as $\tau_{t+1} = \Delta t / (\phi_{t+1} - \phi_t)$, in addition to m_{t+1} and ϕ_{t+1} , on the condition that $\phi_{t+1} > \phi_t$; this can be extended, as detailed in Algorithm 5 (Figure 35), to a more robust methodology by using a general ordered set of sequential inputs $\mathbf{X}^* \in \mathbb{R}^{K \times N_b}$ and performing a weighted average according to discounted similarity $\gamma^{n-1} d^m$, with $n \in \{1, 2, \dots, N_b\}$ and $\gamma < 1$, prioritizing newer high-quality matches.

While Algorithm 4 and 5 have increased complexity when compared to Algorithm 3, and are therefore associated with a greater computational overhead, the need for accurate estimation of any prior context $\mathcal{C}_t$ is bypassed, while still preserving time-invariance of the underlying DMP dynamics; moreover, τ_s and Δt ²⁵ provide a way to linearly scale complexity, albeit at the cost of phase resolution in each step update. For practicality, a lower threshold $d_{\min}$ is established in Algorithm 3 (#12), 4 (#11) and 5, such that $\mathcal{C}_{t+1} = \emptyset$ i.e. *null* or *unknown* context, if $d^m(\mathbf{x}_{t+1}^*) < d_{\min}$ over all $\mathbf{M}$, indicating that an agent has no prior recollection of the perceptual input, and which can be used to e.g. prompt exploration. Similarly, a quality threshold d_{match} defines a practical trade-off in order to avoid extra computational steps once a

²²Single loop i.e. operations #7-14 in Algorithm 2.

²³Stochasticity in sensor data can be accounted by sampling from distribution (60) with a pseudo-random number generator.

²⁴Alternatively, an update of ϕ_t and applying Algorithm 2 in all $\mathcal{M}^m$ in $\mathbf{M}$ with ϕ_{t+1} .

²⁵Valid in Algorithm 4 and 5, which perform virtual updates; Δt denotes the interval between $\mathbf{x}_{t+1}^*$ and $\mathbf{x}_t^*$ in Algorithm 3.




Figure 34: Illustrative example of contextual search (Algorithm 4) on a library $\mathbf{M}$ with $M = 6$ perceptual memories; as highlighted ($\square$), $\mathcal{M}^5$ and ϕ_t indicate maximum $d^m(\mathbf{x}_t^*)$ over $\Phi = [\phi_{\min}, \phi_{\max}]$.




Figure 35: Illustrative example of contextual *bootstrap* (Algorithm 5) on a library $\mathbf{M}$ with $M = 6$ perceptual memories; Algorithm 4 is sequentially repeated for perceptual inputs $\mathbf{x}_t$, with ϕ_t ($\square$) representing maximum $d^m(\mathbf{x}_t^*)$ over Φ in each case; given the input sequence, $\mathcal{M}^3$ represents the best match in $\mathbf{M}$.

Algorithm 3 Pseudo-code describing the context tracking method; a new context $\mathcal{C}_{t+1}$ is estimated preserving τ_t from $\mathcal{C}_t$, by combining a dead reckoning update (#6) of all $\mathcal{M}^m$ in $\mathbf{M}$ with a maximum search over $\mathbf{d}$ (#14). Assuming #14 is performed within the main loop in #9, complexity follows $\mathcal{O}(MK)$.

```

1: procedure TRACK ( $\mathbf{M}, \mathbf{x}_{t+1}^*, \mathcal{C}_t, \mathbf{z}, \Delta t, d_{\text{match}}, d_{\text{min}}$ )
2:    $\{m_t, \phi_t, \tau_t\} \leftarrow \mathcal{C}_t$ 
3:    $\phi_{t+1} \leftarrow \text{UPDATE}(\phi_t, \Delta t, \tau_t)$  ▷ e.g. Euler method
4:   for each  $m \in \{1, 2, \dots, M\}$  do
5:     for each  $k \in \{1, 2, \dots, K\}$  do
6:        $\mathcal{M}_{t+1}^m \leftarrow \text{UPDATE}(\mathcal{M}_t^m, \Delta t, \tau_t)$ 
7:        $\{\hat{x}_{t+1,k}^m, \hat{s}_{t+1,k}^m\} \leftarrow \mathcal{M}_{t+1}^m$ 
8:        $d_k^m(x_{t+1,k}^*) \leftarrow \exp\left(-\frac{(x_{t+1,k}^* - \hat{x}_{t+1,k}^m)^2}{2z_k \hat{s}_{t+1,k}^m}\right)$  ▷ (68)
9:        $d^m(\mathbf{x}_{t+1}^*) \leftarrow \left(\prod_{k=1}^K d_k^m(x_{t+1,k}^*)\right)^{\frac{1}{K}}$  ▷ (69)
10:      if  $d^m(\mathbf{x}_{t+1}^*) > d_{\text{match}}$  then ▷ Optional
11:        return  $\mathcal{C}_{t+1} = \{m, \phi_{t+1}, \tau_t\}$ 
12:      if  $\max(\mathbf{d}(\mathbf{x}_{t+1}^*)) < d_{\text{min}}$  then
13:        return  $\mathcal{C}_{t+1} = \emptyset$ 
14:       $m_{t+1} \leftarrow \arg \max_m (\mathbf{d}(\mathbf{x}_{t+1}^*))$  ▷ Single-pass in #9
15:    return  $\mathcal{C}_{t+1} = \{m_{t+1}, \phi_{t+1}, \tau_t\}$ 

```




Figure 36: Conceptual workflow and hierarchical organization between the different levels of perceptual context identification described; given $\mathcal{C}_t$, Algorithm 3, 4 and 5 are sequentially chained according to the success or failure of context identification, with increasing computational overhead.

satisfactory match is found²⁶ i.e. $d^m \geq d_{\text{match}}$ (#10, #9). Furthermore, d_{min} and d_{match} also provide a way to indirectly modulate the sensitivity and specificity of perceptual context identification. In order to minimize computational overhead in real-time perception, Algorithm 3, 4 and 5 can be applied hierarchically, as illustrated in Figure 36.

During this work, the C++ libraries developed in Chapter 2 were extended to integrate the perceptual memory model described in this section, as well as object-oriented implementations of Algorithm 3, Algorithm 4 and Algorithm 5. For a detailed description of the implementation, including the underlying structure and technical specifications, please refer to Appendix A: Practical Implementation.

²⁶Reasonable values found in $d_{\text{min}} = 0.5$ and $d_{\text{match}} = 0.8$.

Algorithm 4 Pseudo-code describing the context search method illustrated in Figure 34; Algorithm 3 is applied sequentially over Φ a total of $N_s = \tau_s(\phi_{\max} - \phi_{\min})/\Delta t$ times. Assuming #13 and #14 are performed within the main loop in #8 and #7, respectively, complexity follows $\mathcal{O}(MKN_s)$.

```

1: procedure SEARCH ( $\mathbf{M}, \mathbf{x}_{t+1}^*, \mathcal{C}_t, \tau_s, \mathbf{z}, \Delta t, d_{\text{match}}, d_{\text{min}}$ )
2:    $\{m_1, \phi_1, \tau_t\} \leftarrow \mathcal{C}_t$ 
3:    $N_s = \lceil \tau_s(\phi_{\max} - \phi_{\min})/\Delta t \rceil$  ▷ Total number of integration steps over  $\Phi$ 
4:    $\mathcal{C}_1 \leftarrow \{m_1, \phi_1, \tau_s\}$ 
5:   for each  $\phi_n \in \{\phi_2, \phi_3, \dots, \phi_{N_s}\}$  do
6:      $\mathcal{C}_n \leftarrow \text{TRACK}(\mathbf{M}, \mathbf{x}_{t+1}^*, \{m_{n-1}, \phi_{n-1}, \tau_s\}, \mathbf{z}, \Delta t, d_{\text{match}}, d_{\text{min}})$  ▷ Algorithm 3
7:      $\{m_n, \phi_n, \tau_s\} \leftarrow \mathcal{C}_t$ 
8:      $d^{m_n}(\mathbf{x}_{t+1}^*) \leftarrow \left( \prod_{k=1}^K d_k^{m_n}(x_{t+1,k}^*) \right)^{\frac{1}{K}}$  ▷ (69)
9:     if  $d^{m_n}(\mathbf{x}_{t+1}^*) > d_{\text{match}}$  then ▷ Optional
10:      return  $\mathcal{C}_{t+1} = \{m_n, \phi_n, \tau_t\}$ 
11:   if  $\max(\mathbf{d}(\mathbf{x}_{t+1}^*)) < d_{\text{min}}$  then
12:     return  $\mathcal{C}_{t+1} = \emptyset$ 
13:    $m_{t+1} \leftarrow \arg \max_{m_n} (\mathbf{d}(\mathbf{x}_{t+1}^*))$  ▷ Single-pass in #8
14:    $\phi_{t+1} \leftarrow \arg \max_{\phi_n} (\mathbf{d}(\mathbf{x}_{t+1}^*))$  ▷ Single-pass in #7
15:   return  $\mathcal{C}_{t+1} = \{m_{t+1}, \phi_{t+1}, \tau_t\}$ 

```

3.4 Validation in Human Locomotion

Within the scope of **RQ2**, a concrete real-world application of the described perceptual model is provided, demonstrating how the proposed perceptual model (3.3) can accurately and uniquely describe episodic stereotypical experiences in a real-world setting to different practical ends. Human bipedal locomotion is treated as an illustrative case study, whose stereotypical and periodic nature can be leveraged to construct perceptual memories - with or without PPs - for different gait modalities and, employing context identification algorithms (3.3.2), detect behavioral abnormalities in timely fashion.

To that end, an experimental dataset, specifically constructed to study slip-related gait perturbations [Ferreira, 2022], provided experimental data on different types of bipedal walking with multiple individuals. Adopting a cross-validation strategy, multiple episodic perceptual memories were constructed, and used to identify gait type from transient input, as well as detect abnormal behavior such as slip-like events. While context identification is achieved offline, and dynamic online operation is not explored in this work, experimental data is sourced on a sample-by-sample basis, thus mimicking real-time conditions; Moreover, in contrast to previous works [André et al., 2016b; André et al., 2014], it showcases practical applicability beyond the field of robotics, relevant from the multidisciplinary perspective of this thesis.

3.4.1 Slip-Related Disturbances

Slip-induced falls are a significant cause of accidents, especially among the elderly, often leading to hip fractures [Chang et al., 2016]. From a social and health perspective, given the potentially severe consequences of slip-related falls, it becomes of paramount importance to develop a complete comprehension

Algorithm 5 Pseudo-code describing the bootstrap methodology; from set of sequential N_b measurements $\mathbf{X}^* \in \mathbb{R}^{K \times N_b}$, ordered in descending temporal order, Algorithm 4 is applied to each column $\mathbf{x}_n^*$, and $\mathcal{C}_{t+1}$ estimated through a discounted weighted average over all $\mathbf{X}^*$. Complexity follows $\mathcal{O}(MKN_sN_b)$.

```

1: procedure BOOTSTRAP ( $\mathbf{M}, \mathbf{X}^*, \gamma, \tau_s, \mathcal{C}_t, \mathbf{z}, \Delta t, d_{\text{match}}, d_{\text{min}}$ )
2:   for each  $\mathbf{x}_n^*$  in  $\mathbf{X}^*$  do
3:      $\mathcal{C}_n \leftarrow \text{SEARCH}(\mathbf{M}, \mathbf{x}_n^*, \mathcal{C}_t, \tau_s, \mathbf{z}, \Delta t, d_{\text{match}}, d_{\text{min}})$  ▷ Algorithm 4
4:      $\{m_n, \phi_n, \tau_n\} \leftarrow \mathcal{C}_n$ 
5:      $a_n \leftarrow \gamma^{n-1} d^m(\mathbf{x}_n^*)$  ▷ (69)
6:     if  $n > 1$  and  $\phi_n > \phi_{n-1}$  then
7:        $\tau_n \leftarrow \Delta t / (\phi_n - \phi_{n-1})$  ▷  $\tau_t \in \mathbb{R}^+$ 
8:        $m_{t+1} \leftarrow \text{round} \left( \frac{\sum_{n=1}^{N_b} a_n m_n}{\sum_{n=1}^{N_b} a_n} \right)$ 
9:        $\phi_{t+1} \leftarrow \frac{\sum_{n=1}^{N_b} a_n \phi_n}{\sum_{n=1}^{N_b} a_n}$ 
10:       $\tau_{t+1} \leftarrow \frac{\sum_{n=2}^{N_b} a_n \tau_n}{\sum_{n=2}^{N_b} a_n}$ 
11:   return  $\mathcal{C}_{t+1} = \{m_{t+1}, \phi_{t+1}, \tau_{t+1}\}$ 

```

of human response to slipping and design effective strategies for fall detection and prevention. To that end, scientific research has focused extensively on understanding slip biomechanics and developing preventive strategies [Aprigliano et al., 2017; Chou et al., 2011; Marigold & Patla, 2002; Moyer et al., 2009; Nazifi et al., 2017; Parijat & Lockhart, 2008], as well as scrutinizing the intricate interplay between sensory input and motor control [Karlsson et al., 2013; Monaco et al., 2017]. Real-time wearable assistive technologies provide a newfound sense of security, effectively diminishing the fear of falling and empowering individuals to carry out their daily activities with confidence. This, in turn, reduces the need for post-fall rehabilitation, leading to cost savings and a more efficient allocation of human resources.

From a biomechanical lens, slips often occur during heel strike as a result of poor foot-floor friction, leading to Loss of Balance (LoB), and can be aggravated by environmental factors [Lockhart, 2013]. LoB triggers different sensor receptors, which prompt the Central Nervous System (CNS) to engage specific muscular units in stabilizing and recovery mechanisms [Lockhart, 2013]; scientific evidence shows that this biomechanical response generally starts 160-200 ms after slip onset, with compensatory behaviors lasting approximately 150-480 ms, contingent upon age and other pertinent factors [Beschorner et al., 2013; Chambers & Cham, 2007; Martelli et al., 2017; Parijat & Lockhart, 2012; Yang et al., 2008].

3.4.2 Automated Slip-Avoidance Strategies

Practical solutions often target specific joints, such as the hip or knee, to both detect and correct LoB scenarios by restoring a biomechanically favorable position. As reported in the literature, detection times range from as low as 30-300 ms following a slip, while actuation times typically fall within the range of 200-350 ms [Mioskowska et al., 2020; Monaco et al., 2017; Trkov et al., 2017]. The Active Pelvis Orthosis (APO) [Monaco et al., 2017] is one such example, built upon an adaptive-threshold algorithm using gait predictive Adaptive Frequency Oscillators (AFOs)²⁷, that enforces a customized torque-based

²⁷Commonly used to implement CPG-based locomotion [Dzeladini et al., 2019].

assistance at the hip joint for 250 ms in order to counteract the downward displacement of the CoM, reporting LoB detection within 350 ms; alternatively, the Robotic Knee Assistive Device (ROKAD) [Trkov et al., 2017] employs impedance and torque feedback control to actuate on the knee joint during slip events, relying on an IMU and force-based slip detection algorithm; Lastly, Mioskowska et al. [2020] proposed an alternative assistance strategy that employed a cable-driven knee brace mechanism to quickly re-establish ground contact in the event of LoB, effectively expanding the subject's BoS for balance recovery and stability; an average actuation time of 82 ms for a 90 to 0° extension and 72 ms for a 60 to 0° extension was measured. All these represent wearable solutions designed with unobtrusiveness in mind, designed to avoid hindering normal gait under stable conditions; in addition - perhaps most importantly - their optimal performance is reliant on accurate detection of LoB. Nevertheless, it is relevant to note that the practical applicability of these devices goes beyond the scope of gait assistance and clinical rehabilitation; in fact, timely and accurate detection of risk conditions - such as, but not limited to, LoB - is crucial to safe and nominal operation of most robotic devices, but particularly challenging outside industrial settings where environments are routinely managed so that robots can operate within well-modeled task parameters [Mülling et al., 2010].

3.4.3 Walking Trials

In the pursuit of an alternative fall-prevention strategy, Ferreira [2022] conducted human locomotion trials in a controlled environment, reproducing different types of walking with multiple participants, while recording an extensive set of physiological and biometric data. This study utilizes a portion of the resulting dataset - in the form of lower-limb kinematics tracked with a motion capture system - for the validation of the perceptual model and context detection methodology described. Notably, despite the multimodal nature of the experimental data, herein the focus is on motion capture as a perceptual source, due to its applicability both in and beyond robotics and automated systems as part of vision processing stacks; other biosensor information available, such as respiratory, cardiac and electromyographic patterns, is excluded.

3.4.3.1 Participants and Equipment

The participant pool included 11 healthy young individuals (age: 24.55 ± 2.15 ; height: 1.70 ± 0.09 m; weight: 63.25 ± 7.11 kg; males = 6; females = 5) that respected the following inclusion criteria: 1) capable of healthy stable locomotion; 2) exhibited total postural balance; 3) ≥ 18 years old; and with 4) body mass ≤ 135 kg. Participants that disclosed any 1) locomotive or gait-related pathology, or 2) were recently subjected to surgical procedures that affect mobility, were excluded. For a qualitative assessment of foot preference, each participant completed the *Waterloo Footedness Questionnaire* [Elias et al., 1998]. All participants²⁸ willingly agreed to participate in the experimental studies after providing written informed consent. The research adhered to the ethical guidelines defined by the University of Minho Ethics Committee (CEICVS 063/2021), which aligns with the principles outlined in the declaration of Helsinki and the Oviedo Convention.

²⁸Henceforth abbreviated as $\mathbf{S}j$, $j \in \{1, 2, \dots, 11\}$.



Figure 37: Experimental apparatus (**Top**) adopted in the human walking trials, with brief snapshots of a perturbation being induced by the operator (**Bottom**). Adapted from [Ferreira, 2022].

Locomotion trials were conducted indoors, on an motorized incline treadmill (TC290, Domyos, France); a multithreaded nylon rope (20 mm diameter), tied to the ankle, was used to manually induce slip-like perturbations by a trained operator, as depicted in Figure 37. During the trials, kinematic data²⁹ was acquired with a full-body inertial wireless motion capture system (MVN Awinda®, Xsens Technologies, B.V., The Netherlands) through the *MVN Analyze*® software suite (v. 2021.2) at 60 Hz; in addition, video of the experiment was recorded by a color and depth camera (Kinect 2, Microsoft, USA) at 30 fps. For safety reasons, all participants were equipped with a safety harness attached to the ceiling (Figure 37).

3.4.3.2 Experimental Protocol

Prior to the walking trials, participants were outfitted with the full body motion capture system³⁰ which was calibrated per manufacturer instructions; the safety harness was then equipped with the rope length adjusted to register a minimum of 15 cm between the knees and the treadmill belt [Okubo et al., 2019], as illustrated in Figure 37. After instrumentation, participants were encouraged to walk on the treadmill at their own leisure, in order to familiarize themselves with the experimental setup.

As presented in Table 6, locomotion trials included every possible combination of perturbed 1) leg (right and left), and 2) gait event (heel strike and toe off), for 3) two treadmill inclinations (0% and 10%) and 4) two gait velocities (1.8 and 5.4 km/h, or 0.5 and 1.5 m/s, respectively, inspired by relevant literature [Klemetti et al., 2014; Kyrdalen et al., 2019; Quach et al., 2011]), and repeated 5) with and without perturbations³¹; a total of 32 trials per participant were registered, executed in a randomized order.

During perturbation trials, at 3 random moments, a trained operator tugged the ankle rope during

²⁹Following International Society of Biomechanics (ISB) guidelines for joint coordinate systems [Wu et al., 2002].

³⁰A total 17 IMUs attached to lower and upper limbs, torso and head.

³¹Retaining the classification of perturbed leg and event (left/right, heel strike/toe off) for consistency.

Table 6: Overview of experimental conditions in walking trials conducted with each participant, showing all combinations of velocity, treadmill incline, perturbed leg - left (L) and right (R) - and gait event - heel strike (HS) and toe off (TO); each trial was repeated with and without perturbations; a total of 32 trials were conducted per participant, and executed in a blind randomized order.

Velocity	Incline	Unperturbed				Perturbed			
		Leg		Event		Leg		Event	
		L	R	HS	TO	L	R	HS	TO
1.8 km/h	0°	•	•	•	•	•	•	•	•
	10°	•	•	•	•	•	•	•	•
5.4 km/h	0°	•	•	•	•	•	•	•	•
	10°	•	•	•	•	•	•	•	•

heel-strike and toe-off events, causing instability comparable to a slip event [Ferreira et al., 2022]. In order to avoid any prior bias in their biomechanical response, participants were unaware 1) if, 2) when and 3) how often perturbations would occur; additionally, participants were instructed to focus their gaze on a spot at roughly eye level while walking, and to try their best to recover normal gait after each perturbation induced by the operator. Perturbation trials were stopped after roughly 10 s of stable gait following the third perturbation, with the average trial duration ranging from 30 s to 1 min. Conversely, trials without perturbations lasted approximately 30 s on average.

3.4.3.3 Data Processing Pipeline

Using video of the experiment, recorded data was manually labeled, with special care taken in identifying, for each induced perturbation, a *start* sample - identifying when the operator started to pull on the nylon rope - and *end* sample - first heel strike of the perturbed foot after the participant has regained steady gait; in contrast, perturbations that did not cause any visually perceived gait alterations were ignored and considered stable gait during analysis. At this stage, **S2** and **S6** were discarded due to data corruption.

Subsequently, using problem-specific heuristics, lower limb kinematics (comprising a total of $K = 18$ DoFs) in each labeled trial were segmented, resampled and averaged according to the processing pipeline illustrated in Figure 38. Specifically, kinematic patterns of the left knee *flexion/extension* angle $\mathbf{q}_{(\text{lknee})}$ were submitted to a 1) *thresholding* step, followed by a 2) *differencing* operation that highlighted threshold transitions and a 3) *peak detection*³² algorithm that accurately located indexes t_i , $i \in \mathbb{N}^+$ corresponding to the edges of the i^{th} gait cycle (i.e. step) in the trial; then, in all K inputs $\mathbf{q}$, each interval $[t_i, t_{i+1}]$ was 4) *linearly interpolated* at $\hat{t}_l = t_i + l(t_{i+1} - t_i)/L$, $l \in \{1, 2, \dots, L\}$, in order to obtain multiple sequences of fixed length L , which were then 5) *averaged* on a sample-by-sample basis, producing reference trajectories $\bar{\mathbf{x}}^{\text{REF}}, \bar{\mathbf{s}}^{\text{REF}} \in \mathbb{R}^{K \times L}$ with sample average and standard deviation across segmented periods in multiple trials. Herein, participant-specific thresholds $q_{\min}$ were hand-tuned³³ and $L = 100$ was fixed for all trials, normalizing the time scale of trials at different velocities. In order to ensure reasonable segmentation results, two hard requirements were enforced: a minimum of 9 segmented cycles per trial, and standard

³²Using `find_peaks` function from `scipy.signal` Python module, with *height* argument as $0.7 \max(\mathbf{q}_{(\text{lknee})})$.

³³As a general rule of thumb, $q_{\min}$ was set to a value that is intersected *only once* per period/gait cycle.




Figure 38: Conceptual flowchart of the processing steps conducted to extract fixed-length reference periods $\bar{\mathbf{x}}^{\text{REF}}, \mathbf{s}^{\text{REF}} \in \mathbb{R}^{K \times L}$ from periodic kinematic recordings $\mathbf{q}_k$, using left knee flexion/extension angle as reference for segmentation; t_t denotes the time stamp of the t^{th} sample, $\tilde{t}_L^i$ the interpolation time stamps in the i^{th} cycle, $q_{\min}$ the knee angle threshold, and L the number of reference samples.

deviation of step duration should be below 10 samples i.e. $\sqrt{\text{var}(t_{i+1} - t_i)} \leq 10$, $i \in \mathbb{N}^+$. Trials where these conditions were not met were discarded before the *averaging* step (Figure 38).

3.4.4 Memory Consolidation and Practical Validation

Under the assumption that, in the absence of perturbations and given the controlled experimental setting, different participants exhibited similar walking patterns, inter-participant references of normal gait were constructed from unperturbed trials of matching velocity and treadmill inclination, in a Leave-One-Out (LOO) cross-validation scheme, as illustrated in Figure 39; in each of the 9 folds, these constitute $\{\bar{\mathbf{x}}^{\text{REF}}, \mathbf{s}^{\text{REF}}\}$ associated with walking at 1) 1.8 km/h, on flat ground, at 2) 1.8 km/h, on a 10° incline, at 3) 5.4 km/h, on flat ground and at 4) 5.4 km/h, on a 10° incline. In addition, for analysis purposes, auxiliary *global* references $\{\bar{\mathbf{x}}^{\text{GLB}}, \mathbf{s}^{\text{GLB}}\}$ were constructed by pooling data from all participants, bypassing cross-validation partitioning of the dataset.

A parallel correlation analysis was conducted on $\bar{\mathbf{x}}_k^{\text{GLB}}$ in order to quantify the level of monotonic correlation - if any - across the $K = 18$ DoFs, by means of the Spearman's rank correlation coefficient r_s ; in theory, this could allow the identification of redundant dimensions, and thus simplify memory construction, as well as inform PP sharing between different DoFs. On one hand, the proposed memory model is grounded on the statistical independence of different sensor sources (3.3); on the other, considering the synergistic nature of muscular behavior, which is well documented regarding lower limb movements such as walking and running [d'Avella & Bizzi, 2005; Zandvoort et al., 2022], and the homogeneous nature of the dataset, some degree of correlation was to be expected. In practice, while $r_s = 0$ does not imply statistical independence, herein it is understood to support applicability of the proposed model. To this end, r_s is calculated for each pair of rows in $\bar{\mathbf{x}}^{\text{GLB}}$, yielding a symmetric correlation matrix $\mathbf{C} \in \mathbb{R}^{K \times K}$; row-wise averages³⁴ of $\mathbf{C}$ provide an objective measure of interdimensional correlation in $\bar{\mathbf{x}}^{\text{GLB}}$. In addition, $\bar{\mathbf{x}}^{\text{GLB}}$ was also decomposed with Non-negative Matrix Factorization (NMF)³⁵ [Lee & Seung, 1999] for $r_{\text{NMF}} = \{1, 2, 3, 4, 5, 6\}$ components, with the goal of assessing the level of shared information across

³⁴Given that $\text{diag}(\mathbf{C}) = \mathbf{1}_K$, diagonal elements (autocorrelation) are excluded.

³⁵Enforcing NMF non-negativity requirement with a prior *min-max* normalization step.

Table 7: Different learning strategies (2.3.3) adopted during memory consolidation, applied separately to $\bar{\mathbf{x}}^{\text{REF}}$ and $\mathbf{s}^{\text{REF}}$, to populate all $\mathbf{M}^{\mathbf{s}^j}$, and corresponding number of PPs, CMA-ES parameters λ and σ , learning dimensionality n (given $N = 20$ and $K = 18$) as well as the total number of repetitions (#) required throughout the 9 cross-validation folds, to fully consolidate $\{\bar{\mathbf{x}}^{\text{REF}}, \mathbf{s}^{\text{REF}}\}$ into $\mathcal{M}$ (66).

Strategy	Parameters			n	#
	P	λ^{BBO}	σ_0^{BBO}		
Isolated	18	15	50	20	$2 \times 18 \times 9$
Bulk	5	30	50	280	$2 \times 1 \times 9$
Baseline + Tuning	5	30 20	50 10	280	2×1 $2 \times 1 \times 9$

the experimental dataset. This approach borrows from neurophysiology, where NMF has been extensively used to identify primitives in muscular coordination [d’Avella & Bizzi, 2005], and represents a less computationally intensive methodology when compared to repeated execution of the skill learning procedure with increasing P , as adopted in Chapter 2 (2.5.1).

Employing the proposed memory model (3.3), participant-specific $\bar{\mathbf{x}}^{\text{REF}}$ and $\mathbf{s}^{\text{REF}}$ were separately³⁶ encoded into DMP multidimensional systems through the skill learning approach described in Chapter 2 (2.3.3), parametrizing $\mathcal{S}^{\mathbf{x}}$ and $\mathcal{S}^{\mathbf{s}}$ with $\theta^{\mathbf{x}}$ and $\theta^{\mathbf{s}}$, respectively, through least-squares optimization (23 - 24). In all memories, the number of basis was fixed at $N = 20$ (4), and the procedure was conducted with and without shared PPs across the K inputs, according to 3 alternative expansive learning strategies of varying complexity, summarized in Table 7:

- **Isolated** ignoring shared PPs *i.e.* $P = K$ (3.3), allowing each input to be approximated independently in a total of K optimization steps of dimensionality $n = N$;
- **Bulk** *one-shot* learning sharing $P < K$ PPs, increasing dimensionality to $n = 2KP + KN$;
- **Two-stage** learning with $P < K$, where a global *baseline* learned from $\{\bar{\mathbf{x}}^{\text{GLB}}, \mathbf{s}^{\text{GLB}}\}$ was *tuned* for each participant-specific $\{\bar{\mathbf{x}}^{\text{REF}}, \mathbf{s}^{\text{REF}}\}$, with $n = 2KP + KN$.

In all cases, under the adopted folding scheme (Figure 39), a perceptual library $\mathbf{M}^{\mathbf{s}^j}$ was obtained for each non-discarded participant $j \in \{1, 3, 4, 5, 7, 8, 9, 10, 11\}$, describing 4 different types of stereotypical - *i.e.* unperturbed - gait (Table 6). Following the implementation outlined in Chapter 2 (2.4.3), learning was conducted with CMA-ES, manually tuned for each strategy (Table 7), prioritizing exploration by increasing σ_0^{BBO} when learning from scratch, and scaling λ^{BBO} according to n . Weights w_n (4) and $\mathbf{\Lambda}$ were constrained to $[-1000, 1000]$, and, similarly, $\mathbf{\Delta}$ was constrained to $[-\pi, \pi]$ (15);

In each cross-validation fold, data associated with the test participant - left out of memory construction - was used to evaluate how accurately the proposed model was able to encode stereotypical patterns specific to each gait type, in line with **RQ2**. To that end, equipped with $\mathbf{M}^{\mathbf{s}^j}$, individual test samples $\mathbf{q}_t$ where sequentially fed to Algorithm 4 with two practical purposes:

³⁶Shared information between $\bar{\mathbf{x}}^{\text{REF}}$ and $\mathbf{s}^{\text{REF}}$ was not explored in this study.




Figure 39: Illustration of the Leave-One-Out (LOO) cross-validation scheme adopted in the creation of participant-specific perceptual libraries; kinematic patterns $\mathbf{q}$ are sourced from trials matching velocity and inclination in all remaining participants; segmented, interpolated and averaged according to the procedure in Figure 54; and the resulting reference trajectories $\mathbf{x}^{\text{REF}}$, $\mathbf{s}^{\text{REF}}$ fed, separately, to the skill learning procedure described in Chapter 2 (2.3.3) (23 - 24). **S2** and **S6** were discarded due to data corruption, resulting in a total of 9 folds for the remaining participants. As highlighted, in each fold, all trials conducted with the skipped participant, with and without perturbation, compose the test subset.




Figure 40: Kinematic references $\bar{\mathbf{x}}^{\text{REF}} \pm \mathbf{z}^{\text{T}} \mathbf{s}^{\text{REF}}$, with $\mathbf{z} = 1.96 \mathbf{1}_K$, obtained for 6 lower limb DoFs of **S3**.

- **Gait Type Identification** in unperturbed samples³⁷, classifying them according to m in $\mathcal{C}_t$;
- **Detection of Perturbed Gait** by locating *non-stereotypical* samples where $\mathcal{C}_t = \emptyset$.

As both cases are formulated as classification problems, resulting performance is evaluated through the *sensitivity*, *specificity* and *accuracy* of detection; in addition, the *detection time* of perturbed gait is measured as the temporal difference between the start of perturbed gait and the closest sample classified with a null context $\emptyset$. In all trials, $d_{\min} = 0.5$, $d_{\text{match}} = 0.8$, $\tau_s = 2.0$ and $\mathbf{z} = z \mathbf{1}_K$, with $z = 1.96$ ³⁸.

For practicality, data processing and result statistical analysis were implemented in Python, complementing custom code developed for interfacing with the experimental dataset and conducting LOO cross-validation; The general-purpose C++ code developed in Chapter 2 and 3.3 provided the implementation of the underlying dynamical systems, primitive and skill structures, as well as skill learning operations, through high-level Python bindings, preserving computational efficiency.

3.5 Results

3.5.1 Kinematic References

Figure 40 plots an illustrative subset of participant-specific validation references $\bar{\mathbf{x}}^{\text{REF}} \pm \mathbf{z}^{\text{T}} \mathbf{s}^{\text{REF}}$, in the case of **S3**. In all depicted sensor dimensions - denoting lower limb DoFs - kinematic trajectories exhibited considerable similarity, relatively low variance - negligible in localized segments, as seen in foot *abduction/adduction* angles - and a consistent half-period *i.e.* $L/2 = 50$ sample delay between left and

³⁷ Denoting samples in trials without perturbation, as well as samples before and after perturbations in trials with perturbation.

³⁸ Under the normality assumption (3.3), enclosing 95% of observed variations in measurement.




Figure 41: Global $\bar{\mathbf{x}}^{\text{GLB}} \pm \mathbf{z}^{\text{T}} \mathbf{s}^{\text{GLB}}$ and participant-specific $\bar{\mathbf{x}}^{\text{REF}} \pm \mathbf{z}^{\text{T}} \mathbf{s}^{\text{REF}}$ kinematic references, with $\mathbf{z} = 1.96\mathbf{1}_K$, obtained for 3 lower limb DoFs of **S9** in 1.8 km/h walking at 10° incline.

right leg DoFs. Relevant variations between gait modalities were present in the form of higher amplitudes in knee *flexion/extension* and foot *abduction/adduction* for the faster 5.4 km/h gait, coupled with later and earlier peaks in each case, respectively; and positive amplitude offsets in knee *flexion/extension* and ankle *dorsiflexion/plantarflexion* when walking on a 10° incline, constant across the whole movement in the latter. For comparison purposes, Figure 41 presents $\bar{\mathbf{x}}^{\text{GLB}}$, in the case of walking at 5.4 km/h on a 10° incline, overlaid with the previous references for **S3**. Similarity between trajectory shapes was evident, as well as a recurring overlap between both confidence bands; on the other hand, both a marked amplitude offset (as much as 5° in the case of right ankle *abduction/adduction*) and slightly narrower confidence bands (notable in right knee *flexion/extension*), were observed.

Regarding the preliminary descriptive analysis of the experimental dataset, correlation matrices **C** are depicted in Figure 42, with DoF-specific averages listed in Table 8³⁹. With the exception of *moderate* correlation between knee *flexion/extension*, knee *internal/external rotation*, ankle *internal/external rotation* and ankle *abduction/adduction* in both legs, all remaining DoFs exhibited *weak* correlation, with a global average $\bar{r}_s$ of 0.373, suggesting *weak* interdimensional correlation across the dataset. When employing NMF to factorize of $\bar{\mathbf{x}}^{\text{GLB}}$, reconstruction Variance Accounted For (VAF) was found to increase with r_{NMF} ; with a threshold of 0.9, at least $r_{\text{NMF}} = 3$ components were required to accurately reconstruct global references (Figure 43). Based on Chapter 2 (2.3.2), for $N = 20$ basis and given $K = 18$, effective dimensionality reduction is achieved at $P \leq \lfloor \frac{KN}{2K+N} \rfloor \leq 6$. As such, and informed by physiological and kinesiological studies of human locomotion [d’Avella & Bizzi, 2005; Ivanenko et al., 2005], $P = 5$ primitives were assumed to describe the perceptual references with marginal error; in practice, this entails a dimensionality reduction of $\sim 22\%$, when compared to the conventional DMP structure *i.e.* $P = K$. The complete set of kinematic references, for all 9 non-discarded participants, is available in Appendix C: Stereotypical Kinematics in Human Locomotion.

³⁹Values 0.20-0.39 generally imply *weak* correlation, while the 0.40-0.59 range implies *moderate* correlation [Akoglu, 2018].




Figure 42: Illustration of Spearman's rank correlation r_s between different dimensions in different walking modalities; values closer to 1 and -1 imply greater correlation between different dimensions.

Table 8: Average magnitude of Spearman's rank correlation coefficient computed for each dimension in different walking modalities; higher values imply greater average pairwise rank correlation magnitude $|r_s|$.

k		1.8 km/h		5.4 km/h		Avg.
		0°	10°	0°	10°	
1	L Knee Flexion/Extension	0.412	0.519	0.356	0.419	0.426
2	L Knee Abduction/Adduction	0.234	0.278	0.250	0.262	0.256
3	L Knee Internal/External Rotation	0.490	0.570	0.415	0.488	0.491
4	L Ankle Abduction/Adduction	0.490	0.533	0.388	0.445	0.464
5	L Ankle Internal/External Rotation	0.490	0.539	0.419	0.353	0.450
6	L Ankle Dorsiflexion/Plantarflexion	0.313	0.458	0.367	0.418	0.389
7	L Ball Foot Abduction/Adduction	0.242	0.257	0.306	0.303	0.277
8	L Ball Foot Internal/External Rotation	0.295	0.165	0.218	0.227	0.226
9	L Ball Foot Flexion/Extension	0.293	0.223	0.368	0.330	0.303
10	R Knee Abduction/Adduction	0.298	0.448	0.162	0.376	0.321
11	R Knee Internal/External Rotation	0.500	0.567	0.434	0.440	0.485
12	R Knee Flexion/Extension	0.357	0.538	0.306	0.416	0.404
13	R Ankle Abduction/Adduction	0.478	0.555	0.370	0.437	0.460
14	R Ankle Internal/External Rotation	0.483	0.551	0.410	0.382	0.457
15	R Ankle Dorsiflexion/Plantarflexion	0.326	0.507	0.297	0.429	0.390
16	R Ball Foot Abduction/Adduction	0.289	0.287	0.316	0.331	0.306
17	R Ball Foot Internal/External Rotation	0.331	0.253	0.336	0.239	0.290
18	R Ball Foot Flexion/Extension	0.371	0.261	0.367	0.283	0.320
		0.372	0.417	0.338	0.365	0.373




Figure 43: Variance Accounted For (VAF) between original and reconstructed patterns obtained when factorizing $\mathbf{q}$ with $r_{\text{NMF}} \in \{1, 2, 3, 4, 5, 6\}$ components; higher values imply better reconstruction quality; $\text{VAF} \geq 0.9$ was assumed to represent an accurate reconstruction.

3.5.2 Memory Consolidation

Figure 44 depicts the evolution of $J(\theta, \bar{\mathbf{x}}^{\text{REF}})$ and $J(\theta, \mathbf{s}^{\text{REF}})$ (23) and the respective convergence times $I_{\mathbf{x}}^{\text{conv}}$ and $I_{\mathbf{s}}^{\text{conv}}$ with different learning strategies (Table 7) during expansive learning of $\mathbf{M}^{\mathbf{s}^7}$, in the specific case of 5.4 km/h walking on a 0° incline ($\mathcal{M}^3$, Figure 39). Convergence was achieved noticeably faster with the isolated strategy, at around 300 k and 250 k generations⁴⁰ respectively. In contrast, learning the full set of $P = 5$ PPs from scratch in both the bulk approach and the initial baseline learning in the two-stage strategy registered highest $I_{\mathbf{x}}^{\text{conv}}$ and $I_{\mathbf{s}}^{\text{conv}}$, at roughly 450-500 k, while tuning the learned baseline to each test subset required approximately 350 k generations in both cases, with considerably smaller initial costs and flatter decrease in cost. Table 9 presents global averages across all cross-validation folds *i.e.* all 9 non-discarded participants; similar trends were observed *i.e.* lower compound costs and convergence times in the isolated approach, and increasingly higher values in tuning, baseline and bulk strategies, highlighting the association between learning performance, problem dimensionality and level of prior information. In both $\bar{\mathbf{x}}^{\text{REF}}$ and $\mathbf{s}^{\text{REF}}$, at $\alpha = 0.05$, statistically significant⁴¹ decreases in cost and convergence time were observed between the isolated strategy and the alternatives ($p < .00001$ in all cases); similarly, tuning also converged significantly faster than baseline and bulk learning ($p < .00001$ in all cases); on the other hand, no statistical significant difference was observed between final costs obtained with the bulk and two-stage strategies ($p = 0.647$ and $p = 0.058$).

Overall, $J(\theta, \mathbf{s}^{\text{REF}})$ was lower than $J(\theta, \bar{\mathbf{x}}^{\text{REF}})$ for the same walking modality and learning strategy. Higher averages registered, in general, slightly smaller average convergence times, possibly a result of premature convergence; this was particularly evident in 5.4 km/h walking on a 0° incline. Nonetheless, estimated errors per reference sample $e_{\mathbf{x}}$ and $e_{\mathbf{s}}$ were below 0.1° and 0.5° , respectively, in all walking velocities, treadmill inclines and learning strategies (Table 9). Likewise, measured VAF was ≥ 0.88 in all cases, supporting the general observation that memory consolidation was consistently accurate regardless

⁴⁰With each optimization step conducted sequentially; concurrent/parallel learning was not explored.

⁴¹Using Welch's two-tailed two-sample t-test.

of learning strategy and trial variables; isolated learning achieved lowest errors per sample (as low as 0.02°) and highest VAF at 1.0, implying a perfect or near-perfect reconstruction of reference trajectories $\mathbf{x}^{\text{REF}}$ and $\mathbf{s}^{\text{REF}}$. This is illustrated in Figure 45 for the specific case of **S3** walking at 1.8 km/h on a 0° incline; despite high $\text{VAF}_{\mathbf{x}}$ and $\text{VAF}_{\mathbf{s}}$, as well as low $e_{\mathbf{x}}$ and $e_{\mathbf{s}}$, localized deviations were observed when $\mathbf{s}^{\text{REF}}$ is near 0, as is the case of left and right foot *abduction/adduction*. In contrast, a greater overlap between all trajectories was achieved for both knee *flexion/extension* inputs.

For comparison purposes, Figure 46 illustrates the set of $P = 5$ PPs shared across the $K = 18$ inputs in the case of **S3** and **S8**; notably, all PPs exhibited highly erratic and fluctuating patterns, in line with the resulting trajectories in Figure 45; on the other hand, some degree of similarity can be seen between participants, hinting at potential inter-participant PPs sharing. Lastly, individual values in $\mathbf{\Lambda}$ were generally small, contrasting with the higher amplitude ranges $[-150, 150]$ in f^p , while $\mathbf{\Delta}$ were noticeably closer to preset bounds $[-\pi, \pi]$.

The complete set of learning results, for all 9 non-discarded subjects, is available in Appendix C: Stereotypical Kinematics in Human Locomotion.

3.5.3 Gait Classification

Table 10 presents detection *sensitivity*, *specificity* and *accuracy* measures obtained when employing Algorithm 4 in each cross-validation fold i.e. with participant-specific memory libraries $\mathbf{M}^{\mathbf{s}^j}$, $j \in \{1, 3, 4, 5, 7, 8, 9, 10, 11\}$. Generally, sensitivity was found to decrease with velocity and incline, despite varying considerable between participants, with global averages ranging from roughly 0.5 to >0.8 , as well as standard deviation >0.1 in all walking modalities. A noted low sensitivity was measured for 5.4 km/h walking at 10° incline in specific participants, with **S3**, **S4** and **S10** lowering the global average to 0.54. In contrast, very high specificity was observed in all cases, averaging ≥ 0.94 with a standard deviation ≤ 0.07 across different gaits; similarly, high accuracy⁴² was observed in all gait types - with averages ≥ 0.87 and standard deviation ≤ 0.05 . Global averages of 0.70 ± 0.17 , 0.96 ± 0.04 and 0.89 ± 0.04 were obtained for sensitivity, specificity and accuracy respectively, implying a highly specific detection with regard to walking modality, with the perceptual memory model (3.3) able to uniquely encode stereotypical perceptual patterns, and attesting to the practical effectiveness of Algorithm 4 in gait classification.

3.5.4 Detection of Non-Stereotypical Gait

Analogous classification metrics obtained in the detection of perturbed - or otherwise non-stereotypical - walking are presented in Table 11 for different walking modalities, and in Table 12 for different types of perturbation (3.4.3.2). While, in general, performance was worse than in gait classification, in all cases average accuracy was ≥ 0.76 , with a global average of 0.79. Similar to the previous problem, higher specificity than sensitivity was registered, with values ranging from 0.77 to 0.91 in the first case and 0.50 to 0.73 in the second. An increase in specificity and accuracy variance was observed between different

⁴²Due to a high negative-to-positive ratio in the dataset in classification, accuracy is greatly influenced by specificity values.




Figure 44: Evolution of $J(\theta, \bar{\mathbf{x}}^{\text{REF}})$ (**Top**) and $J(\theta, \mathbf{s}^{\text{REF}})$ (**Bottom**) (23) during memory consolidation with the skill learning procedure (Chapter 2, 2.3.3) with **S7** walking at 5.4 km/h and a 0° incline under different learning strategies (Table 7); in the isolated learning case, dimension-specific costs $J_k(\theta, \bar{\mathbf{x}}^{\text{REF}})$ are sequentially compounded; final values presented in Table 9.

Table 9: Memory consolidation results across all 9 cross validation folds, including convergence time I^{conv} , final costs $J(\theta^{\text{best}}, \bar{\mathbf{x}}^{\text{REF}})$ and $J(\theta^{\text{best}}, \mathbf{s}^{\text{REF}})$, estimated errors per sample $\hat{e}_{\mathbf{x}} = 1/KL \times J(\theta^{\text{best}}, \bar{\mathbf{x}}^{\text{REF}})$ and $\hat{e}_{\mathbf{s}} = 1/KL \times J(\theta^{\text{best}}, \mathbf{s}^{\text{REF}})$, and measured VAF $_{\mathbf{x}}$ and VAF $_{\mathbf{s}}$, with different learning strategies (Table 7).

	Velocity	Incline	$I_{\mathbf{x}}^{\text{conv}}$	$J(\theta^{\text{best}}, \bar{\mathbf{x}}^{\text{REF}})$	$\hat{e}_{\mathbf{x}}$	VAF $_{\mathbf{x}}$
Isolated	1.8 km/h	0°	305.10±9.20 k	93.46±4.23	0.05	1.00±0.00
		10°	303.59±7.61 k	67.31±3.38	0.04	1.00±0.00
	5.4 km/h	0°	298.91±16.69 k	141.93±4.95	0.08	1.00±0.00
		10°	296.94±17.44 k	127.74±5.46	0.07	1.00±0.00
Bulk	1.8 km/h	0°	485.46±14.29 k	536.02±93.05	0.30	0.89±0.02
		10°	479.38±7.34 k	450.60±45.25	0.25	0.91±0.02
	5.4 km/h	0°	438.97±109.61 k	743.26±217.18	0.41	0.90±0.02
		10°	479.05±4.86 k	608.54±56.86	0.34	0.90±0.03
Two-stage	Baseline	1.8 km/h	472.77 k	503.42	0.28	0.92
			484.17 k	446.49	0.25	0.93
		5.4 km/h	481.08 k	591.22	0.33	0.91
			488.04 k	475.46	0.26	0.90
	Tuning	1.8 km/h	351.34±4.62 k	542.10±54.65	0.30	0.89±0.01
			352.01±4.39 k	449.86±60.79	0.25	0.92±0.02
		5.4 km/h	347.09±4.49 k	699.78±71.89	0.39	0.88±0.02
			353.00±6.09 k	585.49±63.14	0.33	0.90±0.02
	Velocity	Incline	$I_{\mathbf{s}}^{\text{conv}}$	$J(\theta^{\text{best}}, \mathbf{s}^{\text{REF}})$	$\hat{e}_{\mathbf{s}}$	VAF $_{\mathbf{s}}$
Isolated	1.8 km/h	0°	244.14±9.59 k	35.98±3.08	0.02	1.00±0.00
		10°	251.86±10.99 k	34.85±1.98	0.02	0.99±0.00
	5.4 km/h	0°	233.98±12.69 k	42.71±1.65	0.02	0.99±0.00
		10°	234.82±12.36 k	38.96±2.13	0.02	1.00±0.00
Bulk	1.8 km/h	0°	489.43±9.75 k	102.83±12.53	0.06	0.96±0.01
		10°	487.59±7.90 k	101.88±14.65	0.06	0.95±0.01
	5.4 km/h	0°	483.90±11.46 k	121.05±11.50	0.07	0.95±0.01
		10°	487.57±5.72 k	99.64±11.27	0.06	0.95±0.01
Two-stage	Baseline	1.8 km/h	490.11 k	101.59	0.06	0.96
			477.81 k	116.23	0.06	0.95
		5.4 km/h	499.47 k	108.43	0.06	0.96
			486.78 k	88.80	0.05	0.97
	Tuning	1.8 km/h	358.66±4.53 k	114.60±10.59	0.06	0.95±0.01
			361.31±5.30 k	108.46±12.13	0.06	0.94±0.02
		5.4 km/h	356.14±5.41 k	128.57±16.39	0.07	0.94±0.02
			361.36±4.64 k	102.69±9.27	0.06	0.96±0.01




Figure 45: Comparison between original references $\{\bar{\mathbf{x}}^{\text{REF}}, \mathbf{s}^{\text{REF}}\}$ (Figure 40) and trajectories learned with the isolated (no shared PP, $P = 18$) and two-stage ($P = 5$) strategies (Table 7), for 6 lower limb DoFs of **S3** during walking at 1.8 km/h on a 0° incline.




Figure 46: Resulting set of $P = 5$ PPs and skill $\mathcal{S}^*$ (16, 66) describing $\bar{\mathbf{x}}^{\text{REF}}$ learned with the two-stage strategy (Table 7) for walking at 1.8 km/h on a 0° incline with **S3** (Top) and **S7** (Bottom); dimension labels are provided in Table 8.

Table 10: Classification metrics (*sensitivity*, *specificity* and *accuracy*) of context identification using Algorithm 4 with participant-specific perceptual libraries $\mathbf{M}^{\mathbf{S}^j}$, as well as global averages ($\pm$ std. deviation).

	Velocity	Incline	Participant									Average
			S1	S3	S4	S5	S7	S8	S9	S10	S11	
Sensitivity	1.8 km/h	0°	0.93	0.72	0.82	0.82	0.96	0.98	0.94	0.98	0.57	0.84 $\pm$ 0.14
		10°	0.81	0.89	0.56	0.54	0.53	0.92	0.80	0.95	0.54	0.66 $\pm$ 0.22
	5.4 km/h	0°	0.96	0.74	0.91	0.94	0.86	0.67	0.97	0.73	0.78	0.78 $\pm$ 0.13
		10°	0.81	0.23	0.45	0.82	0.62	0.44	0.92	0.38	0.46	0.54 $\pm$ 0.21
	0.70 $\pm$ 0.17											
Specificity	1.8 km/h	0°	0.98	0.99	0.83	0.98	0.97	0.92	0.99	0.90	0.99	0.94 $\pm$ 0.06
		10°	0.96	0.86	0.99	0.99	1.00	0.87	1.00	0.91	0.99	0.94 $\pm$ 0.07
	5.4 km/h	0°	0.97	0.95	0.94	0.94	0.99	1.00	0.99	1.00	0.97	0.97 $\pm$ 0.03
		10°	0.97	0.99	0.94	0.94	0.98	1.00	0.96	0.99	0.99	0.97 $\pm$ 0.03
	0.96 $\pm$ 0.04											
Accuracy	1.8 km/h	0°	0.96	0.94	0.82	0.94	0.97	0.93	0.98	0.91	0.87	0.92 $\pm$ 0.05
		10°	0.92	0.87	0.87	0.88	0.88	0.88	0.95	0.92	0.87	0.87 $\pm$ 0.04
	5.4 km/h	0°	0.97	0.90	0.94	0.94	0.96	0.92	0.98	0.93	0.92	0.92 $\pm$ 0.03
		10°	0.93	0.75	0.81	0.91	0.89	0.86	0.95	0.84	0.90	0.87 $\pm$ 0.05
	0.89 $\pm$ 0.04											




Figure 47: Illustration of the gait classification procedure (3.4.4) in a trial with **S4** at 5.4 km/h on a 10° incline, with two perturbations during left heel strike ($\square$, 3.4.3.2); captured lower limb kinematics $\mathbf{q}$ are used as perceptual input $\mathbf{x}^*$; upper and lower edges indicate perceptual memory in $\mathcal{C}_t$ (70) when $\mathcal{C}_t \neq \emptyset$ (Algorithm 4); detection matches at **P0** and **P1**, respectively 16 ms after and 16 ms before ground truth, corresponding to 1 sample, at a sampling frequency of 60 Hz.

Table 11: Classification metrics (*sensitivity*, *specificity* and *accuracy*) and *detection time* (s) obtained in the detection of perturbed walking using Algorithm 4 with participant-specific perceptual libraries $\mathbf{M}^{s_j}$, as well as global averages ($\pm$ std. deviation), grouped by walking velocity and terrain incline.

	Velocity	Incline	Participant									Average
			S1	S3	S4	S5	S7	S8	S9	S10	S11	
Sensitivity	1.8 km/h	0°	0.40	0.14	0.57	0.54	0.45	0.58	0.27	0.45	0.49	0.50 $\pm$ 0.22
		10°	0.49	0.64	0.59	0.61	0.83	0.33	0.69	0.71	0.54	0.60 $\pm$ 0.18
	5.4 km/h	0°	0.43	0.81	0.65	0.65	0.54	0.69	0.67	0.74	0.76	0.73 $\pm$ 0.18
		10°	0.49	0.25	0.91	0.83	0.56	0.45	0.84	0.46	0.82	0.70 $\pm$ 0.21
												0.63 $\pm$ 0.20
Specificity	1.8 km/h	0°	0.95	0.97	0.95	0.97	0.96	0.96	0.94	0.91	0.97	0.91 $\pm$ 0.13
		10°	0.96	0.85	0.95	0.93	0.96	0.97	0.88	0.78	0.89	0.78 $\pm$ 0.20
	5.4 km/h	0°	0.91	0.91	0.93	0.89	0.85	0.94	0.88	0.91	0.95	0.81 $\pm$ 0.21
		10°	0.96	0.93	0.72	0.93	0.80	0.77	0.75	0.78	0.66	0.77 $\pm$ 0.13
												0.82 $\pm$ 0.17
Accuracy	1.8 km/h	0°	0.91	0.85	0.92	0.86	0.91	0.89	0.87	0.79	0.90	0.86 $\pm$ 0.10
		10°	0.92	0.81	0.91	0.89	0.95	0.87	0.85	0.78	0.82	0.76 $\pm$ 0.16
	5.4 km/h	0°	0.83	0.89	0.91	0.86	0.83	0.91	0.81	0.88	0.92	0.80 $\pm$ 0.15
		10°	0.91	0.83	0.74	0.91	0.76	0.71	0.76	0.69	0.72	0.76 $\pm$ 0.09
												0.79 $\pm$ 0.13
Det. Time	1.8 km/h	0°	0.02	0.07	-0.02	0.06	0.02	-0.01	0.02	-0.14	-0.02	-0.01 $\pm$ 0.09
		10°	-0.03	-0.07	-0.07	0.08	0.03	0.23	-0.02	-0.23	-0.09	-0.02 $\pm$ 0.14
	5.4 km/h	0°	0.07	-0.01	0.01	0.01	0.00	0.06	-0.10	-0.10	-0.04	-0.01 $\pm$ 0.07
		10°	-0.04	-0.07	-0.14	0.03	-0.01	0.00	-0.01	-0.01	-0.19	-0.06 $\pm$ 0.09
												-0.03 $\pm$ 0.09

Table 12: Classification metrics (*sensitivity*, *specificity* and *accuracy*) and *detection time* (s) obtained in the detection of perturbed walking using Algorithm 4 with participant-specific perceptual libraries $\mathbf{M}^{s_j}$, as well as global averages ($\pm$ std. deviation), grouped by type of perturbation.

	Leg	Event	Participant									Average
			S1	S3	S4	S5	S7	S8	S9	S10	S11	
Sensitivity	R	TO	0.40	0.49	0.43	0.487	0.45	0.83	0.54	0.56	0.45	0.60 $\pm$ 0.23
		HS	0.14	0.64	0.81	0.246	0.33	0.45	0.36	0.70	0.53	0.55 $\pm$ 0.24
	L	TO	0.57	0.59	0.65	0.908	0.58	0.69	0.69	0.84	0.49	0.68 $\pm$ 0.18
		HS	0.54	0.61	0.65	0.828	0.27	0.71	0.67	0.46	0.72	0.69 $\pm$ 0.21
Specificity	R	TO	0.95	0.96	0.91	0.96	0.96	0.96	0.85	0.80	0.91	0.83 $\pm$ 0.15
		HS	0.97	0.85	0.91	0.93	0.97	0.77	0.93	0.75	0.96	0.83 $\pm$ 0.14
	L	TO	0.95	0.95	0.93	0.72	0.96	0.88	0.94	0.75	0.97	0.81 $\pm$ 0.19
		HS	0.97	0.93	0.89	0.93	0.94	0.78	0.88	0.78	0.76	0.79 $\pm$ 0.21
Accuracy	R	TO	0.91	0.92	0.83	0.91	0.91	0.95	0.83	0.76	0.79	0.81 $\pm$ 0.13
		HS	0.85	0.81	0.89	0.83	0.87	0.71	0.83	0.74	0.90	0.79 $\pm$ 0.11
	L	TO	0.92	0.91	0.91	0.74	0.89	0.85	0.91	0.76	0.90	0.80 $\pm$ 0.15
		HS	0.86	0.89	0.86	0.91	0.87	0.78	0.81	0.69	0.75	0.78 $\pm$ 0.15
Det. Time	R	TO	0.02	-0.04	0.07	-0.04	0.02	0.03	0.00	-0.01	-0.14	-0.01 $\pm$ 0.08
		HS	0.07	-0.07	-0.01	-0.07	0.23	0.00	0.13	0.08	0.17	0.07 $\pm$ 0.12
	L	TO	-0.02	-0.07	0.01	-0.14	-0.01	-0.02	0.06	-0.01	-0.02	-0.03 $\pm$ 0.07
		HS	0.06	0.08	0.01	0.03	0.02	-0.23	-0.10	-0.01	-0.08	-0.02 $\pm$ 0.11

participants, with global averages of 0.13. Once again, treadmill incline was found to have a relevant impact on detection performance, with 10° walking exhibiting lower sensitivity and specificity - reflected in lower overall accuracy - than 0° walking; conversely, higher sensitivity and lower specificity were measured in 5.4 km/h walking when compared with slower 1.8 km/h walking, suggesting that in faster locomotion, deviations from normal gait were more effectively identified. Metrics were generally consistent between different types of perturbations, although with a noted sensitivity increase during left leg events (Table 12).

Measured detection times were generally short, often <0 , denoting detections that occurred ahead of time *i.e. before* the ground truth marker in the dataset; on the other hand, high variance was registered between subjects and trials, with standard deviations above average detection time, implying that detection regularly occurred *after* affecting gait; nevertheless, these values were still ≤ 0.14 , with a global average detection time of -0.03 s, representing - at most - a short delay with minimal impact in the majority of practical scenarios. A decrease in detection time was also observed in 5.4 km/h gait, in line with the increased sensitivity (0.70-0.73) in Table 11. Conversely, substantial detection delays were observed in the case of perturbations induced during right heel strike (Table 12); **S7**, **S9** and **S11**, in particular, registered averages of 0.14, 0.13 and 0.17 s respectively, contrasting with values <0 obtained for perturbations induced at right toe off and left leg events. Globally, Algorithm 4 was able to consistently identify gait perturbations from 0.11 s *before* to 0.06 s *after* noticeable impact to sensor patterns, while, on average, detection occurs 0.03 s ahead of time. Figure 47 depicts the detection procedure in a specific trial, wherein both accurate gait classification and detection of perturbed walking is illustrated.

Full classification results are provided in Appendix C: Stereotypical Kinematics in Human Locomotion.

3.6 Discussion

The cross-validation strategy (3.4) employed in this study allowed an objective evaluation of the performance and practicability of the perceptual memory model (3.3) within the constraints of a real-world dynamic task, and demonstrated how such episodic memories, when combined with primitive theory and powerful practical frameworks *i.e.* DMPs, provide a flexible and robust way to model perception in artificial automated systems. Within the context of human locomotion analysis, walking modality - velocity and terrain inclination - was accurately detected using the practical algorithms described in 3.3.2, as were transient gait disturbances induced during the experimental protocol.

The kinematic references $\bar{\mathbf{x}}^{\text{REF}}$ (Figure 40) obtained for the different types of gait reproduced in the experimental trials were within expectations: on one hand, a higher amplitude implies a greater range of motion and consequently larger steps, in line with faster walking; on the other, a positive offset, while preserving the overall motion shape, in DoFs of joints operating on the sagittal plane - knee *flexion/extension* and ankle *dorsiflexion/plantarflexion* - is characteristic in forward incline walking. In addition, this phenomenon was more pronounced in the ankle than the knee or foot joints, suggesting such terrain alterations to be accommodated primarily at the ankle. Overall, beyond the uniformity and periodicity of recorded kinematics $\mathbf{q}_k$, the morphological similarity and temporal synchronicity observed across different $\bar{\mathbf{x}}^{\text{REF}}$, coupled with the overall relatively low variance of $\mathbf{s}^{\text{REF}}$, attest to the the uniformity and periodicity of

the movement, as well as the accuracy and efficiency of the data processing procedure (3.4.3.3), particularly the segmentation stage; notwithstanding the heuristic nature of the approach, these results suggest its applicability can extend beyond the current dataset, to generalized segmentation of periodic and stereotypical patterns. Conversely, despite preserving the overall trajectory shape, some disparities were noted between global and participant-specific references (Figure 41), in the form of variations in amplitude offset - illustrated in the right ankle *abduction/adduction* angle - accompanied by a slight decrease in variability - specifically notable in right knee *flexion/extension*; while the latter was expected when considering the - albeit marginally - smaller data pool used during LOO cross-validation, these suggest that human walking, even when executed in a controlled supervised environment, at a fixed velocity and incline, still exhibits subjective nuances that may be difficult to encode in inter-participant memories.

Overall, despite the offline nature of the approach described in this chapter, $\{\bar{\mathbf{x}}^{\text{REF}}, \mathbf{s}^{\text{REF}}\}$ can be seen as time-invariant working memories of a specific gait type, capturing the stereotypical nature of the underlying motion across multiple repetitions; this parallel with human perception is particularly fitting in situations where input $\mathbf{q}$ represents some form of a short-term data buffer, as is common in real-time robotic control. On the other hand, beyond the practical drawbacks associated with a discretized representation, directly storing $\bar{\mathbf{x}}^{\text{REF}}$ and $\mathbf{s}^{\text{REF}}$ as indexed sequences would require, in practice, at least $2L = 200$ values to be archived, which may pose a practical challenge when dealing with extensive and continuously growing perceptual libraries.

The preliminary analysis conducted on $\bar{\mathbf{x}}^{\text{GLB}}$ demonstrated relatively weak correlation among different perceptual inputs (Figure 42, Table 8), which, despite not expressly implying statistical independence, reinforces the applicability of the perceptual model in this particular dataset, while still hinting at the existence of elemental primitives that can be shared without loss of accuracy in representation. This was evident during the factorization step, where as few as $r_{\text{NMF}} = 3$ components were sufficient to describe $\bar{\mathbf{x}}^{\text{GLB}}$ with over 90% accuracy; although this value could not be directly translated into a number of shared PPs, due to incorporation of DMP dynamics, it illustrated the potential dimensionality reduction achievable with a PP-based representation. Moreover, it aligns with the results obtained in Chapter 2 (2.5.1), where $P = 3$ MPs were able to efficiently reproduce bipedal walking. In contrast, $P = 5$ PPs represented an over-conservative approach; on the other hand, it allowed a more robust and flexible representation of the perceptual input, at the cost of introducing some level of redundancy. Within the context of **RQ2**, the primary objective of this study was to ascertain the capability of the dynamic model to encode stereotypical patterns, rather than to minimize dimensionality and optimize data efficiency. Nonetheless, it's highlighted that further simplification could have been achieved with $P = 3$ or $P = 4$. Additionally, given the nature of the experiment and the modality of the recorded data, an overlap between any observed PPs and natural MPs in human locomotion is plausible; notwithstanding, a general, high-level view of artificial perception frames body kinematics - whether external or self-perceived - as sensor input, which is relevant in dynamic environments shared between humans and machines.

All of the 3 learning strategies (Table 7) employed during memory consolidation were successful at encoding reference patterns to varying degrees of accuracy and efficiency (Table 9), as illustrated in Figure 45. Following a *divide-and-conquer* approach through isolated optimization subproblems produced the

best results, both in terms of convergence time and solution quality, when compared to the alternatives. This is in line with the significant decrease in n - from a total of $2KP + PN$ to N (3.4.4) - achieved by fitting each dimension separately; this simplification was specially effective when taking into account the computational overhead of each update in CMA-ES, which scales according to n^3 [Hansen et al., 2003; Hansen & Ostermeier, 2001]; moreover, while not adopted in this work, the independent nature of each optimization step allows concurrent execution for all K inputs, further streamlining memory consolidation. On the other hand, the isolated approach disregarded the presence of shared PPs among perceptual inputs, making the PP-based structure of the model redundant, and decreasing storage efficiency⁴³, when compared to alternatives defined over a reduced number of PPs. As such, isolated learning is viewed as a practically efficient but ultimately conceptually limited strategy to consolidate $\{\bar{\mathbf{x}}^{\text{REF}}, \mathbf{s}^{\text{REF}}\}$ into perceptual memories $\mathcal{M}^m$, whereby existing knowledge contributes relatively little to memory acquisition; notwithstanding, it represents a fast and straightforward solution to exploit the perceptual model (3) and practical methods (3.3.2) under temporal constraints or limited computational resources; more intricate cognitive strategies can implement methods for further reconsolidation or refinement of learned memories e.g. identifying shared PPs across memories in $\mathbf{M}$ (65), following an initial isolated approach.

Regarding the strategies where $P = 5$, both bulk learning of $\mathbf{M}^{\mathbf{s}^j}$ and the initial baseline step in the two-stage strategy (Table 7) registered, as expected, similar performance, as the underlying optimization steps - in terms of priors (or their absence), parametrization and dimensionality - were effectively identical. Resulting discrepancies can be assumed to have been caused by data variations - as highlighted in Figure 41 - or by stochasticity introduced by the learning algorithm. In the latter case, tuning the global baseline to participant-specific references in each cross-validation fold was shown to be achieved significantly faster (Figure 44, Table 9). In practice, decreasing exploration and focusing on local small-step updates led to a flatter decline in $J(\boldsymbol{\theta}, \bar{\mathbf{x}}^{\text{REF}})$ and $J(\boldsymbol{\theta}, \mathbf{s}^{\text{REF}})$, but the lower initial costs allowed optimization to converge, on average, after around 350 k generations, close to the values registered for the isolated strategy. This reinforces the practical appeal of guiding optimization with *a priori* information; extending this notion to autonomous L2 systems, analogous baselines can be provided by similar (69) existing memories, and the tuning procedure can gradually be replaced with intrinsic learning ($n = 2KP$, 2.3.3) as $\mathbf{M}$ grows, favoring extrapolating from similar experiences over learning new PPs. When looking for a reasonable trade-off between learning dimensionality, representation versatility and computational efficiency, it is understood that while PP-sharing strategies increase the complexity of memory acquisition during expansive learning i.e. early cognitive stages, they constitute powerful long-term alternatives in continually growing perceptual archives, by exploiting existing information and providing more data-efficient representations.

Figure 45 highlights the accuracy of the consolidation procedure, with a consistent overlap between reference patterns and learned DMP trajectories. In the particular case of two-stage learning of foot *abduction/adduction* ($k = 9, 18$) in $\mathbf{S3}$, local inaccuracies in flatter, null-amplitude segments of $\bar{\mathbf{x}}$ align with increased λ_9^2 , λ_9^5 and λ_{18}^2 in $\mathbf{\Lambda}$, as well as the high-frequency nature of f^p (Figure 46, **Top**). The latter can be explained by $N = 20$, which leads to a higher variability in possible shapes, despite allowing higher trajectory resolution; the former, despite suggesting premature convergence to suboptimal solutions, did

⁴³With $P = K$, at least $KN = 360$ values need to be archived, in contrast with $2KP + PN = 280$ with $P = 5$.

not exceed 1° over the whole period, in line with estimated errors per sample and measured VAF (Table 9). Generally, these variations agree with variable ranges of motion in different DoFs: the cumulative nature of the cost function (23) entails that inputs with greater amplitude drive optimization, resulting in varying accuracy across different dimensions; normalizing $\bar{\mathbf{x}}$ and $\mathbf{s}$ prior to learning⁴⁴ can help mitigate this issue. Ultimately, it would allow *one-shot* memory consolidation by extending (23) to:

$$J^*(\boldsymbol{\theta}, \bar{\mathbf{x}}^{\text{REF}}, \mathbf{s}^{\text{REF}}) = J(\boldsymbol{\theta}, \bar{\mathbf{x}}^{\text{REF}}) + J(\boldsymbol{\theta}, \mathbf{s}^{\text{REF}}) \quad (71)$$

Notably, elements in Δ were notably close to preset bounds $[-\pi, \pi]$ (Figure 46), suggesting over-estimation and excessive step sizes during optimization, as a result of the largely unbalanced parameter ranges. In fact, heterogeneous parameter sets such as $\boldsymbol{\theta} = [\mathbf{W}, \Lambda, \Delta]^\top$ (2.3.3) generally require normalization (or alternative transformations e.g. logarithmic scaling) [Deb, 2011; Hansen, 2016]; overlooking this step risks making learning Δ redundant in most composite learning strategies beyond sequential learning where it's the sole target of optimization. It should also be noted that the learning procedure was hand-tuned through trial-and-error until reasonable results were consistently achieved; different parameter sets can provide better results, with a more comprehensive benchmark required for optimal performance [Hansen, 2016]. Furthermore, Algorithm 3, 4 and 5 are independent of the learning strategy and underlying structure i.e. a fixed number of sensor dimensions K was assumed for each memory in $\mathbf{M}$, regardless of primitive sharing (3.3.2). Lastly, some level of similarity was observed between f^p obtained for **S3** and **S7**, as evidenced in Figure 46, suggesting the sharing of PPs among different participants, in line with $r_{\text{NMF}} \leq 3$ (Figure 43) and the nature of the movement [d'Avella & Bizzi, 2005] i.e. described with a reduced number of locomotor MPs.

Gait type identification posed a challenging problem due to the overall similarity and frequent overlap between nominal ranges across different walking modalities (Figure 40). This likely explains the moderate sensitivity measured, decreasing with velocity and incline (Table 10); it is therefore hypothesized that, at higher velocities, it becomes more difficult to identify specific perceptual contexts, although without an increase in false detections, as evidenced by high specificity regardless of gait type; while this was expected due to the added instability of faster gait, increasing z_k (68) may constitute a way to counterbalance this effect, or potentially increasing detection sensitivity at the cost of lower specificity. In practice, z_k can be set in order to fit different application requirements and/or constraints. Overall, the achieved results describe a highly specific detection in regard to gait type, where small deviations from nominal sensor input were sufficient to trigger a *null* match i.e. $\mathcal{C}_t = \emptyset$, which can be relevant in identification of obstacles or disturbances; nevertheless, the detection algorithm was still able to consistently identify a valid memory throughout the majority of the trials; with a global accuracy average of 0.89, it can be concluded that the perceptual model was able to uniquely encode stereotypical perceptual patterns in different gait conditions; moreover, the effectiveness of Algorithm 4 in gait classification is highlighted.

In line with these results, an average accuracy of 0.79 was obtained in the alternative detection of perturbed walking, despite slightly lower sensitivity and specificity values (Table 11). Higher sensitivity and

⁴⁴Followed by dimension-specific rescaling post-optimization.

lower specificity were measured in 5.4 km/h walking, suggesting that in faster locomotion, deviations from normal gait were more effectively identified; this can potentially be seen as a byproduct of lower sensitivity of context detection in such scenarios, as perturbations are assigned in the absence of matching memories, thus expected to follow an inverse trend. On the other hand, despite fast detection in different velocities, inclines and types of perturbation (Table 11, Table 12), the procedure employed to quantify detection times is dependent on labeled ground truth information in the data, thus limited to supervised validation scenarios; practical solutions to automate this detection are needed, such as enforcing a minimum time- or sample-based threshold, as previously implemented in similar works [André et al., 2016b]. In practice, advance detections can be caused anticipatory movements by the participant, noise in sensor data, or, potentially, delays on ground-truth perturbation time stamps, expected to some extent given the visual nature of the labeling procedure (3.4.3.2). Overall, despite the general nature of the approach, it was able to accurately and timely detect perturbed or non-stereotypical gait, with performance on par with automated alternatives in the literature [Mioskowska et al., 2020; Monaco et al., 2017; Trkov et al., 2017].

On top of these results, an answer to **RQ2** can be objectively formulated: repurposing versatile primitive-based motor control frameworks such as DMPs allows stereotypical perceptual information to be encoded into compact, data-efficient and time-invariant episodic memories, as well as accommodate for PPs shared across different dimensions, memories and systems. This is achieved through coupled multidimensional DMP systems capturing sensor variability across multiple repetitive episodes.

3.7 Conclusion

This study proposed a generic, amodal, data-driven and dynamical-system-based model to represent stereotypical perceptual patterns through episodic memories, repurposing the control structure presented in Chapter 2, and exploiting several of the appealing traits of the DMP framework such as time invariance, dynamical stability and parameter linearity. Notably, its high-level modular structure allows PPs to be shared among different dimensions and memories, which can ultimately accelerate memory formation and consolidation, on top of simplifying representation. Practical approaches to memory acquisition based on LfD strategies, as well as straightforward algorithms for perceptual context identification are introduced and validated on an experimental dataset of human locomotion (3.4.3.2). Results demonstrate that the proposed solutions were not only able to identify gait type (velocity and inclination) with high accuracy, but also deviations from nominal or stereotypical behavior in a timely fashion, understood to be crucial in real-time applications where corrective or mitigating measures need to be undertaken.

In essence, this work constitutes a valuable foundation for automated sensor-driven strategies ranging from simple monitoring problems (as demonstrated herein) to complex control scenarios that require effective perceptuomotor coupling, particularly when coupled with the control strategy introduced in the previous chapter; moreover, it is built upon insights from biological organisms and exhibits key features of the Lifelong Learning paradigm [Kudithipudi et al., 2022], such as 1) *transfer and adaptation* and 2) *exploiting task similarity* by using existing memories as initial baselines in memory acquisition or potentially sharing PPs across multiple memories; 3) *overcoming catastrophic forgetting* through general libraries of

perceptual memories in **M** and **L**; 4) *noise tolerance* by representing sensor data as a statistical distribution of values and through stable DMP dynamics; and 5) *resource efficiency* through its parametric PPs-based nature. Additionally, the proposed model was shown to be capable of *automated acquisition through observation*, *managing uncertainty*, and *fast inference*, features highlighted by Tenorth & Beetz [2009] as requirements for flexible reasoning and knowledge processing. Moreover, it partially aligns with Žabkar & Leonardis [2016] requirements for representation of motor memories, by being *compact* and *reusing information across different domains* i.e. *tasks*. Overall, it exemplifies how shared PPs can contribute to more efficient episodic memories.

With a measured accuracy of 89% in gait classification, and an effective dimensionality reduction of $\sim 22\%$ when sharing PPs across different inputs, an objective answer to **RQ2** was obtained: *coupled DMP systems can be harnessed to uniquely encode stereotypical perceptual patterns, as well as efficiently represent PPs in sensor data*. Objective **02** was completed through the use of the motor control strategy introduced in Chapter 2 for motion classification; the construction of perceptual libraries $\mathbf{M}^{sj}$ from motion capture data during validation, particularly when sharing PPs, accomplished Objectives **01**, **04** and **05**.

As general drawbacks, the proposed model is focused on periodic kinesthetic experiences; while the notion of stereotypical behaviors extends to the discrete case, the underlying formulation is built upon rhythmic DMP systems for simplicity. Moreover, despite no assumption being made regarding sensor modality, each input is treated as a unidimensional data stream; as noted in 3.3, a higher level of abstraction would be needed for use with multidimensional inputs e.g. vision data - a ubiquitous perception system in autonomous systems [Kotseruba & Tsotsos, 2018]. In addition, the validation procedure, despite achieving good results in gait classification and perturbation detection, was oversimplified for practical purposes; shared PPs across different walking modalities are not explored, despite the similarity between perceptual inputs.

Nevertheless, in addition to extensive real-time validation - including the bootstrapping mechanism in Algorithm 5 - with a larger and multimodal sensor set, both the proposed memory model and practical algorithms can be improved upon by 1) discarding the independence assumption in (68), and inferring inter-parameter co-relationships through their covariance; 2) accumulating sensor input sequentially in a moving time window, estimating an *experimental* distribution $\mathcal{N}_t^*$ describing transient sensor data, and employing e.g. entropy-based distance measures e.g. Kullback-Leibler divergence $D_{KL}(\mathcal{N}_t^* || \mathcal{M}_t^m)$ to define alternative similarity d^m metrics; 3) implementing dynamic *forgetting* mechanisms to remove irrelevant memories over time [Mohan et al., 2014], improving data efficiency; 3) exploring the use of Probabilistic Movement Primitives (ProMPs), Hidden Markov Models (HMMs) or alternative statistical models (e.g. SARIMA) to model stochasticity and cyclostationarity in sensor data [Guan et al., 2021]; and 4) updating the theoretical formalization with insights from modern DMP extensions e.g. scale invariance [Saveriano et al., 2021]. From a practical standpoint, it can be improved upon by 5) creating perturbation-specific memories that allow discriminating between different sources of disturbances, and can be used as *negative* examples informing high-level planning; 6) discarding null matches (i.e. perturbations) below a minimum sample or time threshold; 7) using large, heterogeneous perceptual libraries to benchmark the computational efficiency of Algorithm 3, 4 and 5; and 8) harnessing the time-invariance of perceptual memories, by

encoding motions at different time scales with the same memory. Conceptually, exploring 9) memory generalization through generic task- and platform-agnostic PPs, as well as establishing 10) higher-level associative perceptuomotor connections between $\mathcal{M}^m$ and $\mathcal{S}^s$ (Chapter 2, 2.3.2) present appealing lines of research in the short term, particularly focusing on contextual RL techniques for motor adaptation.

Impact of Muscle Fatigue in Musculoskeletal Primitives

In this chapter, a complementary study is presented where perceptual information is used as an objective indicator of environmental conditions such as *muscle fatigue* in human collaborators, in particular through primitive elements at a neuromuscular level *i.e.* Muscle Synergies (MSs). In line with the role of perception in artificial cognitive architectures (Chapter 3, 3.2.3), this represents a practical example of perceptual data - or extrapolated perceptual features - proving valuable information to high-level control strategies, with application in different fields of research in engineering, but particularly appealing in collaborative robotics, postural monitoring and ergonomic management. Conceptually, it not only explores a different category of primitive elements and their compositional role in human movement generation, but also how they evolve over time in order to meet dynamic task requirements and overcome environmental limitations, underlining the adaptability of biological behavior.

Objectives **01**, **05** and **06** are accomplished. For validation purposes, >1 synergy-based metrics exhibiting statistically significant variations is assumed to represent a Key Performance Indicator (KPI) of the study and attest to the applicability of the approach in real-world scenarios.

4.1 Introduction

Work-related Musculoskeletal Disorders (WRMSDs) represent a major issue for public health, reportedly representing 53% of all occupational diseases [Bevan, 2015], with estimated annual costs up to 240 B euros [Bevan, 2015; USBJI, 2015]. Muscle fatigue is one of the most relevant risk factors of WRMSDs, affecting roughly 35% of employees in Europe [de Kok et al., 2019], and with 22.8% of cases involving the upper body [Bevan, 2015]. Given these figures, it becomes of the utmost importance to timely identify and prevent the occurrence of muscle fatigue among the workforce, not only from a social and health perspective, but also with economic productivity and prosperity in mind [Bevan, 2015].

Muscle fatigue is commonly defined as the reduction of the force-generating capacity in a muscle, as a cumulative consequence of its sustained recruitment and activity [Liu et al., 2002; Sogaard et al., 2006].

From a neuromechanical standpoint a fatigued muscle manifests external signs such as: (i) slower contraction and relaxation, *i.e.* the conduction velocity of electrical signals decreases; and (ii) longer latency period between stimulation and beginning of contraction, *i.e.* muscle force is reduced and contraction occurs slower [Chang, 2018; Grandjean, 1979], which, particularly during the early stages, may not necessarily imply a reduction in task performance [Søgaard et al., 2006]. Popular methods for systematic fatigue identification rely on Electromyography (EMG) and objective metrics such as Root Mean Square (RMS), Median Frequency (MDF), and Mean Power Frequency (MNF) [Wang et al., 2018]. However, although these indicators work relatively well on isometric exercises, they are notably unreliable in dynamic and repetitive movements in submaximal contexts [Chang, 2018]. Moreover, this type of fatigue detection is typically based on the analysis of subjective patterns (*e.g.* RMS increase or MNF decrease) and not on homogeneous thresholds [Al-Mulla et al., 2011]. Hence, an objective and generalized method for accurate early identification is needed for widespread fatigue management in *e.g.* industrial settings. Recent promising methods have emerged based on the use of NNs, such as attention-based mechanisms [Li et al., 2020; Zhang et al., 2021], Gate Recurrent Unit (GRU) [Feng et al., 2022], and Long Short-term Memory (LSTM) [Jiang et al., 2022], achieving accuracy values between 83.5% and 98.42%.

Alternatively, a considerable amount of research shows that the Central Nervous System (CNS) drives its Musculoskeletal (MSK) counterpart through low-dimensional structures known as Muscle Synergies (MSs) [d’Avella et al., 2006]. While the focus is often in the study of muscle coordination, combining these synergistic properties with other relevant muscle information has provided interesting results in fatigue estimation [Li et al., 2020; Thomas et al., 2023]. Incidentally, over the last decade, relevant scientific work has been directed at understanding how muscular fatigue affects MSs *i.e.* the alterations to synergy structure and recruitment: Ortega-Auriol et al. [2018] studied upper limb movement under isometric conditions, while Hajiloo et al. [2020] focused on running as a case study; both concluded that the synergistic structure was preserved after fatigue onset, consistent with the notion that MSs represent elemental building blocks or neuromuscular primitives at the root of movement generation. Additionally, Hajiloo et al. [2020] found that synergy activation decreased with fatigue, hypothesized to be a strategy of the CNS to maintain optimal function of the motor system while ensuring task compliance. Thomas et al. [2023] focused on the effect of fatigue in shoulder muscles of baseball players, observing a decrease in the capability of MSs to describe the movement post-fatigue, as well as significant alterations in the contributions of individual muscles to the movement; this suggests a functional *restructuring* or *reorganization* of the MSK system in order to deal with a progressive inability of agonist muscles to compensate decreased muscular capacity. Smale et al. [2016], in turn, attempted to identify muscle fatigue during squatting motions: beyond validating wavelet-based algorithms as reliable fatigue indicators, obtained results revealed a significant tendency of increased activation of knee extensor muscles over general co-activation of leg muscles in the presence of fatigue. Notwithstanding all their valuable insights, it is relevant to note that none of these works proposes objective and efficient metrics to identify fatigue in the context of real-world industrial settings.

This chapter presents a study on the synergistic structure of the MSK system and how it is affected by muscle fatigue, focusing in repetitive low-force movements of the upper limbs, common in industrial sce-

narios, and reported to play a considerable role in WRMSDs. It attempts to identify practical approaches for early fatigue detection - herein understood as fatigue occurring before any noticeable kinematic change and/or task performance is compromised. In general, under the assumption that, in accordance with previous findings [Hajiloo et al., 2020; Ortega-Auriol et al., 2018; Thomas et al., 2023], the synergistic structure and internal organization of the MSK system is invariant to the occurrence of fatigue, this study attempts to identify a concrete answer to the following Research Question (RQ):

RQ3: What cumulative effect do low-force repetitive movements have in the activation profile of muscle synergies?

Towards an objective answer to this RQ, this chapter presents a study on such low-force repetitive motions in the upper limbs of human participants, and how the synergistic structure of the MSK system evolves as muscle fatigue gradually develops. Using detailed kinematic and EMG recordings, an extensive analysis is conducted, including classical and synergy-oriented methodologies, wherein selected time- and frequency-domain metrics are evaluated regarding their capacity to identify fatigue onset. To the best of the author's knowledge, this study presents the first attempt at identifying fatigue from muscle synergy activation on repetitive low-force movements, notably commonplace in industrial settings.

The chapter is structured as follows: in section 4.2 a brief overview is provided on the underlying theoretical foundations behind muscular behavior in fatigued and non-fatigued conditions, as well as popular methods and practical challenges (*i.e.* hardware instrumentation and signal processing) in muscular behavior monitoring, in order to contextualize the research problem at hand. Section 4.3 describes the experimental protocol used to construct an experimental dataset of kinematic and electromyographic data across multiple healthy subjects, and developed processing pipeline, in the form of EMG factorization, and subsequent conventional EMG and synergy analysis, whose results are presented in section 4.4. Section 4.5 carefully examines the results achieved and attempts to identify the rationale to explained the experimental observations, in the light of the current body of literature and research. Finally, section 4.6 attempts to extrapolate general conclusions regarding early fatigue detection and provide rough experimental guidelines for practical implementation.

4.2 Background

4.2.1 Electromyography

Muscular behavior is commonly characterized through the electric activity resulting from muscular excitation patterns, measured with Electromyography (EMG), which quantifies the *action potential* of the different *motor units* composing a muscle [Márquez-Figueroa et al., 2020; Nandedkar & Sanders, 1990]. Amplitude, frequency and phase analysis are usually employed in EMG processing [Márquez-Figueroa et al., 2020], and have been shown to be useful in a variety of contexts, from medical diagnosis of neuro- and myopathies [Nandedkar & Sanders, 1990; Reaz et al., 2006] to control in collaborative scenarios such

as assistive robotics [Doshi & Suthar, 2016; Jang et al., 2016; Yun et al., 2017] or prosthetics [Barzilay & Wolf, 2011; Bitzer & van der Smagt, 2006; Cimolatto et al., 2022; Peña et al., 2019].

On the other side of this appeal for research and therapeutic applications, however, EMG is fraught with practical challenges that add complexity to its analysis. In the specific case of Surface Electromyography (sEMG) - a popular modality due to sheer practicality - action potentials traveling through the biological tissues up to the skin-electrode interface are attenuated and filtered (with particular incidence to the higher-frequency components). This requires careful electrode placement and extensive data processing in order to achieve a manageable signal-to-noise ratio [Fuglevand et al., 1992; Nandedkar & Sanders, 1990]. In fact, due to the morphology and physiology of muscular tissues - multiple anisotropic and inhomogeneous fibers - the signal recorded by sEMG systems is technically an *interference pattern* influenced by electrical, chemical and mechanical processes [D'Alessio & Conforto, 2001; Márquez-Figueroa et al., 2020], representing the combined activity of multiple motor units within the muscle [Sörnmo & Laguna, 2005]. and further exacerbated by the possible interference of sensor noise, electrophysiological processes such as cardiac activity, and movement artifacts, that significantly degrade the quality of the measurement, particularly in sEMG [Márquez-Figueroa et al., 2020].

4.2.2 Physiological and Biomechanical Effects of Muscular Fatigue

Muscular fatigue is known to lead to disruptions in the activation and contraction dynamics [Liu et al., 2002; Sogaard et al., 2006]. As a consequence, motor neurons start receiving higher corticospinal input which has the effect of a) inducing unwanted coactivations and b) decreasing force steadiness and inducing tremors [Weavil & Amann, 2018]. While this deterioration of muscular behavior is associated with several physiological phenomena that have been extensively documented in the literature, it can be objectively identified through a measurable drop in peak muscle force and in *time-to-failure* values, as well as an increase in motor variability and changes in motor coordination. Studies have reported that fatigue leads to excessive activation of the neighbor muscles [Matsunaga et al., 2021]. Conversely, not only is there evidence to suggest that the internal intermuscular organization is preserved [Hajiloo et al., 2020; Ortega-Auriol et al., 2018; Thomas et al., 2023], but in redundant systems, goal-relevant aspects of the movement remain mostly unaffected by fatigue [Singh & Latash, 2011]. This suggests the MSK system is able to exploit its own redundancy in an effort to minimize detrimental effects of muscular fatigue. Additionally, previous research has suggested a correlation between the force capability of individual *motor units* and their inherent fatigability [Fuglevand et al., 1999; Potvin & Fuglevand, 2017].

From a practical standpoint, the drop in peak force can be objectively quantified through isometric Maximum Voluntary Contraction (MVC) measurements before and after a fatigue-inducing event or period, as reviewed by Chang [2018]. However, conducting MVC trials continuously or concurrently with demanding tasks is not only impractical, but is also likely to induce additional fatigue [Chang, 2018]. Alternatively, a conventional method to quantitatively assess the fatigue level of an individual is regular self-assessment [Öberg et al., 1994] resorting to techniques such as Borg's Perceived Exertion Scale [Borg, 1998] or the Swedish Occupational Fatigue Inventory [Åhsberg et al., 1997]. On the other hand, this

Table 13: Time- and frequency-domain metrics commonly extracted from EMG [Chang, 2018; Najarian & Splinter, 2012; Phinyomark et al., 2012]; e_i denotes the raw EMG signal, E_t the extracted envelope curve, $S_{ee}[f]$ the estimated Power Spectral Density (PSD) of e_i , f_s the sampling frequency, and T the number of samples in e_i .

	Metric	Definition
Time	RMS	Root Mean Square $\sqrt{\frac{1}{T} \sum_{t=1}^T E_t^2}$
	MAV	Mean Absolute Value $\frac{1}{T} \sum_{t=1}^T E_t $
	Int	Integral $\approx \sum_{t=1}^T E_t$
Frequency	MNF	Mean Power Frequency $\frac{\sum_{f=1}^M f S_{ee}[f]}{\sum_{w=1}^M S_{ee}[f]}$
	MDF	Median Frequency $S_{ee}^{-1}(\frac{1}{2} \text{TTP})$
	PkF	Peak Frequency $\arg \max_f (S_{ee}[f])$
	TTP	Total Power $\approx \sum_{j=1}^{f_s/2} S_{ee}[j]$
	FR	Frequency Ratio $\frac{\sum_{j=1}^{f_s/4} S_{ee}[j]}{\sum_{j=f_s/4}^{f_s/2} S_{ee}[j]}$

approach is often challenging to automate due to its subjective nature and lacks the required detail for comprehensive studies on muscular fatigue [Chang, 2018].

Regarding the electrical profile of muscle fatigue - as measured with EMG - perhaps the most well documented outcome is a *spectral compression* i.e. decrease in dominant frequencies, coupled with an increase in EMG amplitude [Chang, 2018; Gizzi et al., 2021; Poyil et al., 2020].

4.2.3 Fatigue Indicators in Classical EMG Analysis

Conventional EMG analysis usually focuses on amplitude, frequency and phase analysis of the recorded signal [Márquez-Figueroa et al., 2020], although alternative approaches e.g. wavelet transforms have also been used [Najarian & Splinter, 2012]. In amplitude analysis, a popular feature is the EMG *linear envelope*, which refers to the low-frequency curve describing local average with time [Nandedkar & Sanders, 1990; Shiavi, 1990], and has been directly associated with muscular force (in the absence of fatigue) in both isometric and dynamic conditions [D'Alessio & Conforto, 2001]¹.

In practical terms, EMG analysis can help detect the occurrence of muscular fatigue, assessed through the computation of several time- and frequency-domain metrics that characterize the recorded signal (Table 13). Amplitude and frequency analysis become notably relevant in order to identify both the increase in scale and the shift in dominant frequencies characteristic of fatigued behavior; RMS [Chang, 2018; Najarian & Splinter, 2012] and MDF [Chang, 2018], respectively, are commonly used to that end. On the

¹While the relationship between EMG amplitude and muscle force has been theorized to be linear, the anisotropy and inhomogeneity of the underlying tissues impede a linear relationship when using sEMG; moreover, under dynamic conditions the moment arm and rotation center are not constant, and cross-cancellation of EMG amplitude between different motor units increase with force, add complexity to amplitude analysis [D'Alessio & Conforto, 2001; De Luca, 1997; Farina et al., 2008].

other hand, amplitude analysis is generally not reliable in most dynamic movements due to, as previously mentioned, being dependent on force. Moreover, alterations to the power spectrum are often reported as statistically insignificant to draw meaningful conclusions, particularly in low-force tasks and during early stages of fatigue [Chang, 2018; Smale et al., 2016]. Alternatively, Chang [2018] proposes a simple fatigue indicator based on the difference between RMS before and after fatigue onset; however, this is only valid in isotonic and isometric conditions.

4.2.4 Muscle Synergies as Neuromuscular Primitives

The notion of Muscle Synergy (MS) refers to coordinated muscle behavior as a result of shared neural activation patterns across different muscles [d’Avella et al., 2006]. In practical terms, a limited set of MSs can describe the behavior of larger pools of musculotendon actuators (Figure 48), which has been theorized to be a strategy through which the CNS tackles the inherent high-dimensionality of the biological motor system. This phenomenon has been observed in different organisms, and effectively exploited in the fields of neurophysiology [Reaz et al., 2006] and motor control [Rückert et al., 2013] as a dimensionality reduction strategy, in different spatial, temporal and spatiotemporal organizations² [d’Avella, 2016]. Empirical evidence that MSs are not only reused among different instances of a motor task e.g. different gait modalities [Barroso, 2016; Oliveira et al., 2016], but also between different tasks [d’Avella & Bizzi, 2005], hints at an underlying hierarchical and modular MSK with appealing properties and applicability in a variety of contexts [d’Avella, 2016]. A conceptual and functional similarity between MSs and kinematic MPs (Chapter 2) is evident; in fact, given the lack of scientific evidence of kinematic MPs in higher-order cortical centers, they are often viewed as emerging from the interaction between MSs and the MSK system³, rather than separate, independent primitive elements [Singh et al., 2018].

From a mathematical perspective, identifying the synergistic organization behind muscular behavior entails a factorization process that isolates lower-rank components **W** (*feature* matrix with muscle-specific weights) and **H** (*coefficient* matrix i.e. temporal activation profile for each synergy) from the input EMG. Numerical methods such as Non-negative Matrix Factorization (NMF) [Lee & Seung, 1999], exploiting the inherent non-negativity of normalized muscular activation, are commonly used to find a low-dimensional representation of recorded EMG patterns; following this methodology, an input matrix **E**, aggregating muscle-specific recordings, is approximated as the multiplication of two lower rank matrices **W** and **H**:

$$\mathbf{E}_{M \times T} \stackrel{\text{NMF}}{\approx} \mathbf{W}_{M \times N} \times \mathbf{H}_{N \times T} = \sum_{n=1}^N \mathbf{w}^n \times \mathbf{h}^n \quad (72)$$

where **E** $\in \mathbb{R}^{M \times T}$ contains M synchronous time series of length T ; **W** $\in \mathbb{R}^{M \times N}$ aggregates N underlying features i.e. synergies $\mathbf{w}^n = [w_1^n, w_2^n, \dots, w_M^n]^\top$, $\forall n \in \{1, 2, \dots, N\}$, as sets of uni-dimensional muscle-specific weights w_m^n ; and **H** $\in \mathbb{R}^{N \times T}$ encodes time-indexed *activation* vectors $\mathbf{h}^n = [h_1^n, h_2^n, \dots, h_T^n]^\top$ scaling each synergy at each t^{th} sample, $\forall t \in \{1, 2, \dots, T\}$:

²In this work, the spatiotemporal synergy model is assumed i.e. with synergy-specific activation patterns [d’Avella, 2016].

³Under this bottom-up view, MSs are often categorized as neuromuscular MPs [Singh et al., 2018].




Figure 48: Conceptual illustration of the underlying principle behind Muscle Synergies, wherein EMG envelope $\mathbf{E}_m$ is factorized into synergies $[\mathbf{w}^1, \mathbf{w}^2, \mathbf{w}^3]$, recruited according to $[\mathbf{h}^1, \mathbf{h}^2, \mathbf{h}^3]$, often representing functional kinematic primitives e.g. elbow flexion, elbow extension and shoulder extension.

$$\begin{bmatrix} \mathbf{E}_1 \\ \mathbf{E}_2 \\ \vdots \\ \mathbf{E}_M \end{bmatrix} = \begin{bmatrix} E_{1,1} & E_{1,2} & \dots & E_{1,T} \\ E_{2,1} & E_{2,2} & \dots & E_{2,T} \\ \vdots & \vdots & \ddots & \vdots \\ E_{M,1} & E_{M,2} & \dots & E_{M,T} \end{bmatrix} \stackrel{\text{NMF}}{\approx} \begin{bmatrix} w_1^1 & \dots & w_1^N \\ w_2^1 & \dots & w_2^N \\ \vdots & \ddots & \vdots \\ w_M^1 & \dots & w_M^N \end{bmatrix} \times \begin{bmatrix} h_1^1 & \dots & h_T^1 \\ h_1^2 & \dots & h_T^2 \\ \vdots & \ddots & \vdots \\ h_1^N & \dots & h_T^N \end{bmatrix} \quad (73)$$

As long as $N < M$, $\mathbf{H}$ constitutes a lower-order representation of the original recorded EMG; this is particularly valuable when the underlying structure behind muscular activity represented in $\mathbf{W}$ remains constant over time, or across different motions and contexts. Contrasting with the MP representation employed in Chapter 2 and 3, spatiotemporal MSs encode discrete muscle weighting vectors *i.e.* $\mathbf{w}^n \in \mathbb{R}^M$, and not time-based trajectories, and can - conceptually - be viewed as separate muscular units.

As previously discussed (4.2.2), muscle fatigue is known to cause changes in the internal coordination and organization of the MSK system, while limb and joint kinematics are, generally, less affected, particularly during early stages of fatigue [Singh & Latash, 2011]. On the other hand, there is scientific evidence demonstrating that the synergistic structure is preserved even in fatigued conditions [Hajiloo et al., 2020; Ortega-Auriol et al., 2018; Thomas et al., 2023]. Under this assumption, and viewing the occurrence of muscular fatigue as a gradual, cumulative process [O'Sullivan & Gallwey, 2002; Potvin & Fuglevand, 2017], slight and nuanced alterations in MS recruitment patterns $\mathbf{H}$ (72 - 73) are expected

during its initial stages. While studies attempting to link MSs to the early identification of muscle fatigue have been conducted [Hajiloo et al., 2020; Li et al., 2020; Ortega-Auriol et al., 2018; Thomas et al., 2023], there is still no reliable generalized metric that can be used in practical, real-world applications to detect its effects before they impacts task performance, or the well-being and ergonomic profile of the individual. From a practical perspective, MSs allow restricting any detection algorithm to a smaller set of functional units instead of the full array of muscles recruited during a specific task.

4.3 Methodology

4.3.1 Participants

Trial participants consisted of healthy volunteers from the academic community of the University of Minho. The purpose, specifics, methodology, and duration of the study were disclosed to them. The following inclusion criteria were established in order to choose and recruit the participants: i) age equal or greater than 18 years old; ii) no motor impairments and no clinical history of muscle strain, lesions or trauma on the upper body; iii) no strenuous activity performed over the previous 24 h.

A total of 17 participants (11 Male; 6 Female; 23.82 ± 2.79 years old, 69.06 ± 14.75 Kg) were recruited (Table 14). Each participant willingly agreed to participate in the study and gave their signed, informed consent, according to the ethical conduct defined by the University of Minho Ethics Committee in Life and Health Sciences (*CEICVS 006/2020*), that follows the standard set by the declaration of Helsinki and the Oviedo Convention. All participants were right-hand dominant.

4.3.2 Instrumentation

To avoid any constraint to upper limb movement and facilitate instrumentation, each participant was asked to wear upper body loose clothing (e.g. shirts or strap tops which allowed easy access to upper arm and chest muscles), as well as sneakers and tight-fitting lower body clothing (e.g. jeans or leggings).

Firstly, sensor locations were prepared by shaving excess hair and thoroughly cleaning with an alcohol pad. A total of 8 sEMG wireless sensors (*Trigno Avanti*[®], Delsys Incorporated, Boston, MA) were attached to relevant right arm muscles (Figure 49, **Left**), namely: Clavicular Head of the Pectoralis Major (PECM), Deltoid Anterior (DELTA), Deltoid Medial (DELTM), Deltoid Posterior (DELTP), Lateral Head of Triceps Brachii (TRILAT), Long Head of Triceps Brachii (TRILON), Biceps Brachii (BIC), and Brachioradialis (BRR); each electrode was placed on the muscle belly, longitudinally aligned with fiber orientation, following the instructions from scientific guidelines i.e. SENIAM [Stegeman & Hermens, 2007] and relevant literature [Lee, 2019; O’Sullivan & Gallwey, 2002]; to grant a secure hold to the sensors and avoid motion or displacement artifacts, tape (*Tesa Pack*[®]) was placed over all sensors (Figure 49, **Left**). Correct placement of the sensors was validated by performing short MVC trials for each muscle while ensuring the conformity of EMG measurement. All sEMG recordings were sampled at 2048 Hz (16 bits p/sample), employing the *EMGWorks*[®] software suite (v. 4.7).



Figure 49: Participant instrumented with the sEMG sensors (*Trigno Avanti*®) on: 1) PECM, 2) DELTA, 3) DELTM, 4) DELTP, 5) TRILON, 6) TRILAT, 7) BIC, and 8) BRR (**Left**); and the inertial motion capture system (*MVN Awinda*®) in addition to the sEMG sensors (**Center, Right**).

Table 14: Overview of the participant pool recruited for the study, including weekly exercising habits, sleep hours and stimulant consumption on the date of the trials.

Participant	Gender	Age	Weight (kg)	Height (m)	Sleep (h)	Weekly exercise	Stimulants
S1	Male	33	68	1.74	6	-	2 <i>coffee</i>
S2	Female	27	63	1.70	7	1x	1 <i>coffee</i>
S3	Female	23	48	1.54	7	3x	1 <i>coffee</i>
S4	Male	22	65	1.88	7	-	-
S5	Male	23	86	1.84	7	1x <i>futsal</i>	2 <i>coffee</i>
S6	Female	23	60	1.62	7	4x	1 <i>coffee</i>
S7	Male	25	75	1.77	8	5-6x <i>jiujitsu</i>	1 <i>coffee</i>
S8	Male	22	75	1.76	8	1x <i>running</i>	-
S9	Male	22	68	1.80	8	2x	-
S10	Male	22	75	1.77	6	3x	-
S11	Female	22	55	1.72	8	1-2x	-
S12	Male	23	64	1.78	5	-	1 <i>coffee</i>
S13	Male	23	94	1.97	7	3-4x	1 <i>coffee</i>
S14	Male	24	101	1.84	8	2-3x	2 <i>coffee</i>
S15	Female	22	52	1.64	7	3x	2 <i>coffee</i>
S16	Male	22	78	1.80	7.5	4x <i>football</i>	-
S17	Female	27	47	1.55	7	3x	-

After sEMG instrumentation, participants were equipped with a full-body inertial wireless motion capture system (MVN Awinda[®], Xsens Technologies, B.V., The Netherlands). A total of 17 inertial trackers were placed following the manufacturer's instructions *i.e.* on the head, shoulders, chest, arms, forearms, wrist, waist, thighs, shanks, and feet, and secured with velcro straps (Figure 49, **Center** and **Right**). During the execution of the experimental trials, kinematic data was acquired through the MVN Analyze[®] software suite (v. 2021.2) and sampled at 60Hz. Finally, synchronization between EMG and kinematic data was achieved through the hardware trigger mechanisms of the *Trigno* and *Awinda* systems.

4.3.3 Experimental Protocol

Prior to the experiment, participants were asked to fill a preliminary questionnaire to collect demographic data, exercising habits, and consumption of CNS stimulants (Table 14). As conceptualized in Figure 50, trials consisted of four different exercises grouped into two distinct parts; electromyographic and kinematic data was recorded during each exercise and processed for posterior analysis.

During the first part of the trial (**P1**), participants executed, in a fixed order, 3 uni-articular movements of the right elbow and shoulder joints, representing non-fatigued *generic* upper limb motions (Figure 51):

A) elbow counter-gravity *flexion* and gravity-assisted *extension*, with torso tilted forwards;

B) shoulder counter-gravity *abduction* and gravity-assisted *adduction*;

C) shoulder counter-gravity *flexion* and gravity-assisted *extension*;

which are repeated 15 times while holding a 1.5 kg dumbbell, interleaved with rest periods of 3 minutes, in order to prevent muscular fatigue between repetitions. After **P1**, an additional rest period of 10 minutes took place, in order to allow further recovery from any potential cumulative effect.

In the second part of the trial (**P2**), with the same 1.5 kg load, participants were requested to perform a composite sequence of the uni-articular **P1** motions **A**, **B** and **C** (Figure 52), until self-reported exertion (*i.e.* when task compliance was no longer possible) or noticeable changes to the kinematic profile of the movement was detected by the research team.




Figure 50: Conceptual flowchart illustrating the two separate parts of the experimental protocol: **P1**, comprising 3 uni-articular exercises (**A**, **B**, **C**); and **P2**, where a composite motion is repeated until exertion (**D**). Generic synergies W_{ABC} are extracted from **P1**; In **P2**, classical and synergy analysis (4.3.4.5) is conducted on each individual repetition of **D**.




Figure 51: **P1** exercises **A**, **B** and **C** (Top) and **P2** exercise **D** (Bottom).




Figure 52: Sequence of movements composing exercise **D**, as depicted in Figure 51 (Bottom), including uni-articular motions from **A**, **B** and **C**; shoulder extension (**C**) is only performed after the first repetition.

The goal of **P2** is twofold: on one hand, to 1) progressively induce fatigue in the upper limb muscles, allowing a step-by-step analysis of the evolution of muscular recruitment from a rested to fatigued state; and to 2) demonstrate that complex movements can be represented by a combination of simpler MSs by integrating generic motions **A**, **B** and **C** into the sequence.

Special care was taken to ensure that the sequence was easy to follow, in order to obtain similar kinematic patterns among participants and repetitions. The research team demonstrated the correct procedure before each exercise, and instructed subjects to remain focused throughout the duration of the trial, and participants were allowed to practice, without any load, until confident enough to reproduce the target movements without error. A metronome, hand-tuned to 40 bpm (**A**, **B** and **C**) and 60 bpm (**D**), provided a reference to guide execution, and facilitated familiarization with the protocol.

For reference purposes, at the end of each exercise subjects were inquired about their perceived fatigue level according to Borg's Perceived Exertion Scale [Borg, 1998]; any time a fatigue level above 4 was reported, rest time would extend by 5 minutes. Lastly, a preliminary correlation analysis was conducted on the recorded data, with the goal of evaluating the degree of uniformity between different subjects, understood to reflect the consistent execution of the experimental protocol and, potentially, identify outliers.

4.3.4 Data Processing

4.3.4.1 EMG Envelope Extraction

Following the guidelines found in relevant literature [Barzilay & Wolf, 2011; D'Alessio & Conforto, 2001; Márquez-Figueroa et al., 2020; Oliveira et al., 2014; Reaz et al., 2006; Stegeman & Hermens, 2007], a 1) 4th-order Butterworth band-pass filter (20 - 450 Hz) was applied to recorded EMG trajectories $\mathbf{e}_m$ in order to remove external disturbances (e.g. low-frequency sensor drift as well as high-frequency interference *i.e.* movement artifacts or electrical noise) not associated with the myoelectric activity. The envelope $\mathbf{E}_m$ of the filtered EMG was then obtained through a combination of 2) full-wave rectification followed by 3) a secondary 4th-order Butterworth low-pass filter (<1 Hz); the former shifted the signal into the positive range with the latter enhancing the underlying low-frequency component. This approach, conceptualized in Figure 53, is in line with common envelope extraction methods that generally implement some form of signal smoothing, implemented through low-pass filtering and/or integration steps [Barzilay & Wolf, 2011; Olstad, 2015; Zhang et al., 1991], and was found to provide satisfactory results with the recorded data.

The raw experimental dataset is publicly available [André et al., 2024].

4.3.4.2 Kinematic Segmentation

Despite the stereotypical nature of the experimental protocol, inter- and intra-participant variability between repetitions of the same exercise (e.g. range and duration) is expected to some extent. For a consistent and uniform data analysis, each extracted envelope $\mathbf{E}_m$ was divided into I segments, using recorded kinematic data $\mathbf{q}$ as reference. In **P1**, a total of $I = 10$ consecutive segments were selected from the 15 repetitions executed (4.3.3) while in **P2**, I denotes the number of repetitions each subject




Figure 53: Conceptual flowchart describing the extraction of linear envelope $\mathbf{E}_m$ for the m^{th} muscle under study, from raw sEMG $\mathbf{e}_m$ recorded at $f_s = 2048$ Hz; in **P2** (—), classical EMG analysis is conducted on filtered $\mathbf{e}_m$ within a fixed window Δt clipped around the global maximum of $\mathbf{E}_m$.




Figure 54: Conceptual flowchart of the kinematic segmentation procedure; $\mathbf{q}^{\text{REF}}$ denotes an exercise-specific reference kinematic DoF, used to estimate I sets of L time values $\hat{t}_l^i$, relating to the l^{th} sample of the i^{th} repetition executed during the trial, $l \in \{1, 2, \dots, L\}$ and $i \in \{1, 2, \dots, I\}$.

was able to perform before exertion.

Similarly to the procedure in Chapter 3 (3.4.3.3), as outlined in Figure 54, recorded kinematics $\mathbf{q}^{\text{REF}}$ of exercise-specific reference joints⁴ (4.3.3) were 1) *high-pass* filtered (2nd order, Butterworth, >0.1 Hz) in order to remove sensor drift and other low-frequency components; each sample in the resulting filtered curve $\mathbf{q}^{\text{HP}}$ was then used as reference in a 2) *thresholding* step discarding values in the original $\mathbf{q}^{\text{REF}}$:

$$q_t^+ = \begin{cases} q_t^{\text{REF}} & \text{if } q_t^{\text{HP}} > 0 \\ 0 & \text{otherwise} \end{cases} \quad (74)$$

The resulting signal $\mathbf{q}^+$ was then 3) *differentiated*, highlighting zero-crossings in $\mathbf{q}^{\text{HP}}$ as very steep gradients, defining positive and negative peaks, which were located using a 4) *peak detector* algorithm that accurately pinpoints local maxima on the data curve. The resulting array of I time stamps - associated with the edges between consecutive repetitions - was linearly sliced into L equidistant target values:

$$\hat{t}_l^i = t_i + \frac{l}{L} (t_{i+1} - t_i), \quad \forall l \in \{1, 2, \dots, L\} \quad (75)$$

which were used to interpolate $\mathbf{E}_m$ into I fixed-width sets of L samples i.e. $\mathbf{E}_m^i$, as will be described in 4.3.4.4. In addition to temporal invariance between repetitions, this approach has the advantage of balancing the contribution of each segment in subsequent processing and analysis procedures. In all exercises $L = 1000$. The result of the segmentation procedure is illustrated in Figure 55.

⁴For reference, elbow (**A**) and shoulder (**B** and **C**) joints; in all cases, the z axis is used.




Figure 55: Illustrative example of the kinematic segmentation procedure, for exercise **C** with **S8**.

4.3.4.3 Normalization

EMG is typically normalized by linearly scaling the values against a reference value that represents the *amplitude at maximal contraction*, commonly the peak value recording during a MVC trial. However, in the context of this work this comprises a practical challenge due to 1) the difficulty in finding a movement that uniformly recruits all muscles under analysis at maximal contraction, 2) the high likelihood of the MVC trial inducing unwanted muscle fatigue and finally 3) a gradual decrease on the maximal capacity of the muscle assumed during **D**; as such the *peak dynamic method* of normalization [Bitzer & van der Smagt, 2006; Burden, 2010; Sousa & Tavares, 2012] is adopted instead, providing a relative measure of muscle recruitment during each exercise. To that end, for each muscle m , all segments $\mathbf{E}_m^i$ were normalized using the maximum value across all I repetitions:

$$E_{m,l}^i = \frac{E_{m,l}^i}{E_m^{\text{REF}}}, \quad E_m^{\text{REF}} = \max(\mathbf{E}_m) \quad \forall k \in \{0, 1, \dots, K\} \quad (76)$$

This approach allowed direct comparison between repetitions of the same exercise - particularly relevant during **D** - and simplifies the experimental protocol by discarding the need for repeated MVC trials during the experiment.

4.3.4.4 Muscle Synergy Estimation

As presented in 4.2.4, the estimation of MSs generally relies on the factorization of the EMG envelope $\mathbf{E}_m$ into lower rank components $\mathbf{W}$ and $\mathbf{H}$. As was illustrated in Figure 50, herein this was carried out on the segmented envelopes $\mathbf{E}_m^i$ in two distinct phases: in **P1**, 1) using concatenated data from **A**, **B** and **C**,




Figure 56: Cascade of processes applied to the EMG envelope in order to extract the synergistic organization of underlying muscular behavior with and without fatigue; in **P2** (–), each repetition is processed individually (bypassing concatenation) and NMF is constrained to the global synergies $\mathbf{W}_{ABC}$.

with the purpose of identifying generic synergies $\mathbf{W}_{ABC}$ relating to uni-articular motions of the upper limbs; and in **P2**, 2) using individual repetitions of $\mathbf{D}$, while enforcing $\mathbf{W}_D = \mathbf{W}_{ABC}$, in order to detect transient or progressive alterations in *how* each generic MS was employed by the MSK system.

Figure 56 describes the procedure adopted in both cases: in **P1**, the complete set of $I = 10$ segments were 1) *interpolated* from reference time stamps $\hat{t}_l^i$ (4.3.4.2), 2) serially *concatenated*, and 3) *normalized* (4.3.4.3), leading to a reconstructed sequence totaling $3 \times I \times L = 30000$ samples per muscle, for each participant. Notably, despite adding non-negligible computational weight to factorization⁵, data aggregation through concatenation was preferred over individual sample averaging, as the first ensured that each sample was equally taken into account, while the latter may have attenuated variations between different repetitions. The resulting concatenated sequences were used as input to a 4) *factorization* step, employing NMF⁶ for the extraction of $n \in \{1, \dots, M\}$ MSs, with $M = 8$ i.e. number of muscles under study. For each value of n , measured Variance Accounted For (VAF) was used to quantify the quality of the factorization process, and ergo the quality of the reconstructed signal [Barroso, 2016; Saito et al., 2018; Yaserifar et al., 2021]; using a fixed threshold of $VAF \geq 0.9$, the *minimum* number of synergies N_{ABC} required to accurately describe **P1** data was identified, and fixed during **P2**, across all participants. Moreover, for all values of n , in order to mitigate the stochastic and non-convex nature of NMF, $\mathbf{W}_{ABC}$ obtained for **S1** was used as an initial guess in the factorization step in the remaining participants; discarding the need for functional sorting of the resulting synergies without constraining the optimization process⁷ [Shourijeh et al., 2016; Smale et al., 2016]; this not only decreased convergence time, but also lead to consistent MS indexing among different participants, relevant in subsequent analysis.

Conversely, in **P2**, the concatenation step (Figure 56) was bypassed and each sequential repetition of $\mathbf{D}$ factorized individually while constrained to the generic synergies $\mathbf{W}_{ABC}$ identified in **P1**; this effectively found $\mathbf{H}_D^i$ i.e. the set of N_{ABC} activation profiles $\mathbf{h}_D^{n,i}$ that best matched input $\mathbf{E}_m^i$. As discussed, this approach was grounded on the assumption that $\mathbf{W}_{ABC}$ remain unaffected after fatigue onset [Hajiloo et al., 2020; Ortega-Auriol et al., 2018; Thomas et al., 2023], with any alterations in EMG being reflected in $\mathbf{H}_D^i$.

⁵As a result of a larger number of samples as input.

⁶`scikit-learn` implementation of NMF is used.

⁷Functional sorting between MSs is often not straightforward; strategies e.g. *Concatenated* NMF fix $\mathbf{W}$ across participants.

Table 15: Metrics comprising the analysis of the synergy activation $\mathbf{H}_D^i$ in the i^{th} repetition, during **P2**; $h_l^{n,i}$ represents the activation of the n^{th} synergy at the l^{th} sample and $S_{hh}^i[f]$ the estimated PSD of $\mathbf{h}_D^{n,i}$.

	Metric	Definition
Time	Energy	Activation Energy $\approx \sum_{l=1}^L h_l^{n,i} ^2$
	RMS	Root Mean Square $\sqrt{\frac{1}{L} \sum_{l=1}^L (h_l^{n,i})^2}$
	MAV	Mean Absolute Value $\frac{1}{L} \sum_{l=1}^L h_l^{n,i} $
	iPk	Index/Position of Peak $\arg \max_l (h_l^{n,i})$
	P2P	Peak-to-Peak Amplitude $\max(h_l^{n,i}) - \min(h_l^{n,i})$
	WaHP	Width at Half Peak $ l_{\text{pos}} - l_{\text{pre}} , \{l_{\text{pre}}, l_{\text{pos}}\} \approx [h^{-1}]_l^{n,i} [\frac{1}{2} \max(h_l^{n,i})]$
	Int	Integral $\approx \sum_{l=1}^L h_l^{n,i}$
$\mathbf{E}_m^i$	TTP	Total Power $\sum_{j=1}^{f_s/2} S_{hh}^i[j]$

4.3.4.5 Classical and Synergy Analysis

In **P2**, as illustrated in Figure 50, each repetition of **D** was processed individually. On one hand, conventional indicators (Table 13) were computed from a fixed window of $\Delta t = 0.250$ s centered around the recruitment peaks of $\mathbf{e}_m$, identified from the maximum within each segment $\mathbf{E}_m^i$ (Figure 53), in order to discard relaxed intervals and restrict analysis to contracted state in all muscles [Bamman et al., 1997; Bolgla & Uhl, 2005; Burden, 2010; Farina et al., 2004]; in addition, beyond a qualitative comparison of the I activation profiles $\mathbf{H}_D^i$, multiple time- and frequency-domain metrics (Table 15) were employed to objectively scrutinize alterations in how muscle synergies are recruited throughout **P2** (Figure 56), and, potentially, identify gradual meaningful trends over time. These metrics were computed for each individual activation curve $\mathbf{h}_D^{n,i}$, and were inspired by conventional EMG analysis (4.2.3, Table 13), with the occurrence of muscle fatigue hypothesized to be reflected on the evolution of these metrics from a non-fatigued state ($i = 1$) until the end of the exercise ($i = I$). It is relevant to note that, as $\mathbf{H}_D^k$ generally constitutes low frequency curves⁸, frequency-domain metrics other than **Total Power (TTP)** provide no useful information.

4.3.4.6 Statistical Significance Analysis

Lastly, a straightforward analysis is conducted on the evolution of metric values, aimed at understanding which 1) exhibit *statistically significant* variations during **P2**, and - when meaningful - 2) by *how much* they change and 3) which vary *linearly* with time and cumulative effort.

In both classical and synergy analysis, results amount to ordered time series encoding metric values in each repetition of **D**. While the presence of electrical noise is expected, gradual or progressive alterations in muscular recruitment suggest the existence of underlying *monotonic* trends. As such, time- and frequency-

⁸As a consequence of factorizing the EMG envelope $\mathbf{E}_m^i$.

domain metrics that register significant variations can be identified as *potential fatigue indicators*. To that end, two different methods are employed for each subject and dimension (i.e. sensor and synergy in classical and synergy analysis, respectively): the 1) Mann-Kendall (MK) test [Gilbert, 1987; Kendall, 1970] for the detection of (linear or non-linear) monotonic trends; and 2) linear regression analysis in order to assess the linearity of meaningful trends.

Globally, to assert if detected trends occur consistently across different participants, muscles and synergies, as well as discard potential trial- or participant-specific outliers, a *one-sample t-test* over the normalized MK test statistic $\bar{Z}$ ⁹ is conducted:

$$t = \frac{\bar{X} - \mu}{\hat{\sigma}/\sqrt{N}} \rightarrow t_Z = \frac{\bar{Z}}{\hat{\sigma}_Z/\sqrt{D}} \quad (77)$$

where $\bar{Z}$ and $\hat{\sigma}$ represent, respectively, the mean and the *unbiased* standard deviation in the values of Z resulting from each MK test; and D the sample size i.e. $D = 17$ and $D = N_{\mathbf{ABC}}$ for muscle- and synergy-specific analysis, respectively; $D = 17 \times M = 136$ for global muscle-specific analysis; and, similarly, $D = 17 \times N_{\mathbf{ABC}}$ in global synergy-specific analysis. In all cases, H_0 entails that Z values are noise around $\mu = 0$ i.e. there is no meaningful trend consistent across trials. Moreover, this test is conducted in two distinct ways: in *two-tailed* form, wherein only *same-directional* trends are deemed to be relevant (either upward or downward); and in a *one-tailed* test over $|Z|$, as a way to infer the significance of any trend (upward and downward) in the data.

4.4 Results

4.4.1 Self-rated Exertion

Table 16 presents the exertion levels reported by the participants after each exercise. Generally, all 17 participants reported a higher level of exertion (6-9) after **D**, when compared to the low to moderate (3-6) values registered after **A**, **B** and **C**. This was found to be statistically significant¹⁰ at $\alpha = 0.05$, with $t(37) = -14.612$, $p < .00001$. Therefore it is possible to conclude that in all participants 1) no fatigue is perceived at the end of **P1**; and that 2) **P2** leads to significant exertion and/or physical limitations, potentially caused by muscle fatigue.

4.4.2 Kinematic and Electromyographic Uniformity

In all exercises, a correlation analysis was conducted on 1) the EMG envelope $\mathbf{E}_m^i$ and 2) a subset of right arm kinematics $\mathbf{q}$ (*RightT4Shoulder*, *RightShoulder*, *RightElbow* and *RightWrist*), in order to evaluate movement and behavior uniformity among different participants. Figure 57 and Figure 58 illustrate the magnitude of Spearman's rank correlation coefficient $|r_s|$ in EMG and kinematic data across all partici-

⁹ $\bar{Z}$ follows *approximately* a normal distribution [Kendall, 1970], allowing the use of parametric statistical tests.

¹⁰Using Welch's one-tailed *two-sample t-test* due to different sample size; $H_0: \bar{X}_{\mathbf{D}} > \bar{X}_{\mathbf{ABC}}$.

Table 16: Self-perceived fatigue level [Borg, 1998] reported by each participant after each exercise.

Exer.		Participant																
		S1	S2	S3	S4	S5	S6	S7	S8	S9	S10	S11	S12	S13	S14	S15	S16	S17
P1	A	4	3	3	5	3	2	1	2	2	1	2	2	2	3	2	1	1
	B	5	5	3	5	1	5	3	3	3	2	2	2	5	4	4	1	2
	C	5	4	3	6	4	4	1	3	3	2	4	3	3	3	5	1	2
P2	D	9	8	9	9	8	8	7	7	6	6	7	6	8	8	7	7	7




Figure 57: Illustration of Spearman's rank correlation magnitude $|r_s|$ across EMG data recorded with different participants; data from individual muscles is concatenated sequentially; values closer to 1 imply greater correlation.

pants in different exercises, with participant-specific averages presented in Table 17. Overall, *moderate* and *strong* levels of monotonic correlation¹¹ were observed in both sets of data, suggesting a reasonable degree of similarity and uniformity in how each exercise was executed as well and how muscular recruitment occurred in different participants, with no clear outliers.

4.4.3 General Upper Limb Synergies

Figure 59 presents the variation of the reconstruction VAF with an increasing number of synergies $n \in \{1, \dots, M\}$ (4.3.4.4). Overall, reconstruction quality was found to increase with n ; using a threshold of $\text{VAF} \geq 0.9$, at least 4 MSs were required to accurately reconstruct the EMG input in all subjects, which represents a dimensionality reduction of 50%. Notably, in all but **S7**, 3 MSs were enough to accurately

¹¹Values 0.20-0.39 generally imply *weak* correlation, while the 0.40-0.59 range implies *moderate* correlation [Akoglu, 2018].




Figure 58: Illustration of Spearman's rank correlation magnitude $|r_s|$ across kinematic data recorded with different participants; all axis of the *RightT4Shoulder*, *RightShoulder*, *RightElbow* and *RightWrist* joints are concatenated sequentially; values closer to 1 imply greater correlation.

Table 17: Average rank correlation magnitude $|r_s|$ of EMG and kinematic data across different participants i.e. Figure 57 and Figure 58 column/row averages; higher values imply greater pairwise correlation.

Exer.			Participant																
			S1	S2	S3	S4	S5	S6	S7	S8	S9	S10	S11	S12	S13	S14	S15	S16	S17
EMG	P1	A	0.63	0.51	0.48	0.57	0.48	0.65	0.50	0.42	0.61	0.53	0.72	0.72	0.70	0.70	0.72	0.69	0.65
		B	0.80	0.83	0.64	0.68	0.82	0.81	0.79	0.81	0.84	0.79	0.77	0.76	0.83	0.78	0.79	0.77	0.82
		C	0.71	0.79	0.66	0.79	0.73	0.68	0.80	0.70	0.74	0.60	0.76	0.79	0.75	0.59	0.75	0.69	0.77
	P2	D	0.65	0.80	0.69	0.73	0.73	0.78	0.75	0.77	0.74	0.62	0.79	0.82	0.72	0.74	0.74	0.70	0.75
Kinematics	P1	A	0.63	0.72	0.71	0.65	0.66	0.75	0.65	0.76	0.67	0.62	0.71	0.77	0.62	0.59	0.65	0.76	0.73
		B	0.73	0.78	0.77	0.57	0.77	0.74	0.60	0.75	0.78	0.75	0.74	0.78	0.79	0.59	0.72	0.51	0.69
		C	0.77	0.81	0.80	0.71	0.80	0.67	0.78	0.80	0.74	0.76	0.79	0.73	0.70	0.68	0.77	0.74	0.70
	P2	D	0.84	0.89	0.83	0.87	0.84	0.83	0.88	0.87	0.89	0.86	0.85	0.89	0.89	0.83	0.87	0.88	0.86




Figure 59: Subject and average VAF obtained in **P1** for $n \in \{1, \dots, M\}$; values ≥ 0.9 were assumed to represent an accurate reconstruction of the original EMG envelope **E**.

describe muscular recruitment; nevertheless, with a more in-depth analysis of fatigue effects in mind, the remainder of the study will consider $N_{\text{ABC}} = 4$ MSs as the basic primitives in muscular recruitment during upper-limb movement.

As illustrated in Figure 60, the set of MSs $\mathbf{W}_{\text{ABC}} = [\mathbf{w}^1, \mathbf{w}^2, \mathbf{w}^3, \mathbf{w}^4]$ and the associated activation $\mathbf{H}_{\text{ABC}} = [\mathbf{h}^1, \mathbf{h}^2, \mathbf{h}^3, \mathbf{h}^4]^T$ provide relevant insight into the structural organization of the MSK system: while a significant variance between participants was observed, generally, $\mathbf{w}^3$, with greater contribution from lower-arm muscles, was recruited mostly during **A**; in similar fashion, $\mathbf{w}^1$ and $\mathbf{w}^2$ were, respectively, mainly employed during **B** and **C**; conversely, $\mathbf{w}^4$ exhibited a greater inter-participant variation in both muscle contribution and recruitment. Additionally, antagonistic relationships between different muscles were expressed by their recruitment in contrasting synergies (e.g. BIC/BRR vs TRILON/TRILAT), while the opposite was observed in agonistic muscle groupings (e.g. DELTA/DELTM/DELTP and TRILON/TRILAT).

When performing a localized - around peak activation - comparison of relevant upper-arm kinematics




Figure 60: Synergies $\mathbf{W}_{ABC}$ (left) and activation profile $\mathbf{H}_{ABC}$ (right) obtained in **P1**; for each of the $N_{ABC} = 4$ synergies, average and inter-participant variance ($\pm$ std. deviation) is presented.

$\mathbf{q}_k$ with envelopes $\mathbf{E}_m$ and MS activation $\mathbf{H}_{ABC}$, as presented in Figure 61 for the illustrative case of **S14**, the functional role of each synergy is evident: $\mathbf{w}^1$ was responsible for shoulder abduction and adduction, recruiting mostly DELTP and DELTM muscles, whose activation accordingly peaked in **B** synchronously with an increase in all shoulder joints; recruitment of $\mathbf{w}^2$, representing shoulder flexion and extension (**C**), was correlated with peak recruitment of DELTA, as well as pronounced variations in shoulder y and z angles; $\mathbf{w}^3$ was related with lower arm movement, namely elbow flexion and extension (**A**), which reflected in peak activation from BIC, BRR and TRILAT, concurrently with variations in the elbow angle along the z axis; lastly, $\mathbf{w}^4$ exhibited consistent recruitment throughout **P1**, notably of agonist nature during **A** and **B** (local maxima) and antagonist in **C** (local minima).

During **P2**, as presented in Table 18, large discrepancies were observed in I i.e. how long each participant was able to reproduce **D**, and in measured VAF in each repetition i (Figure 62), which exhibited a considerable range (0.43 - 0.96) across participants, and no clear noticeable trends. Notably, average VAF ($\overline{\text{VAF}}$) exhibited *negligible* rank correlation with I ($r_s(15) = -.11, p = .68$) e.g. participants exhibited both low I and high $\overline{\text{VAF}}$ (**S2** and **S10**); high I and high $\overline{\text{VAF}}$ (**S5** and **S13**); and above average I with

Table 18: Number of executed repetitions (I) and average VAF obtained in exercise **D** by each participant.

Subject																	
	S1	S2	S3	S4	S5	S6	S7	S8	S9	S10	S11	S12	S13	S14	S15	S16	S17
I	31	20	16	12	64	24	33	25	25	32	11	31	48	45	9	34	21
$\overline{\text{VAF}}$	0.79	0.94	0.83	0.87	0.88	0.82	0.84	0.89	0.71	0.91	0.83	0.89	0.82	0.63	0.78	0.68	0.87




Figure 61: Comparison between relevant right-arm kinematics q_k , EMG envelope E_m and MSs recruitment H_{ABC} during **P1** with **S14**, around exercise-specific maxima in $E_m(\cdot)$.

low $\overline{VAF}$ (**S14** and **S16**). When decreasing the previous threshold to $VAF \geq 0.7$ - in order to account for additional motions during **D**, given its composite nature - it was observed that only 14 out of the total 17 participants (82.4%) consistently yielded reasonable results. Under the assumption that this decreased recruitment of W_{ABC} downplayed potential fatigue-induced alterations, data from **S9**, **S14** and **S16** was discarded from synergy analysis (4.3.4.5).

On closer inspection of two extreme cases (high and low VAF) at $i = 1$ (Figure 63), it can be seen that all MSs were recruited, and their functional relevance preserved; however, a decrease in the amplitude of H_D was evident at lower VAF (**S14**), suggesting that generic MSs W_{ABC} were not enough to describe muscular activation, and hinting at an immediate functional reorganization of the MSK system in **P2**. This is in line with what is presented in Table 18, and supports the notion that participants where this phenomenon was observed may not exhibit relevant fatigue-related effects in MS recruitment.

Individual W_{ABC} , as well as similar functional analysis for **P1** and **P2**, are provided in Appendix D: Fatigue-related Alterations in Muscle Synergies for all remaining participants.

4.4.4 Classical Analysis

Generally, as depicted in Figure 63, an increasing trend can be identified in the **RMS** values of PECM, DELTA, DELTM and BIC, as well as an apparent decreasing trend in the **MDF** of the DELTM, TRILAT and TRILON muscles. Notwithstanding visual inspection being insufficient to draw meaningful conclusions regarding the statistical relevance of observed variations, these observations are in accordance with the resulting p -values of the aggregate t-tests (4.3.4.6), presented in Table 19. At a significance level of $\alpha = .05$, relevant trends were identified across all classical indicators; additionally, with the exception of **Int** and




Figure 62: VAF obtained with each participant during exercise **D**, as well as resulting second-order polynomial trend (over i) approximation.




Figure 63: Comparison, for the first repetition ($i = 1$) of **D** for **S12** (Left) and **S14** (Right), of arm kinematics $\mathbf{q}_k$ with EMG envelope $\mathbf{E}_m$ and MS activation $\mathbf{H}_D$; in **P2**, $\mathbf{W}_D = \mathbf{W}_{ABC}$.




Figure 63: Scale- and time-normalized (interpolated) variation of **RMS (Top)** and **MDF (Bottom)** (4.2.3) during **P2**; average and inter-participant variation ($\pm$ std. deviation) presented for each muscle.

Table 19: p -values obtained when conducting the MK test over classical metric values computed during **P2**; results presented for each muscle, and for all participants under a one-tailed test (**any** significant trend) and two-tailed test (for **same**-direction trends); values colored according to the rejection (**green** i.e. trend) or non-rejection (**red** i.e. no trend) of H_0 under a significance level of $\alpha = .05$.

		Metric							
		RMS	MAV	Int	MNF	MDF	TTP	PkF	FR
any	PECM	$4.05 \cdot 10^{-05}$	$1.56 \cdot 10^{-06}$	$6.73 \cdot 10^{-07}$	$5.78 \cdot 10^{-04}$	$2.15 \cdot 10^{-04}$	$3.44 \cdot 10^{-05}$	$8.34 \cdot 10^{-05}$	$1.17 \cdot 10^{-04}$
	DELTA	$4.21 \cdot 10^{-06}$	$6.67 \cdot 10^{-06}$	$6.59 \cdot 10^{-04}$	$5.89 \cdot 10^{-04}$	$1.23 \cdot 10^{-03}$	$4.74 \cdot 10^{-06}$	$2.91 \cdot 10^{-04}$	$4.85 \cdot 10^{-05}$
	DELTP	$9.35 \cdot 10^{-05}$	$7.57 \cdot 10^{-05}$	$1.24 \cdot 10^{-03}$	$2.10 \cdot 10^{-03}$	$3.78 \cdot 10^{-04}$	$4.95 \cdot 10^{-05}$	$1.49 \cdot 10^{-04}$	$8.56 \cdot 10^{-07}$
	DELTm	$1.46 \cdot 10^{-06}$	$2.38 \cdot 10^{-06}$	$6.14 \cdot 10^{-05}$	$1.56 \cdot 10^{-08}$	$2.00 \cdot 10^{-07}$	$1.09 \cdot 10^{-06}$	$3.86 \cdot 10^{-06}$	$1.83 \cdot 10^{-05}$
	TRILON	$1.70 \cdot 10^{-05}$	$1.89 \cdot 10^{-05}$	$4.74 \cdot 10^{-04}$	$1.79 \cdot 10^{-06}$	$2.36 \cdot 10^{-05}$	$1.32 \cdot 10^{-05}$	$2.79 \cdot 10^{-04}$	$3.90 \cdot 10^{-04}$
	TRILAT	$7.07 \cdot 10^{-06}$	$1.02 \cdot 10^{-05}$	$1.31 \cdot 10^{-04}$	$1.40 \cdot 10^{-05}$	$1.25 \cdot 10^{-06}$	$7.70 \cdot 10^{-06}$	$7.36 \cdot 10^{-04}$	$2.08 \cdot 10^{-04}$
	BIC	$2.19 \cdot 10^{-05}$	$4.80 \cdot 10^{-05}$	$1.46 \cdot 10^{-03}$	$3.52 \cdot 10^{-04}$	$2.46 \cdot 10^{-04}$	$1.54 \cdot 10^{-05}$	$2.27 \cdot 10^{-04}$	$1.48 \cdot 10^{-03}$
	BRR	$1.61 \cdot 10^{-05}$	$6.99 \cdot 10^{-06}$	$3.60 \cdot 10^{-04}$	$8.37 \cdot 10^{-05}$	$1.89 \cdot 10^{-04}$	$1.43 \cdot 10^{-05}$	$2.84 \cdot 10^{-03}$	$1.16 \cdot 10^{-03}$
		$2.41 \cdot 10^{-32}$	$1.44 \cdot 10^{-32}$	$1.90 \cdot 10^{-22}$	$1.98 \cdot 10^{-26}$	$3.25 \cdot 10^{-27}$	$8.48 \cdot 10^{-33}$	$2.31 \cdot 10^{-25}$	$6.16 \cdot 10^{-25}$
same	PECM	$1.47 \cdot 10^{-02}$	$1.44 \cdot 10^{-02}$	$2.25 \cdot 10^{-01}$	$5.35 \cdot 10^{-02}$	$8.84 \cdot 10^{-02}$	$1.58 \cdot 10^{-02}$	$5.33 \cdot 10^{-02}$	$2.76 \cdot 10^{-01}$
	DELTA	$1.54 \cdot 10^{-04}$	$2.52 \cdot 10^{-04}$	$3.79 \cdot 10^{-01}$	$1.17 \cdot 10^{-02}$	$1.77 \cdot 10^{-02}$	$1.75 \cdot 10^{-04}$	$1.04 \cdot 10^{-01}$	$7.56 \cdot 10^{-01}$
	DELTP	$5.79 \cdot 10^{-01}$	$6.00 \cdot 10^{-01}$	$2.18 \cdot 10^{-02}$	$1.78 \cdot 10^{-01}$	$1.38 \cdot 10^{-01}$	$4.73 \cdot 10^{-01}$	$9.24 \cdot 10^{-02}$	$4.09 \cdot 10^{-01}$
	DELTm	$7.27 \cdot 10^{-05}$	$5.41 \cdot 10^{-05}$	$9.81 \cdot 10^{-01}$	$3.11 \cdot 10^{-08}$	$4.00 \cdot 10^{-07}$	$6.68 \cdot 10^{-05}$	$9.15 \cdot 10^{-04}$	$9.36 \cdot 10^{-01}$
	TRILON	$4.07 \cdot 10^{-01}$	$4.44 \cdot 10^{-01}$	$1.17 \cdot 10^{-01}$	$3.57 \cdot 10^{-06}$	$4.72 \cdot 10^{-05}$	$4.31 \cdot 10^{-01}$	$8.79 \cdot 10^{-03}$	$9.17 \cdot 10^{-01}$
	TRILAT	$7.63 \cdot 10^{-01}$	$7.59 \cdot 10^{-01}$	$1.44 \cdot 10^{-01}$	$2.00 \cdot 10^{-04}$	$9.16 \cdot 10^{-04}$	$8.16 \cdot 10^{-01}$	$5.61 \cdot 10^{-01}$	$1.07 \cdot 10^{-01}$
	BIC	$2.07 \cdot 10^{-02}$	$3.33 \cdot 10^{-02}$	$4.27 \cdot 10^{-01}$	$3.19 \cdot 10^{-01}$	$7.60 \cdot 10^{-01}$	$2.59 \cdot 10^{-02}$	$7.75 \cdot 10^{-01}$	$7.17 \cdot 10^{-01}$
	BRR	$5.22 \cdot 10^{-02}$	$1.90 \cdot 10^{-01}$	$3.48 \cdot 10^{-01}$	$4.46 \cdot 10^{-01}$	$4.14 \cdot 10^{-01}$	$5.14 \cdot 10^{-02}$	$3.63 \cdot 10^{-01}$	$8.82 \cdot 10^{-01}$
		$7.84 \cdot 10^{-09}$	$5.65 \cdot 10^{-08}$	$8.64 \cdot 10^{-01}$	$3.81 \cdot 10^{-13}$	$6.91 \cdot 10^{-11}$	$8.04 \cdot 10^{-09}$	$1.30 \cdot 10^{-05}$	$9.74 \cdot 10^{-01}$

FR, there is agreement in trend direction across all metrics. In each case, trend magnitude and direction can be inferred from the sign of the normalized MK test statistic Z , as listed in Table 20: increasing trends were prevalent in **RMS**, **MAV** and **TTP** values, particularly in DELTA and DELTM; conversely, decreasing trends in frequency-domain metrics **MNF**, **MDF** and **PkF** were detected in all muscles under observation, notably in DELTP, TRILON and TRILAT; on average, the magnitude of the detected trends was greater in frequency-domain metrics.

Regarding the nature of each of these variations, low determination coefficient R^2 values were registered in the polynomial models obtained during linear regression (Table 21), highlighting significant non-linearity in relevant trends; this was observed across all classical metrics. Notwithstanding, a slightly better fit was obtained for **RMS** and **MAV** in DELTA, and for **MNF** and **MDF** in DELTM. On the other hand, **Int** and **FR** - which did not exhibit significant trends during **P2** (Table 19) - resulted in the lowest values.

Detailed participant-specific results (p , Z and R^2) obtained for each classical EMG metric are provided in Appendix D: Fatigue-related Alterations in Muscle Synergies.

Table 20: Average Mann-Kendall normalized test statistic Z across all participants for each muscle and classical metric (Table 13); higher absolute magnitude (**green**) implies higher variation, while values closer to 0 (**red**) suggest no significant trend; t-tests where $p > .05$ (Table 19) greyed out.

	Metric							
	RMS	MAV	Int	MNF	MDF	TTP	PkF	FR
PECM	4.24	3.41	1.12	-4.27	-3.35	3.96	-2.63	-1.70
DELTA	8.27	8.33	1.64	-5.54	-4.97	8.16	-2.17	-0.47
DELTP	1.05	1.00	-4.03	-3.64	-2.66	1.25	-3.08	-1.34
DELTM	5.95	6.08	-0.49	-9.97	-9.19	5.75	-3.88	0.31
TRILON	0.98	1.01	1.69	-6.15	-6.22	0.88	-3.65	0.61
TRILAT	0.46	0.47	-2.53	-7.29	-5.24	0.33	-1.32	2.24
BIC	4.76	4.36	-4.09	-3.02	-0.94	4.31	0.84	1.52
BRR	3.08	2.07	0.83	-1.16	-1.08	3.10	-2.12	-0.78
	3.60	3.34	-0.73	-5.13	-4.21	3.47	-2.25	0.05

Table 21: Average determination coefficient R^2 obtained for the polynomial models (4.3.4.6) for all classical metrics; higher values (**green**) imply a better model fit; values where $p > .05$ (Table 19) greyed out.

	Metric							
	RMS	MAV	Int	MNF	MDF	TTP	PkF	FR
PECM	0.27	0.29	0.06	0.23	0.19	0.26	0.08	0.06
DELTA	0.60	0.63	0.09	0.28	0.22	0.58	0.07	0.05
DELTP	0.28	0.27	0.05	0.24	0.18	0.27	0.11	0.06
DELTM	0.44	0.44	0.04	0.74	0.62	0.44	0.14	0.06
TRILON	0.25	0.24	0.05	0.36	0.34	0.25	0.13	0.05
TRILAT	0.33	0.32	0.07	0.52	0.38	0.32	0.11	0.04
BIC	0.33	0.30	0.03	0.27	0.23	0.31	0.01	0.06
BRR	0.26	0.30	0.04	0.18	0.12	0.23	0.11	0.09
	0.34	0.35	0.05	0.35	0.29	0.33	0.11	0.06




Figure 64: Scale- and time-normalized (interpolated) variation of **RMS** (Left) and **WaHP** (Right) values computed from $\mathbf{H}_D$ during **P2** across all subjects. Average and inter-participant variation ($\pm$ std. deviation) presented for each MS.

4.4.5 Synergy Analysis

As illustrated in Figure 64, across all different participants, the distribution of **RMS** values suggests a slight amplitude increase in MS recruitment; conversely, **WaHP** did not exhibit any visually consistent trend during **P2**.

A superficial analysis of potential cumulative effects muscle fatigue had on synergy recruitment during **P2** is presented in Figure 65 for the case of **S5** (participant with highest K , Table 18), as an overlap of $\mathbf{H}_D^i$ obtained during **D** for $i \in \{1, 2, \dots, I\}$. In general, it can be seen that the overall shape of $\mathbf{H}_D^i$ was preserved across all repetitions. However, an amplitude increase is noticeable, in line with what was expected from the results in time-domain metrics observed in conventional indicators (4.4.4). In addition, peak recruitment (local maxima) appeared to gradually occur *earlier* with each repetition, which has been shown to occur in postural control [Monjo & Forestier, 2016; Strang & Berg, 2007], although it is not clear if this phenomenon was consistent throughout the whole dataset.

The results of the aggregate t-tests conducted over Z in **P2** for all synergy metrics (Table 15) are presented in Table 22 and Table 23. Statistically relevant variations were observed in all metrics, consistently in the same upward or downward direction with the exception of **iPk** and **WaHP**. On the other hand, no-table inter-MS differences were observed: $\mathbf{w}^2$ exhibited significant and consistent directional **WaHP** trends throughout the data, while in $\mathbf{w}^4$ this only occurred with **P2P**; this aligns with the greater contribution of DELTA, which registered greater magnitude trends in conventional indicators (4.4.4), in $\mathbf{w}^2$ (Figure 60). Values of the MK normalized test statistic Z (Table 23) suggest strong increasing trends in most studied metrics *i.e.* **Energy**, **RMS**, **MAV**, **P2P**, **Int** and **TTP**; in particular, $\mathbf{w}^2$ registered not only the greatest trend magnitude in **Energy**, **RMS**, **P2P** and **TTP**, but also decreasing trends in **WaHP**.



 Figure 65: Synergy recruitment $\mathbf{H}_D^i$ for $i \in \{1, 2, \dots, 64\}$ observed during $\mathbf{D}$ in the case of $\mathbf{S5}$.

 Table 22: p -values obtained when conducting the MK test over synergy metric values computed during $\mathbf{P2}$; Results presented for each synergy, and for all subjects under a one-tailed test (**any** significant trend) and two-tailed test (for **same**-direction trends); Values colored according to the rejection (**green** i.e. trend) or non-rejection (**red** i.e. no trend) of H_0 under a significance level of $\alpha = 0.05$.

		Metric							
		Energy	RMS	MAV	iPk	P2P	WaHP	Int	TTP
any	w^1	$6.70 \cdot 10^{-07}$	$1.12 \cdot 10^{-06}$	$3.11 \cdot 10^{-05}$	$1.95 \cdot 10^{-03}$	$8.00 \cdot 10^{-09}$	$7.73 \cdot 10^{-04}$	$4.35 \cdot 10^{-05}$	$3.87 \cdot 10^{-06}$
	w^2	$1.38 \cdot 10^{-05}$	$1.33 \cdot 10^{-05}$	$6.11 \cdot 10^{-05}$	$4.00 \cdot 10^{-04}$	$1.11 \cdot 10^{-05}$	$8.21 \cdot 10^{-05}$	$2.30 \cdot 10^{-04}$	$7.03 \cdot 10^{-06}$
	w^3	$5.80 \cdot 10^{-05}$	$4.40 \cdot 10^{-05}$	$1.14 \cdot 10^{-04}$	$2.37 \cdot 10^{-06}$	$3.84 \cdot 10^{-05}$	$3.21 \cdot 10^{-04}$	$3.42 \cdot 10^{-04}$	$1.06 \cdot 10^{-04}$
	w^4	$1.23 \cdot 10^{-05}$	$2.86 \cdot 10^{-05}$	$6.28 \cdot 10^{-05}$	$4.45 \cdot 10^{-05}$	$1.26 \cdot 10^{-04}$	$4.49 \cdot 10^{-04}$	$1.93 \cdot 10^{-04}$	$1.08 \cdot 10^{-05}$
		$4.39 \cdot 10^{-17}$	$2.78 \cdot 10^{-17}$	$3.64 \cdot 10^{-14}$	$1.46 \cdot 10^{-13}$	$1.69 \cdot 10^{-16}$	$1.12 \cdot 10^{-12}$	$2.54 \cdot 10^{-12}$	$4.74 \cdot 10^{-16}$
same	w^1	$2.56 \cdot 10^{-02}$	$2.45 \cdot 10^{-02}$	$2.16 \cdot 10^{-02}$	$8.71 \cdot 10^{-02}$	$1.38 \cdot 10^{-01}$	$4.91 \cdot 10^{-01}$	$1.52 \cdot 10^{-02}$	$2.64 \cdot 10^{-02}$
	w^2	$2.75 \cdot 10^{-05}$	$2.66 \cdot 10^{-05}$	$1.22 \cdot 10^{-04}$	$3.42 \cdot 10^{-01}$	$1.07 \cdot 10^{-04}$	$4.76 \cdot 10^{-02}$	$4.60 \cdot 10^{-04}$	$1.41 \cdot 10^{-05}$
	w^3	$1.73 \cdot 10^{-02}$	$1.50 \cdot 10^{-02}$	$1.27 \cdot 10^{-02}$	$8.03 \cdot 10^{-01}$	$5.74 \cdot 10^{-02}$	$3.48 \cdot 10^{-01}$	$1.93 \cdot 10^{-02}$	$2.08 \cdot 10^{-02}$
	w^4	$5.09 \cdot 10^{-02}$	$7.84 \cdot 10^{-02}$	$1.07 \cdot 10^{-01}$	$2.87 \cdot 10^{-01}$	$2.85 \cdot 10^{-02}$	$2.69 \cdot 10^{-01}$	$2.98 \cdot 10^{-01}$	$6.52 \cdot 10^{-02}$
		$7.60 \cdot 10^{-08}$	$9.42 \cdot 10^{-08}$	$6.51 \cdot 10^{-07}$	$4.17 \cdot 10^{-01}$	$1.25 \cdot 10^{-06}$	$1.14 \cdot 10^{-01}$	$6.15 \cdot 10^{-06}$	$2.07 \cdot 10^{-07}$

Table 23: Average Mann-Kendall normalized test statistic Z across all subjects for each MS and metric (Table 15); higher absolute magnitude (green) implies higher variation, while values closer to 0 (red) suggest no significant trend; t-tests where $p > .05$ (Table 22) greyed out.

	Metric							
	Energy	RMS	MAV	iPk	P2P	WaHP	Int	TTP
w^1	3.08	3.18	4.24	-3.37	2.02	1.34	3.76	3.19
w^2	6.77	6.67	5.43	-1.98	6.81	-3.90	5.42	5.80
w^3	5.43	5.47	5.62	0.24	3.62	-1.55	5.80	4.98
w^4	2.94	2.88	2.16	1.94	3.43	-1.71	1.21	2.65
	4.557	4.55	4.36	-0.79	3.97	-1.45	4.05	4.15

Table 24: Average determination coefficient R^2 obtained for the polynomial models (4.3.4.6) for all synergy metrics; higher values (green) suggest a better model fit; values where $p > .05$ (Table 22) greyed out.

	Metric							
	Energy	RMS	MAV	iPk	P2P	WaHP	Int	TTP
w^1	0.35	0.35	0.39	0.10	0.30	0.15	0.34	0.36
w^2	0.43	0.45	0.38	0.10	0.47	0.21	0.35	0.42
w^3	0.43	0.45	0.50	0.17	0.20	0.14	0.45	0.457
w^4	0.32	0.33	0.34	0.12	0.23	0.11	0.32	0.33
	0.38	0.39	0.40	0.12	0.30	0.15	0.36	0.39

Additionally, with regard to the linearity of observed trends, low R^2 averages were obtained for all metrics across all participants; it can thus be stated that, in all synergies, metric values varied nonlinearly with time and thus the development of muscle fatigue.

Detailed participant-specific synergy analysis results (p , Z and R^2 for each metric) are provided in Appendix D: Fatigue-related Alterations in Muscle Synergies.

4.5 Discussion

Overall, the experimental methodology adopted in this study, including the fatigue-inducing experimental protocol and the data processing pipeline, provided an objective analysis of the cumulative impact of low-force repetitive movements on the MSK behavior and performance in healthy individuals, focusing on the occurrence of muscle fatigue. Beyond demonstrating the compositional nature of the human MSK system and how primitive elements can emerge as a strategy to circumvent its natural high-dimensionality and redundancy, the results allowed the identification of multiple objective metrics as accurate fatigue indicators with appeal in practical, real-world environments, and which can be exploited in the development of effective long-term fatigue-management solutions.

Recorded experimental data, as well as the output of the processing pipeline, was in line with expectations, considering the theoretical background on upper-limb muscular function and agonist-antagonist

relationships [Gabriel, 1997; Sella, 2021; Wierzbicka et al., 1986]. Considerable uniformity was observed in the dataset, and muscular fatigue was shown to occur through statistically significant variations in both self-reported levels of exertion (4.4.1) as well as in conventional time- and frequency-domain fatigue indicators (4.4.4). Subjective differences in the behavior of the MSK system during **P1** and **P2** were expected, which can be explained by different degrees of muscular development between participants, as a result of age, different exercising habits, rest state or stimulant consumption. Nonetheless, in general, low levels of exertion were reported after **P1** (4.4.1), despite the different levels of fatigue observed between participants (Figure 62) in **P2**. While none of the exercises reproduced in this study is isometric - and therefore some of the theoretical background (4.2.2) may not be applicable - care was taken to ensure all participants reproduced them in a stereotypical and repetitive manner, minimizing variability; notwithstanding, the use of a constant load in all participants was expected to influence fatigue onset during the experiment.

The generic MSs in $\mathbf{W}_{ABC}$, identified during **P1**, were consistent across the whole dataset (Figure 60), with $N_{ABC} = 4$ (representing a dimensionality reduction of 50% from $M = 8$) shown to be sufficient to accurately ($VAF > 0.9$) describe muscular behavior in all participants; while, in the particular case of **S7**, slightly lower values were obtained when compared to the remaining subjects, it should be noted that reconstruction quality was still reasonable ($VAF > 0.8$) with $n = 3$; From a functional standpoint, as demonstrated in Figure 60 and Figure 61, $\mathbf{W}_{ABC}$ agreed with the motions being performed; particularly, each **P1** exercise relied mostly on a single MS - potentially related with how gravity assists the downward portion of each motion, decreasing muscular recruitment and overall effort - with some overlap observed between $\mathbf{w}^1$ and $\mathbf{w}^2$, which was anticipated considering that **B** and **C** both focus on the shoulder joint. Conversely, higher participant variability in $\mathbf{w}^4$, both in muscle contribution and overall activation during **P1**, implies a general support role in upper limb movement; it can be hypothesized that an additional synergy in $\mathbf{w}^4$ allows for greater functional consistency in $\mathbf{w}^1$, $\mathbf{w}^2$ and $\mathbf{w}^3$, by focusing on participant-specific particularities (e.g. fitness) and residual noise. Moreover, antagonist and agonist relationships between synergies can be identified in $\mathbf{H}_{ABC}$, as, respectively, symmetrical and asymmetrical concurrent activation peaks e.g. $\mathbf{w}^1$ and $\mathbf{w}^2$ are recruited concurrently during **C**, which correlated with local minima in the activation of $\mathbf{w}^3$ (Figure 61).

When fitting $\mathbf{W}_{ABC}$ to **P2**, varying degrees of reconstruction quality were obtained for different participants, as presented in Figure 62 and Table 18. Generally, VAF was lower than during **P1**, which, considering the added complexity in the composite motion in **D** (Figure 52), was - to some extent - expected. Nevertheless, as the simpler motions executed during **P1** and reproduced during **P2** were intended to be represented by $\mathbf{W}_{ABC}$, relatively high VAF values were expected across all subjects throughout the exercise. Moreover, under the assumption that the underlying synergistic structure is preserved once muscle fatigue occurs, reconstruction quality should not decrease as i increases. In practice, high VAF was indeed measured in most cases (Table 18) during **D**. From a functional standpoint, although activation is decreased at lower VAF (Figure 63) all 4 MSs are employed at least once during each segment, and concurrently with the kinematic primitives previously discussed. Therefore, it can be stated that *simpler uni-articular MSs are repurposed - fully or partly - in complex compound movements.*

On the other hand, contrasting behavior was observed among trial participants, such as 1) consistently

high or low VAF; 2) increasing, stationary or decreasing trends in VAF (Figure 62); as well as 3) large variance in I (Table 18). While, as previously mentioned, such variability can be explained in part by different muscular fitness or force generating capability among the subject pool - which should in theory correlate with variable fatigue resistance - the observed disparity in reconstruction patterns was not expected when taking into account the moderate to high levels of kinematic and electromyographic correlation measured during **D**. Objectively, higher VAF implies that $\mathbf{W}_{ABC}$ was able to better describe muscular recruitment during **P2** and, conversely, a lower value entails a partial or complete restructuring of the MSK system, whether through additional (*i.e.* $N_D > N_{ABC}$) or completely different set of MSs [Ivanenko et al., 2005]. While there is scientific basis for such a reorganization under fatiguing conditions [Smale et al., 2016; Thomas et al., 2023], this contradicts the underlying premise of the present study; instead, structural changes to $\mathbf{W}_{ABC}$ that may occur as a result of *e.g.* compensatory movements gradually adopted by participants under significant exertion, were not expected in a rested state or during the early stages of muscle fatigue. In this study, low VAF was observed in specific participants right from the start of **D** ($i = 1$); furthermore, no clear decreasing trend was present across different participants. As such, any potential MSK restructuring was presumed to not be fatigue-induced, but rather the result of participant-specific compensatory mechanisms or fatigue-management strategies employed by the MSK system. Alternative interpretations of this phenomenon include 1) additional motions executed during **P2** that were unaccounted for in **P1**, either in the transition between elemental primitives or by deviations from the experimental protocol (*e.g.* distractions, poor mind-muscle connection, unreported joint mobility or muscle tightness issues); 2) the movement order (Figure 52), as well as the dynamic nature of the composite motion, influences CNS and MSK organization *e.g.* transitions too fast to recruit specific MSs (as illustrated in Figure 63 (**Right**) as decreased activation of $\mathbf{w}^2$ in contrast with its agonist $\mathbf{w}^1$), or 3) alternative synergistic relationships provide better task performance. Particularly, in the case of **S7**, a noted lower VAF for $n = 2$ and $n = 3$ in **P1** (Figure 59) hints at the relevance of additional MSs (increase in N_{ABC}) employed during **P2** which may be associated with increased fatigue resistance. All these observations highlight the subjective nature of MSK behavior *i.e.* how each subject counteracts muscular fatigue, and how this impacts how long can a specific motion be sustained (I). It is thus established that fatigue resistance goes beyond muscular capacity, with MSK systems adopting alternative strategies to counterbalance or accommodate fatigue, as evidenced in the literature [Fuglevand, 1996; Taylor et al., 2016].

In accordance with the objective of the experimental protocol, the results of the classical analysis (4.4.4) demonstrated that significant alterations occurred in all participants throughout **P2**. This supports the hypothesis that muscles were fatigued towards by the end of the trials (4.4.1). As discussed in 4.2.2, the expected compression of the EMG frequency spectrum would, in theory, be associated with a slower response of the muscles, together with a higher activation peak; as shown in Figure 63, a noted increase in average **RMS** amplitude and a simultaneous decrease in average **MDF** values, measured during **P2** across all participants, attest to this phenomenon; statistically, increasing trends in **RMS**, **MAV** and **TTP** and decreasing trends in **MNF**, **MDF** and **PkF** (Table 20) further substantiate this theoretical impact of fatigue in EMG recordings. While the magnitude and significance of these trends differed between muscles - foreseeable due to varying degrees of fitness, capacity and exertion - generally, DELTA and DELTM exhibited

a larger variation, in agreement with the pronounced average slopes depicted in Figure 63, and which can be explained by the focus of the experimental protocol in the shoulder joint. In contrast, no significant directional trend was observed in BRR, implying variable recruitment level in trial participants. Moreover, a notable divide between time- and frequency-domain metrics was evident. Frequency-domain metrics, namely **MNF** and **MDF**, exhibited the highest absolute trend magnitude (Table 20) and, contrary to the time-domain metrics, presented a meaningful directional variation in the TRILON and TRILAT muscles. Conversely, in PECM and BIC muscles, the opposite is observed since relevant trends were observed only in time-domain metrics. Despite the fact that an EMG amplitude increase is often described as byproduct of its spectral compression in the presence of muscle fatigue [D'Alessio & Conforto, 2001; Farina & Falla, 2008; Hogrel et al., 2008], these results suggest that both phenomena may be dependent on muscle fiber type, the dynamics of the executed task, as well as force generating capabilities and exertion level; Notwithstanding, this contrasts with the notion that spectral alterations are described as often not significant or statistically relevant [Chang, 2018], and highlights **MNF** and **MDF** as potential fatigue indicators. Lastly, the low R^2 values (Table 21) measured during linear regression analysis were, on the surface, not surprising when considering that fatigue models found in the literature often follow exponential capacitor-like models [Chang, 2018; Peternel et al., 2019].

Notably, synergy metrics (Table 15) were informed by the expected alterations to the temporal profile of EMG recordings *i.e.* smaller slopes - implying slower activation of muscle fibers - and higher amplitude peaks (4.2.2). In fact, when analyzing the evolution of values throughout **D**, as presented in Figure 64, a slight increase in **RMS** average amplitude was observed in all MSs. On the other hand, no upwards or downwards trend in **WaHP** was easily discernible. The first is particularly evident, for instance, in the case of **S5** (highest recorded I), as shown when overlapping activation curves H_b^i (Figure 65). From a statistical standpoint, the results of the MK test (Table 22 and Table 23), demonstrating non-negligible variations in all metrics across the different participants, support the notion that *muscle fatigue leads to significant alterations in MS activation*. However, in the case of **iPk** and **WaHP**, such variations did not occur in the same direction across different MSs, hinting at synergy- or participant-specific covariates (e.g. prioritizing certain MSs or different fatigue management behaviors) that influence the position and width of activation peaks, respectively; the latter case is in accordance with the lack of visible trends in Figure 64 (**Right**). While increasing tendencies were observed in all relevant metrics across different MSs (Table 23), in the specific case of **WaHP**, a significant decreasing trend was observed in w^2 . In fact, w^2 exhibited more pronounced variations in most metrics, being therefore hypothesized that PECM and DELTA - its main contributors - are less fatigue-resistant than the remaining muscles under study. This is supported by the outcome of classical analysis (4.4.4), wherein DELTA displayed highest magnitude trends in time-domain metrics. By contrast, in the case of w^4 only **P2P** varied consistently in all participants, supporting the previous allegations about its subjective support role in upper limb movement; this is coherent with the lower recruitment observed throughout **P1** and **P2** - which, in theory, presupposes less pronounced fatigue-induced alterations. Overall, the increase in amplitude-related metrics can be hypothesized to be reflective of an increase in muscle coactivation, which has been shown to be correlated with an increase in central (*i.e.* corticospinal) effort to counteract the increased difficulty in fatigued muscle recruitment [Tallent

et al., 2021]. In that sense, while statistical analysis does not imply meaningful changes in both the **iPk** and **WaHP** across trial participants, the increase of the corticospinal volley, which increases the likelihood of recruitment of adjacent interneurons, may lead to *wider* and *earlier* peaks in the synergy activation in some cases (Figure 65). These results suggests that, generally, in movements characterized by burst contraction, muscle fatigue leads to an increased recruitment of generic MSs *i.e.* peaks in $\mathbf{H}_D^i$ become gradually higher; as products of the envelope factorization process and thus reflective of alterations in the original EMG, this is in line with the theoretical background and the outcome of classical analysis (4.4.4). In addition, while it is not possible to draw general conclusions as to *when* or *how long* peak recruitment occurs as muscle fatigue develops, the results presented herein reflect a localized decrease in peak duration in MSs recruiting more fatigued or less fatigue-resistant units *i.e.* effort is sustained for less time, potentially reflecting spectral alterations in EMG patterns. Nevertheless, from a practical perspective, this phenomenon could potentially be seen as an indicator of extreme fatigue or near-failure conditions, and validate the hypothesis that *objective MS-based metrics can be used to identify muscle fatigue*. Lastly, in line with the nonlinearity measured in EMG classical metrics (4.4.4, Table 21), linear models were found to not fit the evolution of synergy metrics, as evidenced in the low R^2 values in Table 24 *i.e.* the gradual cumulative effect of fatigue in $\mathbf{H}_D^i$ varied nonlinearly with time and effort.

Globally, all of these insights provide an objective answer to **RQ3**: repetitive, low-force motions induce a nonlinear increase in the overall recruitment peaks of MSs, as objectively measured with **Energy**, **RMS**, **MAV**, **P2P**, **Int** and **TTP** (Table 15) - exceeding the established KPI of >1 relevant marker. Notwithstanding, practical applications, as well as a more extensive benchmark between different metrics, should be evaluated on a problem- or task-specific basis.

4.6 Conclusion

In this work, the occurrence of muscular fatigue was investigated in the context of repetitive non-isometric small-load movements of the upper limb, focusing on its impact on the synergistic and functional structure of the MSK system. Under the assumption that, as fatigue gradually develops, the multiple Muscle Synergies (MSs) underlying generic upper-limb movement are preserved, while primary motor cortexes dynamically adjust their recruitment to compensate the diminishing force-generating capacity of fatigued muscles, the procedure described in this manuscript successfully demonstrated that: 1) MSs responsible for uni-articular motions are reused in compound movements, while preserving their functional relevance, both *before* and *after* fatigue onset; 2) synergy recruitment suffers significant alterations with fatigue; and 3) these changes can be accurately and consistently identified with objective metrics.

Employing a total of 17 healthy individuals, a two-stage experimental protocol gradually induced muscle fatigue in right-arm muscles through repetitive motions until self-reported exertion or visible kinematic degradation, while recording detailed kinematic and EMG data. MSs were estimated from simple uni-articular motions in a non-fatigued state, then fit to gradually changing EMG patterns as fatigue develops, and the resulting activation profile in each proceeding repetition characterized through multiple time- and frequency-domain metrics. Non-parametric statistical testing and linear regression analysis were used

to evaluate the significance of measured metric trends, and thus objectively assess the usability of each metric as a muscle fatigue indicator: a nonlinear increase in **Energy, RMS, MAV, P2P, Int** and **TTP** was observed, while classical EMG analysis validated the occurrence of muscle fatigue in trial participants.

Given the results, it was concluded that *low-force, repetitive motions lead to a nonlinear increase in the overall recruitment peaks of MSs*, objectively answering **RQ3**. In a practical sense, the identification of 6 synergy-based metrics as indicators of cumulative muscle fatigue went beyond the preset KPIs and represents a notable scientific contribution of this study; moreover, framing synergy metrics (Table 15) as contextual primitive-based markers, objective **06** was accomplished, and the estimation of similar synergistic structures across multiple participants is understood to also achieve **01** and **05**.

Conceptually, within the context of this thesis, this chapter explored a new category of primitive elements in MSs, complementing the primitive paradigms explored in Chapter 2 and Chapter 3. Results not only illustrate the highly modular and compositional nature of biological movement - reasserting the general appeal of primitive theory at different levels of motor generation and control e.g. kinematic, dynamic, or neuromuscular - but also the generality of primitive elements, often task- and context-agnostic, herein reflected in MSs reused in different exercises, and persistent after fatigue onset; most notably, it is demonstrated that the adaptive nature of vertebrate MSK systems can be harnessed for the design of better collaborative, assistive or rehabilitative automated systems, such as intelligent fatigue management strategies, able to reliably predict fatigue onset, optimize long-term effort of human collaborators in industrial settings, and - ultimately - reduce incidence of WRMSDs.

Nevertheless, the study described herein still presents just an initial step towards such an automated fatigued detection and management solution; despite the relevant insights regarding the effect of muscle fatigue in MS recruitment, and the validation on low-force non-isometric movements, the presented analysis was performed offline and restricted to simulated right-arm tasks in an academic, structured environment. In the short term, research should focus on 1) the construction of a more extensive dataset, potentially exhibiting a greater demographic range, MSK pathologies, physical limitations and active lifestyles; 2) developing and cross-validating a real-time fatigue detection strategy employing the MS metrics identified in this work, providing an objective measure of *detection time*; 3) accurate modeling of the detected trends, through higher-order regression analysis and different exponential, logarithmic or auto-regressive models; 4) exploring accurate surrogate MSK models, discarding the need for invasive and uncomfortable EMG instrumentation; 5) validating fatigue detection in collaborative scenarios and industrial tasks.

Conclusion

This thesis provided a multidisciplinary and practical overview of primitive theory across different areas of research and engineering. To that end, three studies, tackling different practical scientific and engineering problems, demonstrated how primitive elements - whether as Motor Primitives (MPs), Perceptual Primitives (PPs) or Muscle Synergies (MSs) - can be leveraged by researchers and engineers in the development of novel, automated and intelligent systems, whose application ranges from domestic, industrial and collaborative robotics to rehabilitation in clinical settings. With a particular focus on bipedal locomotion (Chapter 2 and 3), robotic control (Chapter 2) and human behavioral monitoring (Chapter 3 and 4), the concepts and solutions presented herein exploit and repurpose dynamical systems theory, flexible representations frameworks e.g. Dynamical Movement Primitives (DMPs) and learning techniques e.g. Black Box Optimization (BBO) to develop adaptive and versatile strategies, models and tools.

In the remainder of this chapter, each study is summarily described, highlighting main takeaways and outcomes, and presenting relevant scientific contributions in the overarching context of this thesis; in addition, major shortcomings and limitations are identified, and framed as appealing avenues for future scientific research.

5.1 Summary

Generic Primitive-based Legged Control

Tackling the problem of adaptive and stable walking in humanoid robots, Chapter 2 aimed to answer **RQ1**: *How can a generic primitive-based controller generate functional locomotion patterns in different platforms?* To that end, building upon previous research, a generic high-level control structure was proposed, designed on top of Dynamical Movement Primitives (DMPs), and with a particular focus on MP sharing across different robotic platforms, exploiting primitive theory to simplify and generalize motor generation and skill acquisition. Using bioinspired locomotive limit cycles based on Central Pattern Generators (CPGs), a versatile Imitation Learning (IL) strategy is developed employing Black Box Optimization (BBO), and used to teach functional, stable locomotion on flat terrain to two simulated anthropomorphic plat-

forms. The notable aspect of the solution was that in both robots lower-limb control patterns were defined as simple combinations of only $P = 3$ generic MPs, with the resulting gait exhibiting over 90% average similarity in selected gait metrics including distance traveled, average velocity, total energetic expenditure and Cost of Transport (CoT)¹. In addition to highlighting how MPs can effectively simplify representation and dimensionality without sacrificing stability and performance, the generic nature of the control architecture - formalizing generic MP and robot-specific *skill* libraries - fits generalized adoption in other robotic tasks, as well as in L2 strategies during the design of autonomous adaptive machines. While the design of fast and stable *i.e.* highly-performant locomotion was not the primary focus of the study, the methodology presented provided an objective answer to **RQ1**: *generic bioinspired MP-based control, exploiting non-linear dynamical systems for motor generation, can effectively be used to imitate complex, highly-dynamic and high-dimensional behaviors such as walking in multiple platforms*. Within the broader context of this thesis, the study accomplished **O1** through an IL approach to motor skill learning from CPG presets, **O3** by achieving stable walking in different platforms, and provided the general structure for a unified DMP-based knowledge database that would be repurposed in Chapter 3, thus partially fulfilling **O2**.

Dynamical Perceptual Memory Model

Focusing on identifying an objective answer to **RQ2**: *How can dynamical systems encode stereotypical episodic information associated with different perceptual contexts, while accounting for primitive elements?*, the DMP-based trajectory introduced in Chapter 2 was repurposed for a different - albeit complementary - application, providing the basis for a dynamical episodic memory model targeting *stereotypical* perceptual input in a generic sensor-agnostic (*i.e.* amodal) representation, while borrowing insights from decades of research in human cognition. The appealing properties of DMP are carried over, leading to a data-driven, compact and time-invariant representation of perception which can complement high-level cognitive architectures. Most notably, the existence of shared Perceptual Primitives (PPs) across different sensor dimensions, modalities as well as - potentially - different tasks and agents/platforms. The previous motor skill learning approach is reframed as a memory consolidation procedure and, additionally, practical algorithms are proposed for memory recall, memory matching and perceptual context identification. Cross-validation was conducted in a kinematic dataset of human locomotion; in practice, when modeling motion capture data as perceptual input, the proposed setup was able to accurately identify walking velocity and terrain incline with, on average, 89% accuracy¹ and over 95% specificity; and perturbed, non-stereotypical gait was detected with 79% accuracy and up to 30 ms in advance. At the same time, substantial dimensionality reduction of over 20%¹ was achieved when sharing $P = 5$ PPs across a total of $K = 18$ perceptual inputs. Beyond verifying established KPIs, the resulting performance demonstrates how primitive structures can be exploited in perceptual modeling, and how stereotypical experiences can be efficiently represented with dynamical systems, answering **RQ2**: *coupled DMP systems can be harnessed to uniquely encode stereotypical perceptual patterns, as well as efficiently represent PPs in sensor data..* **O1** is once again realized by creating perceptual memories from recorded kinematic demonstrations; **O2**

¹ Established Key Performance Indicator (KPI).

is completed by adopting a similar structure to encode motor and perceptual patterns, **04** through the construction of a perceptual memory library and **05** by integrating generic PPs into the dynamics underlying episodic memories.

Impact of Muscle Fatigue in Musculoskeletal Primitives

A complementary study was conducted on repetitive and stereotypical human upper-limb movement, and primitives at neuromuscular level *i.e.* Muscle Synergies (MSs) were exploited in the identification of muscular fatigue from Electromyography (EMG) recordings in dynamic, low-force movements, investigating **RQ3**: *What cumulative effect do low-force repetitive movements have in the activation profile of muscle synergies?* Muscular fatigue is contextualized as a relevant socioeconomical issue in industrial settings, where human-robot collaborative strategies can contribute to minimize long-lasting impact in work performance and Musculoskeletal (MSK) disorders. To that end, an experimental protocol was designed with the aim of gradually inducing muscular fatigue in a total of 17 healthy individuals, as well as validating the compositional nature of the human MSK system. Operating under the well-documented premise that the synergistic structure of the MSK system is preserved after fatigue onset, an elaborate processing pipeline was able to comprehensively profile muscular activity of trial participants over time *i.e.* as fatigue developed. This stage included conventional EMG markers, functional and kinematic assessment of different exercises, and detailed statistical and regression analysis over multiple time- and frequency-domain metrics, regarding the level of significant alterations observed from a non-fatigued to a fatigued state. The work concluded that **Energy, RMS, MAV,P2P, Int** were able to identify fatigued state¹; in addition, observed variations were non-linear across all tested metrics, formulating a concrete answer to **RQ3**: *low-force, repetitive motions lead to a nonlinear increase in the overall recruitment peaks of MSs*. In a practical sense, these results can be used in the development of dynamic fatigue models and exploited by high-level fatigue management strategies, with particular appeal in adaptive robots closely interacting or coexisting with human collaborators. At a higher, conceptual level, this study complements the previous explored primitive modalities (MPs and PPs) by further reinforcing the appeal of primitive elements in the design of intelligent adaptive machines. Through the use of experimental data in the extraction of upper-limb MSs, **01** and **05** are achieved, and their use to identify muscle fatigue accomplishes **06**.

Overall, all practical objectives were fulfilled, with **01** and **05** achieved in multiple chapters. At a global, conceptual level, the research conducted within the context of this thesis exemplifies multiple features of bioinspired Lifelong Learning (L2) systems [Kudithipudi et al., 2022], as summarized in Table 25:

- **Transfer and Adaptation** when fitting learned behaviors and preset motor policies to different robotic platforms, illustrated in the skill learning (Chapter 2, 2.4.3) and two-stage memory consolidation (Chapter 3, 3.4.4) procedures, and when extracting similar MSs across different individuals (Chapter 4, 4.3.4.4);
- **Exploiting Task Similarity** through the use of a shared MP library in different platforms (Chapter 2, 2.4.3), two-stage memory consolidation (Chapter 3, 3.4.4), and reusing uni-articular MSs in

composite movements (Chapter 4, 4.3.4.4);

- **Overcoming Catastrophic Forgetting** by implementing compact, data-efficient episodic memories of stereotypical experiences (Chapter 3, 3.3), allowing the incorporation of associative semantics between motor skills $\mathcal{S}$ and perceptual memories $\mathcal{M}$;
- **Task-agnostic Learning** which, while not directly addressed in this work, can be conducted through offline extrinsic learning (Chapter 2, 2.3.3) of generic primitive elements used in different tasks and platforms/agents.
- **Noise Tolerance** through the use of stable attractor dynamics as the base for motor control (Chapter 2, 2.3.1) and perceptual memories (Chapter 3, 3.3), as well as modeling stochasticity in episodic, stereotypical perceptual input (Chapter 3, 3.3);
- **Resource Efficiency** as a result of primitive sharing, achieving notable dimensionality reduction in motor and perceptual representation structures, and modular learning processes; moreover, employing time-invariant parametric representations allow continually archiving motor skills and episodic memories in an open-ended format.

Concurrently, other perspectives are highlighted, particularly when complementing the motor control strategy formalized in Chapter 2 with the episodic memory model proposed in Chapter 3, whereby this work provides a practical foundation - or a powerful complement - for **flexible reasoning and knowledge processing**, as highlighted by Tenorth & Beetz [2009], by addressing i) *action-centered representation* based on MP combinations, ii) *automated acquisition through observation* in the use of IL methods for skill learning and memory consolidation, iii) *managing uncertainty* by encoding stochasticity and variability observed in sensor patterns, and exploiting stable dynamical systems for online robustness, and iv) *fast inference* achieved with low-dimensional dynamic representations; and, similarly, for **robust representation of motor memory**, namely through the features identified by Žabkar & Leonardis [2016] i.e. i) being *compact* i.e. resource-efficient, and ii) *reusing information across different domains* i.e. exploiting task (and platform/agent) similarity; ultimately, iii) *fast online learning* as a result of reduced dimensionality and versatile learning strategies, and iv) *association of low-level motor commands with high-level symbolic abstract knowledge*, by means of high-level $\mathcal{S} \leftrightarrow \mathcal{M}$ coupling, can be implemented. From the viewpoint of human cognition and high-level cognitive architectures, the motor libraries introduced in Chapter 2 can be viewed as *procedural* memories, while knowledge stored with the memory model proposed in Chapter 3 represent *episodic* memories, which, in *egocentric* implementations, encode stereotypical *proprioceptive* information. In general, associative semantics between both represent a practical, data-driven, modular representation of implicit, long-term, grounded Associative Skill Memories (ASMs).

Overall, the most relevant takeaway of the present thesis is how primitive-based approaches, illustrated in the tools and models proposed herein, can contribute to and complement existing cognitive architectures at different levels of abstraction (semantic, emergent and hybrid) and, in the long term, the creation of truly autonomous L2 machines, able to coexist, interact and grow safely and dynamically with and around humans.

Table 25: Overview of hallmark features in Lifelong Learning (L2) systems [Kudithipudi et al., 2022] and where each is addressed in this work; task-agnostic learning, while not directly addressed, is achievable with the outcomes of Chapter 2 and Chapter 3.

	Chapter 2	Chapter 3	Chapter 4
Transfer and Adaptation	•		•
Exploit Task Similarity	•		•
Overcoming Catastrophic Forgetting		•	
Task-agnostic Learning	•	•	
Noise Tolerance	•	•	
Resource Efficiency	•	•	•

5.2 Outcomes and Contributions

In line with what was initially presented in Chapter 1 (1.2), relevant scientific contributions of the present thesis are reflected in the form of different theoretical frameworks, algorithms, practical applications and resulting experimental datasets.

Algorithms and Theoretical Models

A generic, multi-level and modular primitive-based control structure, employed for motor control (Chapter 2, 2.3) and perceptual modeling (Chapter 3, 3.3) incorporate and different scientific concepts explored in previous works; Exploiting this structure, i) a generalized and versatile IL strategy (specifically as a form of BC), employed for motor skill learning and transfer, as well as offline memory consolidation, and ii) different methods for perceptual context tracking (Algorithm 3), search (Algorithm 4) and bootstrap (Algorithm 5) were developed. Practical implementations, described in Appendix A, are provided as generic, extensible C++ libraries, in addition to high-level Python bindings, front-end packages for the ROS ecosystem, and extensive documentation of their public API. Lastly, iii) multiple time- and frequency-domain fatigue indicators constitute a valuable foundation for the design of robust and efficient fatigue management solutions.

Applications

The adoption of existing methodologies and frameworks in novel applications and/or engineering problems was illustrated in the generalized use of DMPs for i) robotic bipedal locomotion, notably not common given the complexity and dynamic nature of the movement (Chapter 2, 2.4); ii) efficient data-storage, exploited in episodic memory modeling (Chapter 3, 3.3); and iii) detection of slip-like perturbations in human walking (Chapter 3, 3.4). In addition, a novel BC strategy for learning bipedal walking from preset dynamical trajectories, as opposed to the common kinesthetic teaching approach often adopted in manipulation tasks (Chapter 2, 2.4.3). And, lastly, in the MS-based analysis of dynamic, low-force movements (Chapter 4, 4.3.3), while research mostly focus on isometric conditions.

Datasets

An extensive i) dataset including right-arm kinematic and electromyographic patterns, for a total of 17 participants executing 1) multiple short, uni-articular, upper-limb motions and 2) composite movements until self-perceived exertion or kinematic non-conformity (Chapter 4, 4.3.3), was recorded during the study of fatigue-induced alterations in MS recruitment (Chapter 4), is publicly available ([hosted in Zenodo](#)). The dataset includes segmented² raw³ recordings, as well as self-perceived fatigue levels at the end of each exercise reported by trial participants, according to Borg's Perceived Exertion Scale [Borg, 1998].

5.2.1 Publications

From a dissemination perspective, the research conducted within the context of this doctoral thesis led to the following publications in peer-reviewed scientific journals:

- *Using dynamical systems to accurately detect gait context and perturbations*, André, J., Ribeiro, N., Santos, C., [Journal...](#), 2025 [...]
- *Exploring the impact of fatigue in muscle synergies*, André, J., Cerqueira, S., Gizzi, L., Santos, C., *Journal of Neural Engineering*, 2025, [...]

5.3 Limitations and Future Work

Broadly, given the focus on bioinspired (Chapter 2) and biological (Chapter 3 and 4) movement throughout this work, the different categories of primitive elements studied - MPs, PPs and MSs - can be viewed as different manifestations of the same underlying core elements *i.e.* functional MPs within biological MSK systems; while this does not invalidate the results, discussion and conclusions presented herein, it warrants further research into a wider range of practical applications, including comprehensive validation in different motor skills, tasks and/or platforms and systems, ideally in a real-time context. Particularly, a unified approach to grounded skill learning and representation, through the associative combination of motor and perceptual libraries and contextual perceptuomotor coupling *e.g.* with RL techniques, is a natural next step in the research work described in this thesis.

Regarding robotic control (Chapter 2), the i) validation of the proposed strategy outside simulation *i.e.* in real platforms and in more challenging terrains, would further attest to the practical appeal of MP-based control; furthermore, motor skill learning can be improved with the ii) generalization of learned behaviors with contextual RL *e.g.* using perceptual context $\mathcal{C}_t$ (Chapter 3, 3.3.2) as input; and in the specific case of robotic walking, iii) conventional and/or cutting-edge DRL techniques would be able to autonomously improve performance *i.e.* stability, velocity and/or energetic footprint in real time;

The perceptual memory model (Chapter 3), in its turn, was only validated with kinematic input, despite its sensor-agnostic nature. As such, iv) validation with multimodal perceptual inputs and modalities

²Following the kinematic segmentation procedure (Chapter 4, 4.3.4.2).

³Acquired with the experimental setup (4.3.2), including post-processing of body kinematics and preliminarily filtered EMG.

e.g. tactile, auditory and visual perception - beyond motion capture where the overlap with kinematic MPs is plausible - is necessary, with particular focus on the dynamic nature of contextual identification and the *tracking* and *bootstrapping* mechanisms in Algorithm 3 and Algorithm 5, respectively. Moreover, its theoretical formulation can be extended to v) incorporate scale-invariance through existing DMP extensions (Chapter 2, 2.3.1), further generalizing episodic memories and increasing data efficiency.

Lastly, detection of muscle fatigue (Chapter 4) can be improved with vi) dynamical models that can accurately describe fatigue onset over time, predicting physical impairment and thus allowing long-term fatigue management, in line with similar generalized solutions in the literature. In parallel, the vii) validation of identified fatigue metrics in different motions and limbs would further reinforce the validity of the results.

In summary, this thesis has advanced the understanding and application of primitive theory in distinct areas of research, with a particular emphasis on robotic control, but whose generality ultimately allows application beyond robotics and automated systems, as was demonstrated in the validation procedures presented. It bridges theoretical constructs to practical implementations, providing tools for dynamic and adaptable motor control and behavioral monitoring. Going forward, it is hoped that these findings will inspire continued innovation in both robotics and interdisciplinary studies of human and artificial movement, which will enhance the capability of robots to operate in shared human-machine environments.

Practical Implementation

This chapter provides a brief high-level summary of the practical implementation developed in this thesis, supplementing the studies described in Chapter 2 and 3, comprised of generic C++ libraries and a specialized Robot Operating System (ROS) environment used to control locomotion in simulated humanoid platforms.

A.1 Core Libraries

As illustrated in Figure A.1, the core implementation comprises a set of 3 generic libraries, organized in a modular structure:

- **mplibrary** Base implementations of different dynamical systems, discrete and rhythmic multidimensional DMPs, and compound structures for convenience *i.e.* skill and primitive libraries;
- **mplearn** Generic learning-oriented implementations to parametrize DMP classes in **mplibrary**;
- **mpmemory** Higher-level extensions to **mplibrary** providing episodic memory classes and context identification algorithms.

Development followed a generic OOP paradigm, specifically through polymorphic derivations of a minimal set of base abstract interfaces. For computational efficiency, a *pass-by-reference* approach was adopted, and all implementations are thread-safe. Moreover, for additional convenience, high-level bindings for python are also provided using *pybind11*¹ [REF].

A.1.1 mplibrary

As illustrated in Figure A.2, different dynamical systems are implemented as derivations of the *DynamicalSystem* interface, describing generic time-scalable first- and second-order dynamics. By default,

¹v. 2.13.6.




Figure A.1: Conceptual flowchart of the C++ implementation developed in this work; core classes and definition as defined on `mplibrary`, with higher-level extensions provided by `mplearn` and `mpmemory`; core libraries written in C++, with high-level Python bindings provided through `pybind11`.




Figure A.2: Illustrative UML diagram describing OOP implementation of different dynamical systems.




Figure A.3: Illustrative UML diagram describing OOP implementation of different RBFs.

DynamicalSystem operates on a step-by-step basis, wherein states are synchronously updated - using Euler's method - according to a given time step. Among other general-purpose systems, *QuadraticDecay* and *LinearPacer* classes implement canonical systems for discrete and rhythmic DMP (6) used in the classical formulation [Ijspeert et al., 2013, 2003]; in addition, the *SpringDamper* class provides second-order attractor dynamics in DMP transformation systems (1), with a separate *Damping* structure encoding the system damping.

Similarly, different Radial Basis Functions (RBFs) are implemented through a minimal abstract *BasisFunction* interface, employing a generic *Parameters* class aggregating center and width (Figure A.3). Of note are the *Gaussian* and *VonMises* derivations, describing the DMP kernels in the discrete and rhythmic cases (5), respectively.

At a higher level, in a hierarchical, compositional structure, the *Primitive*, *Skill* and *DMP* classes provide the core implementations in *mplibrary* (Figure A.4) as unique elements in motor libraries (2.3.2): *Primitive* as a generic collection of RBFs (12), *Skill* as a combination of different *Primitive* instances (20), and *DMP* as a specialization of the *Multi* class template implementing a container of multiple *SpringDamper* instances, while associated with a specific *Skill* instance. These classes derive from *UniqueType*, a minimal interface generating Universally Unique Identifiers (UUIDs) on construction as well as allowing custom text-based labels for convenience.

Convenience containers are also provided (Figure A.5) as both a generic *Archive* template aggregating multiple same-type *UniqueType* instances *i.e.* *Primitive*, *Skill* and *DMP*, and a top-level multimodal *Library* class holding specialized *Archive* instances for each unique type, as *primitives*, *skills*, *policies* members, respectively.

Lastly, serialization of these structures through their parametrization to/from a custom Extensible Markup Language (XML) format (*.mpx*) is provided, allowing instances of *Library*, *Archive*, and different *UniqueType* derivations to be stored in portable, human-readable local files.




Figure A.4: Illustrative UML diagram describing OOP implementation of modular unique elements in motor libraries.




Figure A.5: Illustrative UML diagram describing OOP implementation of high-level containers for unique elements.




Figure A.6: Illustrative UML diagram describing OOP implementation of learning-oriented structures.

A.1.2 mplearn

Built on top of `mplibrary` structures, `mplearn` provides a uniform interface for optimization and learning implementations, by defining virtual abstract *BasicPolicyLearner* and *BasicPolicyPerformer* interfaces (Figure A.6), from which different learning techniques and problem-specific cost functions can be derived, respectively. In the specific case of the works described in Chapter 2 and 3, a *CMAESLearner* was derived from *BasicPolicyLearner*, wrapping around `libcmaes`'s [Beniz, nd] implementation of Covariance Matrix Adaptation - Evolutionary Strategy (CMA-ES). Concurrently, *BasicDMPRollout* extended *BasicPolicyPerformer* to implement customizable parametrization of a multidimensional DMP system, according to a specific *Parameters* subclass. *DMPReferenceMatch*, in its turn, implemented Root Mean Square Error (RMSE) computation between a sequence of DMP states and a reference trajectory, according to (23). Notably, `mplearn` allows compounding multiple *PolicyPerformer* instances for multi-objective, chained implementations, similar to what is proposed in (25), as depicted in Figure A.7. Such a generic interface allows implementation of different learning paradigms, such as BBO and RL, while also being inherently extensible, with ongoing development in other robotic problems, whereby new *BasicPolicyLearner* and *BasicPolicyPerformer* derived classes can be integrated.

A.1.3 mpmemory

Also extending `mplibrary`, `mpmemory` defines an additional *UniqueType* derivation in a *Footprint* class coupling two *dmp* instances (66) and implementing the memory matching through (69) (Figure A.8). Separate *Context* and *Tracker* classes provide public and static member functions implementing the *update* (i.e. *tracking* Algorithm 3) operation as well as the *search* (Algorithm 4) and *bootstrap* (Algorithm 5)




Figure A.7: Conceptual flowchart illustrating possible learning strategies implemented in `mplearn`; learner types can be attached to single (a) or multiple (b) performer instances for multi-objective optimization.




Figure A.8: Illustrative UML diagram describing OOP implementation of generic perceptual memories.

algorithms (Figure A.9). *Library* is extended to *ExtLibrary* to also include a *memories* member aggregating different perceptual memories. Retroactively compatible with the structures in `mplibrary`, namely the learning approach provided by `mplearn`, used for memory consolidation (Chapter 3, 3.4.4), and the generic serialization implementations for convenience.

A.2 ROS Environment

A ROS² workspace was constructed for the walking simulations described in Chapter 2, comprised of different C++ packages exploiting `mplibrary` in the control of humanoid platforms, as illustrated in

²v. Noetic Ninjemys.




Figure A.9: Illustrative UML diagram describing OOP implementation of perceptual context identification.




Figure A.10: Conceptual flowchart of the ROS C++ packages developed in this work; `roswalker` employs the generic implementations in `mplibrary` to load and/or generate walking command trajectories, while `roscontrol_utils` and `roshumanoid` allow abstraction from platform-specific models and drivers.

Figure A.10:

- **`roscontrol_utils`** Generic tools to simplify and streamline the development of control-related ROS packages, wrapping ROS topic and service interfaces into a high-level code API;
- **`roshumanoid`** Package built on top of `roscontrol_utils`, establishing a uniform humanoid topic and service interface, allowing a *plug-and-play* approach to the control of different humanoid - i.e. morphologically similar - platforms; configuration presets for the OP and the NAO are included;
- **`roswalker`** Locomotion-oriented, platform-agnostic package employing `mplibrary`, that loads and publishes high-level control patterns to leg actuators, using `roshumanoid` to interface with different robot models; CPG and DMP controllers are included;
- OP simulation model and ROS drivers (low-level PID implementation):
 - `darwin_description`
 - `darwin_control`
- NAO simulation model and ROS drivers (low-level PID implementation):
 - `nao_description` (as part of `nao_robot`)
 - `nao_control` (as part of `nao_virtual`)
 - `nao_meshes`

A.2.1 `roscontrol_utils`

Commonly, joint control in ROS is implemented resorting to the `controller_manager` package that handles low-level PID control and advertises topics for control commands and joint state; similarly, sensor plugins (in simulation) or hardware drivers publish data to sensor-specific topics (at a preset sampling rate). For general convenience, `roscontrol_utils` wraps ROS topic interface through *Actuator* and




Figure A.11: Illustrative UML diagram describing OOP implementation of high-level general-purpose ROS control utilities.

Sensor classes, aggregated in *MultiActuator* and *MultiSensor* container-like types (Figure A.11), providing a minimal, high-level code API to streamline the design of generic control strategies, abstracted from platform-specific implementations.

A.2.2 roshumanoid

Making use of the generic types provides by `roscontrol_utils`, `roshumanoid` introduces a *Humanoid* class (Figure A.12) derived from *MultiActuator* and coupled with a *MultiSensor* instance. In order to abstract control from robot morphology, a general uniform structure is assumed across different humanoid platforms (Figure A.13) as a collection of biologically-plausible spherical joints. Up to 3 DoFs can be defined at each joint (*i.e.* description of DoF-specific ROS topic), and platform-specific mappings are provided in local configuration files (in the case of the OP and the NAO), which are loaded into the ROS parameter server at runtime. Recursively, nested *LimbInterface* instances allow independent limb control. Additionally, *BasicHumanoidControl* provides a base interface for generalized control of *Humanoid* instances.




Figure A.12: Illustrative UML diagram describing OOP implementation of a general *Humanoid* class for control of humanoid platforms, as well as a minimal control interface.




Figure A.13: Default humanoid structure expected by the *Humanoid* class; different humanoid configurations can define up to 3 DoFs per articulated joint, as well as multiple sensor inputs.




Figure A.14: Illustrative UML diagram describing OOP implementation of different force-base feedback mechanisms.




Figure A.15: Illustrative UML diagram describing OOP implementation of CPG dynamics.

A.2.3 roswalker

Given the *Humanoid* interface provided by *roshumanoid*, *roswalker* defines a locomotion-oriented *DynamicalWalker* class derived from *BasicHumanoidControl*, further extended in *CPGWalker* and *DMPWalker* in the case of CPG-based (2.4.1) and DMP-based (2.4.2) walking, respectively. Helper types *Lateral* and *Descriptor* describe, respectively, the contralateral and hierarchical³ organization of a humanoid robot. In the particular case of CPG control, the *BipedGaitCPG* class was developed using the base types in *mplib*. Both walking controllers prescribe open-loop trajectories to the *left_leg* and *right_leg* limb interfaces, according to the methodologies described in Chapter 2. Closed-loop control is achieved though a *WalkingPacer* class extending *DynamicalSystem* implementing the coupled phase system with force feedback (47 - 48) used in both *CPGWalker* and *DMPWalker*. Simultaneously, the goal states of the ankle *pitch* and *roll* joints are adjusted dynamically by the *PitchCorrection* and *RollCorrection* classes implementing the ankle feedback mechanisms (53 - 54).

³Through the recursive design of *Descriptor* i.e. each instance as a collection of multiple *Descriptor* instances allows extension beyond humanoid and/or anthropomorphic organizations.




Figure A.16: Illustrative UML diagram describing OOP implementation of CPG and DMP walking controllers.

Humanoid Platforms

B.1 Robotis OP

The Robotis OP humanoid, previously Dynamic Anthropomorphic Robot with Intelligence - Open Platform (DARwIn-OP), is mobile anthropomorphic robot manufactured by Robotis, with a total of 20 DoFs (Figure B.1): 6 on each leg, 3 on each arm, and 2 on the neck, actuated by DC servos (*Dynamixel MX28*, with a 193:1 reduction ratio, TTL interface, and encoder resolution of 12 bits $\sim 0.1^\circ$), which are controlled by an *ARM Cortex M3* (CM-730) board. In addition, a 3-axis Inertial Measurement Unit (IMU) is installed on the central frame, and 4 FSRs sensors (0.1-10 N) in each foot. It is 46 cm tall and weighs nearly 3 kg (Table B.1), with a detailed overview of its physical dimensions illustrated in Figure B.1. Additionally, it features an onboard computer with an Intel Atom Z530 CPU and a 4 GB SSD, which facilitates standalone and/or autonomous operation.

B.2 Softbank NAO

The NAO robot, developed by Aldebaran (currently part of Softbank Robotics), is a 58 cm tall humanoid platform with 25 DoFs (Figure B.2): 6 on each leg - with mechanically coupled hip yaw DoFs - 5 on each arm, and 2 on the neck, actuated by brushless DC motors. It features a variety of sensors including 2 cameras, 4 microphones, an IMU, magnetic rotary encoders with 12 bit precision $\sim 0.1^\circ$, and 8 foot FSRs (0-25 N). Similarly to the OP, low-level PID control is executed on an *ARM Cortex M3* STM32 board, and the robot is equipped with an onboard computer (Intel Atom Z530 CPU, 1 GB RAM, and a total of 2 GB internal storage, extensible to 8 GB through a micro SDHC card slot) enabling autonomous movement.

	OP	NAO
Mass (kg)	2.9	5.3
Height (m)	0.4597	0.5791
L (m)	0.2195	0.2481

Table B.1: Mass, height and leg length (L) of both humanoid platforms used in this work.




Figure B.1: Physical overview of the Robotis OP [ROBOTIS, 2019].




Figure B.2: Physical overview of the Softbank NAO V5 [Robotics, 2019].

Stereotypical Kinematics in Human Locomotion

C.1 Kinematic References

Figures C.1 to C.36 depict the lower-limb kinematic references $\bar{\mathbf{x}}^{\text{REF}}$ and associated standard deviation $\mathbf{s}^{\text{REF}}$, for all participants, velocities and inclines, using the methodology described in 3.4.3.3 and 3.4.4.

C.2 Shared Perceptual Primitives

The set of $P = 5$ PPs identified in all participants, for the different velocities and inclines, is presented in Figures C.37 to C.72.




Figure C.1: Lower-limb kinematic references $\bar{\mathbf{x}}^{\text{REF}}$ and $\mathbf{s}^{\text{REF}}$ obtained with **S1**, for 5.4 km/h and 0° incline.




Figure C.2: Lower-limb kinematic references $\bar{\mathbf{x}}^{\text{REF}}$ and $\mathbf{s}^{\text{REF}}$ obtained with **S3**, for 5.4 km/h and 0° incline.




Figure C.3: Lower-limb kinematic references $\bar{\mathbf{x}}^{\text{REF}}$ and $\mathbf{s}^{\text{REF}}$ obtained with **S4**, for 5.4 km/h and 0° incline.




Figure C.4: Lower-limb kinematic references $\bar{\mathbf{x}}^{\text{REF}}$ and $\mathbf{s}^{\text{REF}}$ obtained with **S5**, for 5.4 km/h and 0° incline.




Figure C.5: Lower-limb kinematic references $\bar{\mathbf{x}}^{\text{REF}}$ and $\mathbf{s}^{\text{REF}}$ obtained with **S7**, for 5.4 km/h and 0° incline.




Figure C.6: Lower-limb kinematic references $\bar{\mathbf{x}}^{\text{REF}}$ and $\mathbf{s}^{\text{REF}}$ obtained with **S8**, for 5.4 km/h and 0° incline.




Figure C.7: Lower-limb kinematic references $\bar{\mathbf{x}}^{\text{REF}}$ and $\mathbf{s}^{\text{REF}}$ obtained with **S9**, for 5.4 km/h and 0° incline.




Figure C.8: Lower-limb kinematic references $\bar{\mathbf{x}}^{\text{REF}}$ and $\mathbf{s}^{\text{REF}}$ obtained with **S10**, for 5.4 km/h and 0° incline.




Figure C.9: Lower-limb kinematic references $\bar{\mathbf{x}}^{\text{REF}}$ and $\mathbf{s}^{\text{REF}}$ obtained with **S11**, for 5.4 km/h and 0° incline.




Figure C.10: Lower-limb kinematic references $\bar{\mathbf{x}}^{\text{REF}}$ and $\mathbf{s}^{\text{REF}}$ obtained with **S1**, for 5.4 km/h and 10° incline.




Figure C.11: Lower-limb kinematic references $\bar{\mathbf{x}}^{\text{REF}}$ and $\mathbf{s}^{\text{REF}}$ obtained with **S3**, for 5.4 km/h and 10° incline.




Figure C.12: Lower-limb kinematic references $\bar{\mathbf{x}}^{\text{REF}}$ and $\mathbf{s}^{\text{REF}}$ obtained with **S4**, for 5.4 km/h and 10° incline.




Figure C.13: Lower-limb kinematic references $\bar{\mathbf{x}}^{\text{REF}}$ and $\mathbf{s}^{\text{REF}}$ obtained with **S5**, for 5.4 km/h and 10° incline.




Figure C.14: Lower-limb kinematic references $\bar{\mathbf{x}}^{\text{REF}}$ and $\mathbf{s}^{\text{REF}}$ obtained with **S7**, for 5.4 km/h and 10° incline.




Figure C.15: Lower-limb kinematic references $\bar{\mathbf{x}}^{\text{REF}}$ and $\mathbf{s}^{\text{REF}}$ obtained with **S8**, for 5.4 km/h and 10° incline.




Figure C.16: Lower-limb kinematic references $\bar{\mathbf{x}}^{\text{REF}}$ and $\mathbf{s}^{\text{REF}}$ obtained with **S9**, for 5.4 km/h and 10° incline.




Figure C.17: Lower-limb kinematic references $\bar{\mathbf{x}}^{\text{REF}}$ and $\mathbf{s}^{\text{REF}}$ obtained with **S10**, for 5.4 km/h and 10° incline.




Figure C.18: Lower-limb kinematic references $\bar{\mathbf{x}}^{\text{REF}}$ and $\mathbf{s}^{\text{REF}}$ obtained with **S11**, for 5.4 km/h and 10° incline.




Figure C.19: Lower-limb kinematic references $\bar{\mathbf{x}}^{\text{REF}}$ and $\mathbf{s}^{\text{REF}}$ obtained with **S1**, for 5.4 km/h and 0° incline.




Figure C.20: Lower-limb kinematic references $\bar{\mathbf{x}}^{\text{REF}}$ and $\mathbf{s}^{\text{REF}}$ obtained with **S3**, for 5.4 km/h and 0° incline.




Figure C.21: Lower-limb kinematic references $\bar{\mathbf{x}}^{\text{REF}}$ and $\mathbf{s}^{\text{REF}}$ obtained with **S4**, for 5.4 km/h and 0° incline.




Figure C.22: Lower-limb kinematic references $\bar{\mathbf{x}}^{\text{REF}}$ and $\mathbf{s}^{\text{REF}}$ obtained with **S5**, for 5.4 km/h and 0° incline.




Figure C.23: Lower-limb kinematic references $\bar{\mathbf{x}}^{\text{REF}}$ and $\mathbf{s}^{\text{REF}}$ obtained with **S7**, for 5.4 km/h and 0° incline.




Figure C.24: Lower-limb kinematic references $\bar{\mathbf{x}}^{\text{REF}}$ and $\mathbf{s}^{\text{REF}}$ obtained with **S8**, for 5.4 km/h and 0° incline.




Figure C.25: Lower-limb kinematic references $\bar{\mathbf{x}}^{\text{REF}}$ and $\mathbf{s}^{\text{REF}}$ obtained with **S9**, for 5.4 km/h and 0° incline.




Figure C.26: Lower-limb kinematic references $\bar{\mathbf{x}}^{\text{REF}}$ and $\mathbf{s}^{\text{REF}}$ obtained with **S10**, for 5.4 km/h and 0° incline.




Figure C.27: Lower-limb kinematic references $\bar{\mathbf{x}}^{\text{REF}}$ and $\mathbf{s}^{\text{REF}}$ obtained with **S11**, for 5.4 km/h and 0° incline.




Figure C.28: Lower-limb kinematic references $\bar{\mathbf{x}}^{\text{REF}}$ and $\mathbf{s}^{\text{REF}}$ obtained with **S1**, for 5.4 km/h and 10° incline.




Figure C.29: Lower-limb kinematic references $\bar{\mathbf{x}}^{\text{REF}}$ and $\mathbf{s}^{\text{REF}}$ obtained with **S3**, for 5.4 km/h and 10° incline.




Figure C.30: Lower-limb kinematic references $\bar{\mathbf{x}}^{\text{REF}}$ and $\mathbf{s}^{\text{REF}}$ obtained with **S4**, for 5.4 km/h and 10° incline.




Figure C.31: Lower-limb kinematic references $\bar{\mathbf{x}}^{\text{REF}}$ and $\mathbf{s}^{\text{REF}}$ obtained with **S5**, for 5.4 km/h and 10° incline.




Figure C.32: Lower-limb kinematic references $\bar{\mathbf{x}}^{\text{REF}}$ and $\mathbf{s}^{\text{REF}}$ obtained with **S7**, for 5.4 km/h and 10° incline.




Figure C.33: Lower-limb kinematic references $\bar{\mathbf{x}}^{\text{REF}}$ and $\mathbf{s}^{\text{REF}}$ obtained with **S8**, for 5.4 km/h and 10° incline.




Figure C.34: Lower-limb kinematic references $\bar{\mathbf{x}}^{\text{REF}}$ and $\mathbf{s}^{\text{REF}}$ obtained with **S9**, for 5.4 km/h and 10° incline.




Figure C.35: Lower-limb kinematic references $\bar{\mathbf{x}}^{\text{REF}}$ and $\mathbf{s}^{\text{REF}}$ obtained with **S10**, for 5.4 km/h and 10° incline.




Figure C.36: Lower-limb kinematic references $\bar{\mathbf{x}}^{\text{REF}}$ and $\mathbf{s}^{\text{REF}}$ obtained with **S11**, for 5.4 km/h and 10° incline.







Figure C.37: $P = 5$ PPs shared across all $K = 18$ lower-limb DoFs of **S1**, for 5.4 km/h and 0° incline.


 Figure C.38: $P = 5$ PPs shared across all $K = 18$ lower-limb DoFs of **S3**, for 5.4 km/h and 0° incline.

 Figure C.39: $P = 5$ PPs shared across all $K = 18$ lower-limb DoFs of **S4**, for 5.4 km/h and 0° incline.

 Figure C.40: $P = 5$ PPs shared across all $K = 18$ lower-limb DoFs of **S5**, for 5.4 km/h and 0° incline.




Figure C.41: $P = 5$ PPs shared across all $K = 18$ lower-limb DoFs of **S7**, for 5.4 km/h and 0° incline.




Figure C.42: $P = 5$ PPs shared across all $K = 18$ lower-limb DoFs of **S8**, for 5.4 km/h and 0° incline.




Figure C.43: $P = 5$ PPs shared across all $K = 18$ lower-limb DoFs of **S9**, for 5.4 km/h and 0° incline.




Figure C.44: $P = 5$ PPs shared across all $K = 18$ lower-limb DoFs of **S10**, for 5.4 km/h and 0° incline.




Figure C.45: $P = 5$ PPs shared across all $K = 18$ lower-limb DoFs of **S11**, for 5.4 km/h and 0° incline.




Figure C.46: $P = 5$ PPs shared across all $K = 18$ lower-limb DoFs of **S1**, for 5.4 km/h and 10° incline.




Figure C.47: $P = 5$ PPs shared across all $K = 18$ lower-limb DoFs of **S3**, for 5.4 km/h and 10° incline.




Figure C.48: $P = 5$ PPs shared across all $K = 18$ lower-limb DoFs of **S4**, for 5.4 km/h and 10° incline.







Figure C.49: $P = 5$ PPs shared across all $K = 18$ lower-limb DoFs of **S5**, for 5.4 km/h and 10° incline.


 Figure C.50: $P = 5$ PPs shared across all $K = 18$ lower-limb DoFs of **S7**, for 5.4 km/h and 10° incline.

 Figure C.51: $P = 5$ PPs shared across all $K = 18$ lower-limb DoFs of **S8**, for 5.4 km/h and 10° incline.

 Figure C.52: $P = 5$ PPs shared across all $K = 18$ lower-limb DoFs of **S9**, for 5.4 km/h and 10° incline.




Figure C.53: $P = 5$ PPs shared across all $K = 18$ lower-limb DoFs of **S10**, for 5.4 km/h and 10° incline.




Figure C.54: $P = 5$ PPs shared across all $K = 18$ lower-limb DoFs of **S11**, for 5.4 km/h and 10° incline.







Figure C.55: $P = 5$ PPs shared across all $K = 54$ lower-limb DoFs of **S1**, for 5.4 km/h and 0° incline.


 Figure C.56: $P = 5$ PPs shared across all $K = 54$ lower-limb DoFs of **S3**, for 5.4 km/h and 0° incline.

 Figure C.57: $P = 5$ PPs shared across all $K = 54$ lower-limb DoFs of **S4**, for 5.4 km/h and 0° incline.

 Figure C.58: $P = 5$ PPs shared across all $K = 54$ lower-limb DoFs of **S5**, for 5.4 km/h and 0° incline.




Figure C.59: $P = 5$ PPs shared across all $K = 54$ lower-limb DoFs of **S7**, for 5.4 km/h and 0° incline.




Figure C.60: $P = 5$ PPs shared across all $K = 54$ lower-limb DoFs of **S8**, for 5.4 km/h and 0° incline.




Figure C.61: $P = 5$ PPs shared across all $K = 54$ lower-limb DoFs of **S9**, for 5.4 km/h and 0° incline.




Figure C.62: $P = 5$ PPs shared across all $K = 54$ lower-limb DoFs of **S10**, for 5.4 km/h and 0° incline.




Figure C.63: $P = 5$ PPs shared across all $K = 54$ lower-limb DoFs of **S11**, for 5.4 km/h and 0° incline.




Figure C.64: $P = 5$ PPs shared across all $K = 54$ lower-limb DoFs of **S1**, for 5.4 km/h and 10° incline.




Figure C.65: $P = 5$ PPs shared across all $K = 54$ lower-limb DoFs of **S3**, for 5.4 km/h and 10° incline.




Figure C.66: $P = 5$ PPs shared across all $K = 54$ lower-limb DoFs of **S4**, for 5.4 km/h and 10° incline.




Figure C.67: $P = 5$ PPs shared across all $K = 54$ lower-limb DoFs of **S5**, for 5.4 km/h and 10° incline.




Figure C.68: $P = 5$ PPs shared across all $K = 54$ lower-limb DoFs of **S7**, for 5.4 km/h and 10° incline.




Figure C.69: $P = 5$ PPs shared across all $K = 54$ lower-limb DoFs of **S8**, for 5.4 km/h and 10° incline.




Figure C.70: $P = 5$ PPs shared across all $K = 54$ lower-limb DoFs of **S9**, for 5.4 km/h and 10° incline.




Figure C.71: $P = 5$ PPs shared across all $K = 54$ lower-limb DoFs of **S10**, for 5.4 km/h and 10° incline.




Figure C.72: $P = 5$ PPs shared across all $K = 54$ lower-limb DoFs of **S11**, for 5.4 km/h and 10° incline.

Fatigue-related Alterations in Muscle Synergies

This chapter presents the extensive set of results obtained during the study of how Muscle Synergies (MSs) are impacted by muscular fatigue (Chapter 4), namely as the sequential computation of conventional and MS-based metrics during **D** for all 17 trial participants.

D.1 Muscle Synergy Estimation

Figures D.1 to D.17 illustrate generic synergies $\mathbf{W}_{ABC}$ and their recruitment $\mathbf{H}_{ABC}$ during **P1**, for all participants, using the methodology described in 4.3.4.4.

D.2 Classical Analysis

The evolution of the conventional metrics (Chapter 4, 4.2.3) for all participants is illustrated in Figures D.18 to Figure D.34.




Figure D.1: Generic synergies $\mathbf{W}_{ABC}$ and activation profile $\mathbf{H}_{ABC}$ extracted from **P1** with **S1**.




Figure D.2: Generic synergies $\mathbf{W}_{ABC}$ and activation profile $\mathbf{H}_{ABC}$ extracted from **P1** with **S2**.




Figure D.3: Generic synergies $\mathbf{W}_{ABC}$ and activation profile $\mathbf{H}_{ABC}$ extracted from **P1** with **S3**.




Figure D.4: Generic synergies W_{ABC} and activation profile H_{ABC} extracted from **P1** with **S4**.




Figure D.5: Generic synergies W_{ABC} and activation profile H_{ABC} extracted from **P1** with **S5**.




Figure D.6: Generic synergies W_{ABC} and activation profile H_{ABC} extracted from **P1** with **S6**.




Figure D.7: Generic synergies $\mathbf{W}_{ABC}$ and activation profile $\mathbf{H}_{ABC}$ extracted from **P1** with **S7**.




Figure D.8: Generic synergies $\mathbf{W}_{ABC}$ and activation profile $\mathbf{H}_{ABC}$ extracted from **P1** with **S8**.




Figure D.9: Generic synergies $\mathbf{W}_{ABC}$ and activation profile $\mathbf{H}_{ABC}$ extracted from **P1** with **S9**.




Figure D.10: Generic synergies $\mathbf{W}_{ABC}$ and activation profile $\mathbf{H}_{ABC}$ extracted from **P1** with **S10**.




Figure D.11: Generic synergies $\mathbf{W}_{ABC}$ and activation profile $\mathbf{H}_{ABC}$ extracted from **P1** with **S11**.




Figure D.12: Generic synergies $\mathbf{W}_{ABC}$ and activation profile $\mathbf{H}_{ABC}$ extracted from **P1** with **S12**.




Figure D.13: Generic synergies $\mathbf{W}_{ABC}$ and activation profile $\mathbf{H}_{ABC}$ extracted from **P1** with **S13**.




Figure D.14: Generic synergies $\mathbf{W}_{ABC}$ and activation profile $\mathbf{H}_{ABC}$ extracted from **P1** with **S14**.




Figure D.15: Generic synergies $\mathbf{W}_{ABC}$ and activation profile $\mathbf{H}_{ABC}$ extracted from **P1** with **S15**.




Figure D.16: Generic synergies W_{ABC} and activation profile H_{ABC} extracted from **P1** with **S16**.




Figure D.17: Generic synergies W_{ABC} and activation profile H_{ABC} extracted from **P1** with **S17**.

S1




Figure D.18: Classical metrics measured for each individual repetition of **D** with **S1**.

S2





Figure D.19: Classical metrics measured for each individual repetition of **D** with **S2**.

S3Figure D.20: Classical metrics measured for each individual repetition of **D** with **S3**.

S4



Figure D.21: Classical metrics measured for each individual repetition of **D** with **S4**.

S5






Figure D.22: Classical metrics measured for each individual repetition of **D** with **S5**.



S7Figure D.24: Classical metrics measured for each individual repetition of **D** with **S7**.




Figure D.25: Classical metrics measured for each individual repetition of **D** with **S8**.

S9




Figure D.26: Classical metrics measured for each individual repetition of **D** with **S9**.

S10




Figure D.27: Classical metrics measured for each individual repetition of **D** with **S10**.

S11




Figure D.28: Classical metrics measured for each individual repetition of **D** with **S11**.

S12





Figure D.29: Classical metrics measured for each individual repetition of **D** with **S12**.

S13Figure D.30: Classical metrics measured for each individual repetition of **D** with **S13**.

S14




Figure D.31: Classical metrics measured for each individual repetition of **D** with **S14**.

S16




Figure D.33: Classical metrics measured for each individual repetition of **D** with **S16**.

D.3 Synergy Analysis

Figures D.35 to D.51 depict synergy recruitment $\mathbf{H}_{\mathbf{ABC}}$ during **P2** for all participants, with $N_{\mathbf{ABC}} = 4$. The resulting evolution of the synergy metrics (Chapter 4, 4.3.4.5) in each case is illustrated in Figures D.52 to Figure D.68.

S17




Figure D.34: Classical metrics measured for each individual repetition of **D** with **S17**.




Figure D.35: Recruitment of generic synergies $\mathbf{W}_{ABC}$ during **P2** with **S1**.




Figure D.36: Recruitment of generic synergies $\mathbf{W}_{ABC}$ during **P2** with **S2**.




Figure D.37: Recruitment of generic synergies $\mathbf{W}_{ABC}$ during **P2** with **S3**.




Figure D.38: Recruitment of generic synergies $\mathbf{W}_{ABC}$ during **P2** with **S4**.




Figure D.39: Recruitment of generic synergies $\mathbf{W}_{ABC}$ during **P2** with **S5**.




Figure D.40: Recruitment of generic synergies $\mathbf{W}_{ABC}$ during **P2** with **S6**.




Figure D.41: Recruitment of generic synergies $\mathbf{W}_{ABC}$ during **P2** with **S7**.




Figure D.42: Recruitment of generic synergies $\mathbf{W}_{ABC}$ during **P2** with **S8**.




Figure D.43: Recruitment of generic synergies $\mathbf{W}_{ABC}$ during **P2** with **S9**.




Figure D.44: Recruitment of generic synergies $\mathbf{W}_{ABC}$ during **P2** with **S10**.




Figure D.45: Recruitment of generic synergies $\mathbf{W}_{ABC}$ during **P2** with **S11**.




Figure D.46: Recruitment of generic synergies $\mathbf{W}_{ABC}$ during **P2** with **S12**.




Figure D.47: Recruitment of generic synergies $\mathbf{W}_{ABC}$ during **P2** with **S13**.




Figure D.48: Recruitment of generic synergies $\mathbf{W}_{ABC}$ during **P2** with **S14**.




Figure D.49: Recruitment of generic synergies $\mathbf{W}_{ABC}$ during **P2** with **S15**.




Figure D.50: Recruitment of generic synergies $\mathbf{W}_{ABC}$ during **P2** with **S16**.





Figure D.51: Recruitment of generic synergies $\mathbf{W}_{ABC}$ during **P2** with **S17**.

S1Figure D.52: Classical metrics measured for each individual repetition of **D** with **S1**.

S2





Figure D.53: Classical metrics measured for each individual repetition of **D** with **S2**.

S3Figure D.54: Classical metrics measured for each individual repetition of **D** with **S3**.

S4





Figure D.55: Classical metrics measured for each individual repetition of **D** with **S4**.

S5Figure D.56: Classical metrics measured for each individual repetition of **D** with **S5**.

S6






Figure D.57: Classical metrics measured for each individual repetition of **D** with **S6**.

S7Figure D.58: Classical metrics measured for each individual repetition of **D** with **S7**.

S8Figure D.59: Classical metrics measured for each individual repetition of **D** with **S8**.

S9




Figure D.60: Classical metrics measured for each individual repetition of **D** with **S9**.

S10





Figure D.61: Classical metrics measured for each individual repetition of **D** with **S10**.

S11Figure D.62: Classical metrics measured for each individual repetition of **D** with **S11**.

S12







Figure D.63: Classical metrics measured for each individual repetition of **D** with **S12**.

S13Figure D.64: Classical metrics measured for each individual repetition of **D** with **S13**.

S14Figure D.65: Classical metrics measured for each individual repetition of **D** with **S14**.

S15Figure D.66: Classical metrics measured for each individual repetition of **D** with **S15**.

S16




Figure D.67: Classical metrics measured for each individual repetition of **D** with **S16**.

S17




Figure D.68: Classical metrics measured for each individual repetition of **D** with **S17**.

Bibliography

- Ajoudani, A., Zanchettin, A. M., Ivaldi, S., Albu-Schäffer, A., Kosuge, K., & Khatib, O. (2018). Progress and prospects of the human–robot collaboration. *Autonomous Robots*, 42(5), 957–975.
- Akkaladevi, S. C., Plasch, M., Hofmann, M., & Pichler, A. (2021). Semantic knowledge based reasoning framework for human robot collaboration. *Procedia CIRP*, 97, 373–378.
- Akoglu, H. (2018). User's guide to correlation coefficients. *Turkish Journal of Emergency Medicine*, 18(3), 91–93.
- Al-Mulla, M. R., Sepulveda, F., & Colley, M. (2011). A review of non-invasive techniques to detect and predict localised muscle fatigue. *Sensors*, 11(4), 3545–3594.
- Alarie, S., Audet, C., Gheribi, A. E., Kokkolaras, M., & Le Digabel, S. (2021). Two decades of blackbox optimization applications. *EURO Journal on Computational Optimization*, 9, 100011.
- Alexander, R. M. (1991). Energy-Saving Mechanisms in Walking and Running. *Journal of Experimental Biology*, 160(1), 55–69.
- Ames, A. D., Cousineau, E. A., & Powell, M. J. (2012). Dynamically stable bipedal robotic walking with NAO via human-inspired hybrid zero dynamics. In *Proceedings of the 15th ACM international conference on Hybrid Systems: Computation and Control* (pp. 135–144). Beijing China: ACM.
- Ananthanarayanan, A., Azadi, M., & Kim, S. (2012). Towards a bio-inspired leg design for high-speed running. *Bioinspiration & Biomimetics*, 7(4), 046005. Publisher: IOP Publishing.
- Anderson, J. R. (1993). Production Systems and the ACT-R Theory. In *Rules of the Mind*. Psychology Press.
- André, J., Faria, B. M., Santos, C., & Reis, L. P. (2016a). A data mining approach to predict falls in humanoid robot locomotion. In *Robot 2015: Second Iberian Robotics Conference* (pp. 273–285).: Springer.
- André, J., Santos, C., & Costa, L. (2013). Path integral learning of multidimensional movement trajectories. In *AIP Conference Proceedings*, volume 1558 (pp. 1025–1028).: AIP.

- André, J., Santos, C., & Costa, L. (2016b). Skill memory in biped locomotion. *Journal of Intelligent & Robotic Systems*, 82(3-4), 379–397.
- André, J., Teixeira, C., Santos, C. P., & Costa, L. (2015). Adapting biped locomotion to sloped environments. *Journal of Intelligent & Robotic Systems*, 80(3-4), 625–640.
- André, J., Cerqueira, S., & of Minho, U. (2024). Kinematic and electromyographic recordings during dynamic, repetitive, low-force movements.
- André, J., Costa, L., & Santos, C. (2014). Sensor-based failure prediction in biped locomotion. In *2014 IEEE International Conference on Autonomous Robot Systems and Competitions (ICARSC)* (pp. 170–174). Espinho, Portugal: IEEE.
- Apgar, T., Clary, P., Green, K., Fern, A., & Hurst, J. (2018). Fast Online Trajectory Optimization for the Bipedal Robot Cassie. In *Robotics: Science and Systems XIV: Robotics: Science and Systems Foundation*.
- Aprigliano, F., Martelli, D., Tropea, P., Pasquini, G., Micera, S., & Monaco, V. (2017). Aging does not affect the intralimb coordination elicited by slip-like perturbation of different intensities. *Journal of neurophysiology*, 118(3), 1739–1748.
- Arena, P., Fortuna, L., Frasca, M., & Sicurella, G. (2004). An Adaptive, Self-Organizing Dynamical System for Hierarchical Control of Bio-Inspired Locomotion. *IEEE Transactions on Systems, Man and Cybernetics, Part B (Cybernetics)*, 34(4), 1823–1837.
- Aslin, R. N. & Smith, L. B. (1988). Perceptual development. *Annual review of psychology*, 39, 435–473.
- Atkeson, C. G., Moore, A. W., & Schaal, S. (1997). Locally Weighted Learning. *Artificial Intelligence Review*, 11(1/5), 11–73.
- Bahl, S., Mukadam, M., Gupta, A., & Pathak, D. (2020). Neural Dynamic Policies for End-to-End Sensorimotor Learning. arXiv:2012.02788.
- Bajaj, I., Arora, A., & Hasan, M. M. F. (2021). Black-Box Optimization: Methods and Applications. In P. M. Pardalos, V. Rasskazova, & M. N. Vrahatis (Eds.), *Black Box Optimization, Machine Learning, and No-Free Lunch Theorems*, volume 170 (pp. 35–65). Cham: Springer International Publishing. Series Title: Springer Optimization and Its Applications.
- Bamman, M. M., Ingram, S. G., Caruso, J. F., & Greenisen, M. C. (1997). Evaluation of surface electromyography during maximal voluntary contraction. *The Journal of Strength & Conditioning Research*, 11(2), 68–72.
- Barroso, F. A. O. (2016). *Combining neuromuscular and biomechanical features to assess sensorimotor impairments after spinal cord injury or stroke*. PhD, Universidade do Minho.

- Barsalou, L. W. (2016). On Staying Grounded and Avoiding Quixotic Dead Ends. *Psychonomic Bulletin & Review*, 23(4), 1122–1142.
- Barzilay, O. & Wolf, A. (2011). A fast implementation for EMG signal linear envelope computation. *Journal of Electromyography and Kinesiology*, 21(4), 678–682.
- Beetz, M., Mösenlechner, L., & Tenorth, M. (2010). CRAM — A Cognitive Robot Abstract Machine for everyday manipulation in human environments. In *2010 IEEE/RSJ International Conference on Intelligent Robots and Systems* (pp. 1012–1017).
- Bellman, R. (1957). *Dynamic programming*. Princeton, NJ: Princeton Univ. Pr.
- Beniz (n.d.). libcmaes. <https://github.com/beniz/libcmaes>. Accessed: 2024-10-08.
- Beranek, R., Karimi, M., & Ahmadi, M. (2022). A Behavior-Based Reinforcement Learning Approach to Control Walking Bipedal Robots Under Unknown Disturbances. *IEEE/ASME Transactions on Mechatronics*, 27(5), 2710–2720. Conference Name: IEEE/ASME Transactions on Mechatronics.
- Beschorner, K. E., Redfern, M. S., & Cham, R. (2013). Earliest Gait Deviations During Slips: Implications For Recovery. *IIE Transactions on Occupational Ergonomics and Human Factors*, 1(1), 31–37.
- Bevan, S. (2015). Economic impact of musculoskeletal disorders (msds) on work in europe. *Best Practice & Research Clinical Rheumatology*, 29(3), 356–373. Occupation and Musculoskeletal Disorders.
- Bitzer, S. & van der Smagt, P. (2006). Learning EMG control of a robotic hand: towards active prostheses. In *Proceedings 2006 IEEE International Conference on Robotics and Automation, 2006. ICRA 2006*. (pp. 2819–2823). Orlando, FL, USA: IEEE.
- Bolgla, L. A. & Uhl, T. L. (2005). Electromyographic Analysis of Hip Rehabilitation Exercises in a Group of Healthy Subjects. *Journal of Orthopaedic & Sports Physical Therapy*, 35(8), 487–494.
- Bölöni, L. (2011). An investigation into the utility of episodic memory for cognitive architectures. In *AAAI Fall Symposium - Technical Report*, volume FS-11-01 (pp. 42–49).
- Borg, G. (1998). *Borg's Perceived exertion and pain scales*. Champaign, IL: Human Kinetics.
- Boyd, S. P. & Vandenberghe, L. (2004). Regularization. In *Convex Optimization* (pp. 306–310). Cambridge, UK: Cambridge University Press.
- Buchli, J. & Ijspeert, A. J. (2008). Self-organized adaptive legged locomotion in a compliant quadruped robot. *Autonomous Robots*, 25(4), 331–347.
- Burden, A. (2010). How should we normalize electromyograms obtained from healthy participants? What we have learned from over 25years of research. *Journal of Electromyography and Kinesiology*, 20(6), 1023–1035.

- Byrne, F. (2015). Symphony from Synapses: Neocortex as a Universal Dynamical Systems Modeller using Hierarchical Temporal Memory.
- Büschges, A., Akay, T., Gabriel, J. P., & Schmidt, J. (2008). Organizing network action for locomotion: Insights from studying insect walking. *Brain Research Reviews*, 57(1), 162–171.
- Cahen, A. & Tacca, M. C. (2013). Linking perception and cognition. *Frontiers in Psychology*, 4.
- Calandra, R., Gopalan, N., Seyfarth, A., Peters, J., & Deisenroth, M. P. (2014). Bayesian Gait Optimization for Bipedal Locomotion. In P. M. Pardalos, M. G. Resende, C. Vogiatzis, & J. L. Walteros (Eds.), *Learning and Intelligent Optimization* (pp. 274–290). Cham: Springer International Publishing.
- Calinon, S. (2016). A tutorial on task-parameterized movement learning and retrieval. *Intelligent Service Robotics*, 9(1), 1–29.
- Calinon, S., D'halluin, F., Sauser, E., Caldwell, D., & Billard, A. (2010). Learning and Reproduction of Gestures by Imitation. *IEEE Robotics & Automation Magazine*, 17(2), 44–54.
- Calinon, S., Li, Z., Alizadeh, T., Tsagarakis, N. G., & Caldwell, D. G. (2012). Statistical dynamical systems for skills acquisition in humanoids. In *2012 12th IEEE-RAS International Conference on Humanoid Robots (Humanoids 2012)* (pp. 323–329). Osaka, Japan: IEEE.
- Calinon, S., Pistillo, A., & Caldwell, D. G. (2011). Encoding the time and space constraints of a task in explicit-duration Hidden Markov Model. In *2011 IEEE/RSJ International Conference on Intelligent Robots and Systems* (pp. 3413–3418). ISSN: 2153-0866.
- Celik, O., Zhou, D., Li, G., Becker, P., & Neumann, G. (2022). Specializing Versatile Skill Libraries using Local Mixture of Experts. In *Proceedings of the 5th Conference on Robot Learning* (pp. 1423–1433).: PMLR. ISSN: 2640-3498.
- Chakrabarty, D. & Elhilali, M. (2019). A Gestalt inference model for auditory scene segregation. *PLOS Computational Biology*, 15(1), e1006711.
- Chambers, A. J. & Cham, R. (2007). Slip-related muscle activation patterns in the stance leg during walking. *Gait & posture*, 25(4), 565–572.
- Chang, J. (2018). *The Risk Assessment of Work-related Musculoskeletal Disorders based on OpenSim*. Theses, École centrale de Nantes. Issue: 2018ECDN0042.
- Chang, W. R., Leclercq, S., Lockhart, T. E., & Haslam, R. (2016). State of science: occupational slips, trips and falls on the same level. *Ergonomics*, 59(7), 861–883.
- Chatzilygeroudis, K., Rama, R., Kaushik, R., Goepp, D., Vassiliades, V., & Mouret, J.-B. (2017). Black-box data-efficient policy search for robotics. In *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)* (pp. 51–58). Vancouver, BC: IEEE.

- Chatzilygeroudis, K., Vassiliades, V., Stulp, F., Calinon, S., & Mouret, J.-B. (2020). A Survey on Policy Search Algorithms for Learning Robot Controllers in a Handful of Trials. *IEEE Transactions on Robotics*, 36(2), 328–347.
- Chen, C., Yang, C., Zeng, C., Wang, N., & Li, Z. (2017). Robot learning from multiple demonstrations with dynamic movement primitive. In *2017 2nd International Conference on Advanced Robotics and Mechatronics (ICARM)* (pp. 523–528).
- Chen, J., Zhong, S., Kang, E., & Qiao, H. (2019). Realizing human-like manipulation with a musculoskeletal system and biologically inspired control scheme. *Neurocomputing*, 339, 116–129.
- Chen, L. (1982). Topological Structure in Visual Perception. *Science*, 218(4573), 699–700.
- Chen, L. (2005). The topological approach to perceptual organization. *Visual Cognition*, 12(4), 553–637.
- Chestnutt, J., Lau, M., Cheung, G., Kuffner, J., Hodgins, J., & Kanade, T. (2005). Footstep Planning for the Honda ASIMO Humanoid. In *Proceedings of the 2005 IEEE International Conference on Robotics and Automation* (pp. 629–634). Barcelona, Spain: IEEE.
- Chevallereau, C., Abba, G., Aoustin, Y., Plestan, Franck, F., Westervelt, E., Canudas de Wit, C., & Grizzle, J. (2003). RABBIT: a testbed for advanced control theory. *IEEE Control Systems*, 23(5), 57–79.
- Chevallier, S., Jan Ijspeert, A., Ryczko, D., Nagy, F., & Cabelguen, J.-M. (2008). Organisation of the spinal central pattern generators for locomotion in the salamander: Biology and modelling. *Brain Research Reviews*, 57(1), 147–161.
- Chou, Y.-L., You, J.-Y., Lin, C.-J., Su, F.-C., & Chou, P.-H. (2011). The Dynamics of Balance Control During Slipping. *The Chinese Journal of Mechanics. Series A*, 17(2), 69–78.
- Christensen, A. L., O’Grady, R., Birattari, M., & Dorigo, M. (2008). Fault detection in autonomous robots based on fault injection and learning. *Autonomous Robots*, 24(1), 49–67.
- Cimolato, A., Driessen, J. J. M., Mattos, L. S., De Momi, E., Laffranchi, M., & De Michieli, L. (2022). EMG-driven control in lower limb prostheses: A topic-based systematic review. *Journal of NeuroEngineering and Rehabilitation*, 19(1), 43.
- Clocksins, W. F. & Mellish, C. S. (2003). *Programming in Prolog: Using the ISO Standard*. Berlin, Heidelberg s.l.: Springer Berlin Heidelberg, fifth edition edition.
- Coelho, J., Piater, J., & Grupen, R. (2001). Developing haptic and visual perceptual categories for reaching and grasping with a humanoid robot. *Robotics and Autonomous Systems*, 37(2-3), 195–218.
- Collins, S. & Ruina, A. (2005). A Bipedal Walking Robot with Efficient and Human-Like Gait. In *Proceedings of the 2005 IEEE International Conference on Robotics and Automation* (pp. 1983–1988). ISSN: 1050-4729.

- Collins, S., Ruina, A., Tedrake, R., & Wisse, M. (2005). Efficient Bipedal Robots Based on Passive-Dynamic Walkers. *Science*, 307(5712), 1082–1085. Publisher: American Association for the Advancement of Science.
- Collins, S. H., Wisse, M., & Ruina, A. (2001). A Three-Dimensional Passive-Dynamic Walking Robot with Two Legs and Knees. *The International Journal of Robotics Research*, 20(7), 607–615. Publisher: SAGE Publications Ltd STM.
- Craddock, M., Augustine, E., Konerman, S., & Shin, M. (2022). Biorobotics: An Overview of Recent Innovations in Artificial Muscles. *Actuators*, 11(6), 168.
- Crawford, E., Gingerich, M., & Eliasmith, C. (2016). Biologically Plausible, Human-Scale Knowledge Representation. *Cognitive Science*, 40(4), 782–821.
- Crespi, A., Karakasiliotis, K., Guignard, A., & Ijspeert, A. J. (2013). Salamandra Robotica II: An Amphibious Robot to Study Salamander-Like Swimming and Walking Gaits. *IEEE Transactions on Robotics*, 29(2), 308–320. Conference Name: IEEE Transactions on Robotics.
- Cristiano, J., Puig, D., & García, M. A. (2014). Locomotion Control of a Biped Robot through a Feedback CPG Network. In M. A. Armada, A. Sanfeliu, & M. Ferre (Eds.), *ROBOT2013: First Iberian Robotics Conference*, volume 252 (pp. 527–540). Cham: Springer International Publishing. Series Title: Advances in Intelligent Systems and Computing.
- Curby, K. M. & Teichmann, L. (2022). The time course of holistic processing is similar for face and non-face Gestalt stimuli. *Attention, Perception, & Psychophysics*, 84(4), 1234–1247.
- Dalal, M., Pathak, D., & Salakhutdinov, R. (2021). Accelerating Robotic Reinforcement Learning via Parameterized Action Primitives. arXiv:2110.15360.
- D'Alessio, T. & Conforto, S. (2001). Extraction of the envelope from surface EMG signals. *IEEE Engineering in Medicine and Biology Magazine*, 20(6), 55–61.
- Dallali, H. (2012). *Modelling and dynamic stabilisation of a compliant humanoid robot, CoMan*. PhD, The University of Manchester, Manchester, United Kingdom.
- d'Avella, A. (2016). Modularity for Motor Control and Motor Learning. In J. Laczko & M. L. Latash (Eds.), *Progress in Motor Control: Theories and Translations* (pp. 3–19). Cham: Springer International Publishing.
- d'Avella, A. & Bizzi, E. (2005). Shared and specific muscle synergies in natural motor behaviors. *Proceedings of the National Academy of Sciences of the United States of America*, 102(8), 3076–3081.
- d'Avella, A., Portone, A., Fernandez, L., & Lacquaniti, F. (2006). Control of fast-reaching movements by muscle synergy combinations. *Journal of Neuroscience*, 26(30), 7791–7810.

- de Kok, J., Vroonhof, P., Snijders, J., Roullis, G., Clarke, M., Peereboom, K., & van Dorst, P. (2019). *Work-related MSDs: prevalence, costs and demographics in the EU*. Technical report.
- De Luca, C. J. (1997). The Use of Surface Electromyography in Biomechanics. *Journal of Applied Biomechanics*, 13(2), 135–163.
- de-Wit, L. & Wagemans, J. (2015). Individual differences in local and global perceptual organization. In J. Wagemans (Ed.), *The Oxford Handbook of Perceptual Organization*. Oxford University Press.
- Deb, K. (2011). Multi-objective Optimisation Using Evolutionary Algorithms: An Introduction. In L. Wang, A. H. C. Ng, & K. Deb (Eds.), *Multi-Objective Evolutionary Optimisation for Product Design and Manufacturing* (pp. 3–34). London: Springer.
- Degallier, S., Righetti, L., Gay, S., & Ijspeert, A. (2011). Toward simple control for complex, autonomous robotic applications: combining discrete and rhythmic motor primitives. *Autonomous Robots*, 31(2-3), 155–181.
- Deisenroth, M., Englert, P., Peters, J., & Fox, D. (2014). Multi-Task Policy Search for Robotics.
- Deniša, M., Gams, A., Ude, A., & Petrič, T. (2016). Learning Compliant Movement Primitives Through Demonstration and Statistical Generalization. *IEEE/ASME Transactions on Mechatronics*, 21(5), 2581–2594.
- Deniša, M., Petrič, T., Asfour, T., & Ude, A. (2013). Synthesizing compliant reaching movements by searching a database of example trajectories. In *2013 13th IEEE-RAS International Conference on Humanoid Robots (Humanoids)* (pp. 540–546). ISSN: 2164-0580.
- Dev Anand, M., Selvaraj, T., & Kumanan, S. (2012). Fault detection and fault tolerance methods for industrial robot manipulators based on hybrid intelligent approach. *Advances in Production Engineering & Management*, 7(4), 225–236.
- Deverin, A., Gies, V., Marzetti, S., Barchasz, V., & Glotin, H. (2024). Reducing AI Complexity Using Ear Bio-Inspired Primitives. In *2024 32nd European Signal Processing Conference (EUSIPCO)* (pp. 711–715).
- Dimitrijevic, M. R., Gerasimenko, Y., & Pinter, M. M. (1998). Evidence for a spinal central pattern generator in humans. *Annals of the New York Academy of Sciences*, 860, 360–376.
- Dominici, N., Ivanenko, Y. P., Cappellini, G., d'Avella, A., Mondì, V., Cicchese, M., Fabiano, A., Silei, T., Di Paolo, A., Giannini, C., Poppele, R. E., & Lacquaniti, F. (2011). Locomotor primitives in newborn babies and their development. *Science (New York, N.Y.)*, 334(6058), 997–999.
- Doshi, P. & Suthar, V. (2024-01/2024-06). Ab. No. 54 Effect of EMG Driven Robotic-Assisted Rehabilitation on Hand Motor Function in Individuals With Stroke – Systematic Review. *Journal of Society of Indian Physiotherapists*, 8(1), 56.

- Dragan, A. D., Muelling, K., Andrew Bagnell, J., & Srinivasa, S. S. (2015). Movement primitives via optimization. In *2015 IEEE International Conference on Robotics and Automation (ICRA)* (pp. 2339–2346). ISSN: 1050-4729.
- Droniou, A., Ivaldi, S., & Sigaud, O. (2014). Learning a repertoire of actions with deep neural networks. In *4th International Conference on Development and Learning and on Epigenetic Robotics* (pp. 229–234). Genoa, Italy: IEEE.
- Dzeladini, F., Ait-Bouziad, N., & Ijspeert, A. (2019). CPG-Based Control of Humanoid Robot Locomotion. In A. Goswami & P. Vadakkepat (Eds.), *Humanoid Robotics: A Reference* (pp. 1099–1133). Dordrecht: Springer Netherlands.
- Ekeberg, O. & Grillner, S. (1999). Simulations of neuromuscular control in lamprey swimming. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, 354(1385), 895–902.
- Elias, L. J., Bryden, M. P., & Bulman-Fleming, M. B. (1998). Footedness is a better predictor than is handedness of emotional lateralization. *Neuropsychologia*, 36(1), 37–43.
- Eliasmith, C. (2013). Theories of Cognition. In C. Eliasmith (Ed.), *How to Build a Brain: A Neural Architecture for Biological Cognition*. Oxford University Press.
- End, F., Akrou, R., Peters, J., & Neumann, G. (2017). Layered direct policy search for learning hierarchical skills. In *2017 IEEE International Conference on Robotics and Automation (ICRA)* (pp. 6442–6448). Singapore, Singapore: IEEE.
- Endress, A. D., Dehaene-Lambertz, G., & Mehler, J. (2007). Perceptual constraints and the learnability of simple grammars. *Cognition*, 105(3), 577–614.
- Endress, A. D., Nespore, M., & Mehler, J. (2009). Perceptual and memory constraints on language acquisition. *Trends in Cognitive Sciences*, 13(8), 348–353.
- Englsberger, J., Ott, C., Roa, M. A., Albu-Schaffer, A., & Hirzinger, G. (2011). Bipedal walking control based on Capture Point dynamics. In *2011 IEEE/RSJ International Conference on Intelligent Robots and Systems* (pp. 4420–4427). San Francisco, CA: IEEE.
- Englsberger, J., Werner, A., Ott, C., Henze, B., Roa, M. A., Garofalo, G., Burger, R., Beyer, A., Eiberger, O., Schmid, K., & Albu-Schaffer, A. (2014). Overview of the torque-controlled humanoid robot TORO. In *2014 IEEE-RAS International Conference on Humanoid Robots* (pp. 916–923). Madrid, Spain: IEEE.
- Epstein, S. L., Passonneau, R., Gordon, J. B., & Ligorio, T. (2011). The Role of Knowledge and Certainty in Understanding for Dialogue. In *AAAI Fall Symposium: Advances in Cognitive Systems*.
- Eskinazi, I. & Fregly, B. J. (2018). A computational framework for simultaneous estimation of muscle and joint contact forces and body motion using optimization and surrogate modeling. *Medical Engineering & Physics*, 54, 56–64.

- Faraji, S., Pouya, S., Atkeson, C. G., & Ijspeert, A. J. (2014). Versatile and robust 3D walking with a simulated humanoid robot (Atlas): A model predictive control approach. In *2014 IEEE International Conference on Robotics and Automation (ICRA)* (pp. 1943–1950). Hong Kong, China: IEEE.
- Farina, D., Cescon, C., Negro, F., & Enoka, R. M. (2008). Amplitude Cancellation of Motor-Unit Action Potentials in the Surface Electromyogram Can Be Estimated With Spike-Triggered Averaging. *Journal of Neurophysiology*, 100(1), 431–440. Publisher: American Physiological Society.
- Farina, D. & Falla, D. (2008). Estimation of muscle fiber conduction velocity from two-dimensional surface EMG recordings in dynamic tasks. *Biomedical Signal Processing and Control*, 3(2), 138–144.
- Farina, D., Pozzo, M., Merlo, E., Bottin, A., & Merletti, R. (2004). Assessment of Average Muscle Fiber Conduction Velocity From Surface EMG Signals During Fatiguing Dynamic Contractions. *IEEE Transactions on Biomedical Engineering*, 51(8), 1383–1393.
- Feng, Z., Wang, Y., & Liu, L. (2022). Muscle fatigue detection method based on feature extraction and deep learning. *2022 7th International Conference on Intelligent Computing and Signal Processing, ICSP 2022*, (pp. 97–100).
- Ferreira, R. J. N. (2022). Fall Prevention Strategy for an Active Orthotic System. Master's thesis, Universidade do Minho, Guimarães, Portugal.
- Ferreira, R. N., Ribeiro, N. F., Figueiredo, J., & Santos, C. P. (2022). Provoking Artificial Slips and Trips towards Perturbation-Based Balance Training: A Narrative Review. *Sensors*, 22(23).
- Flash, T. & Hochner, B. (2005). Motor primitives in vertebrates and invertebrates. *Current Opinion in Neurobiology*, 15(6), 660–666.
- Flash, T. & Hogan, N. (1985). The coordination of arm movements: an experimentally confirmed mathematical model. *Journal of Neuroscience*, 5(7), 1688–1703. Publisher: Society for Neuroscience Section: Articles.
- Forsyth, D., Malik, J., Fleck, M., & Ponce, J. (1997). Primitives, perceptual organization and object recognition. *Vision Research*.
- Forte, D., Gams, A., Morimoto, J., & Ude, A. (2012). On-line motion synthesis and adaptation using a trajectory database. *Robotics and Autonomous Systems*, 60(10), 1327–1339.
- Frégnac, Y. & Bathellier, B. (2015). Cortical Correlates of Low-Level Perception: From Neural Circuits to Percepts. *Neuron*, 88(1), 110–126.
- Fuglevand, A. J. (1996). Neural Aspects of Fatigue. *The Neuroscientist*, 2(4), 203–206.
- Fuglevand, A. J., Macefield, V. G., & Bigland-Ritchie, B. (1999). Force-frequency and fatigue properties of motor units in muscles that control digits of the human hand. *Journal of Neurophysiology*, 81(4), 1718–1729.

- Fuglevand, A. J., Winter, D. A., Patla, A. E., & Stashuk, D. (1992). Detection of motor unit action potentials with surface electrodes: influence of electrode size and spacing. *Biological Cybernetics*, 67(2), 143–153.
- Fung, M. L., Chen, M. Z. Q., & Chen, Y. H. (2017). Sensor fusion: A review of methods and applications. In *2017 29th Chinese Control And Decision Conference (CCDC)* (pp. 3853–3860).
- Gabriel, D. A. (1997). Shoulder and elbow muscle activity in goal-directed arm movements:. *Experimental Brain Research*, 116(2), 359–366.
- Gams, A., Nemec, B., Ijspeert, A. J., & Ude, A. (2014). Coupling Movement Primitives: Interaction With the Environment and Bimanual Tasks. *IEEE Transactions on Robotics*, 30(4), 816–830.
- Gams, A. & Ude, A. (2009). Generalization of example movements with dynamic systems. In *2009 9th IEEE-RAS International Conference on Humanoid Robots* (pp. 28–33). ISSN: 2164-0580.
- Gauthier, I., Curran, T., Curby, K. M., & Collins, D. (2003). Perceptual interference supports a non-modular account of face processing. *Nature Neuroscience*, 6(4), 428–432.
- Geyer, H. & Herr, H. (2010). A Muscle-Reflex Model That Encodes Principles of Legged Mechanics Produces Human Walking Dynamics and Muscle Activities. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 18(3), 263–273.
- Gibson, J. J. (2014). The Theory of Affordances. In *The Ecological Approach to Visual Perception*. Psychology Press.
- Gilbert, R. O. (1987). *Statistical Methods for Environmental Pollution Monitoring*. Technical Report PNL-4660, Pacific Northwest National Lab. (PNNL), Richland, WA (United States).
- Ginesi, M., Meli, D., Roberti, A., Sansonetto, N., & Fiorini, P. (2021). Dynamic Movement Primitives: Volumetric Obstacle Avoidance Using Dynamic Potential Functions. *Journal of Intelligent & Robotic Systems*, 101(4), 79.
- Giszter, S. F., Moxon, K. A., Rybak, I. A., & Chapin, J. K. (2001). Neurobiological and neurorobotic approaches to control architectures for a humanoid motor system. *Robotics and Autonomous Systems*, 37(2), 219–235.
- Gizzi, L., Yavuz, U. c., Hillerkuss, D., Geri, T., Gneiting, E., Domeier, F., Schmitt, S., & Röhrle, O. (2021). Variations in Muscle Activity and Exerted Torque During Temporary Blood Flow Restriction in Healthy Individuals. *Frontiers in Bioengineering and Biotechnology*, 9, 557761.
- Gomez-Gonzalez, S., Neumann, G., Schölkopf, B., & Peters, J. (2016). Using probabilistic movement primitives for striking movements. In *2016 IEEE-RAS 16th International Conference on Humanoid Robots (Humanoids)* (pp. 502–508). ISSN: 2164-0580.

- Gong, Y., Hartley, R., Da, X., Hereid, A., Harib, O., Huang, J.-K., & Grizzle, J. (2019). Feedback Control of a Cassie Bipedal Robot: Walking, Standing, and Riding a Segway. In *2019 American Control Conference (ACC)* (pp. 4559–4566). ISSN: 2378-5861.
- Grandjean, E. (1979). Fatigue in industry. *Occupational and Environmental Medicine*, 36(3), 175–186.
- Grillner, S. & El Manira, A. (2020). Current Principles of Motor Control, with Special Reference to Vertebrate Locomotion. *Physiological Reviews*, 100(1), 271–320. Publisher: American Physiological Society.
- Grizzle, J., Abba, G., & Plestan, F. (2001). Asymptotically stable walking for biped robots: analysis via systems with impulse effects. *IEEE Transactions on Automatic Control*, 46(1), 51–64. Conference Name: IEEE Transactions on Automatic Control.
- Grizzle, J. W., Chevallereau, C., Sinnet, R. W., & Ames, A. D. (2014). Models, feedback control, and open problems of 3D bipedal robotic walking. *Automatica*, 50(8), 1955–1988.
- Grossberg, S. & Myers, C. W. (2000). The resonant dynamics of speech perception: Interword integration and duration-dependent backward effects. *Psychological Review*, 107(4), 735–767.
- Grotz, M., Kaiser, P., Aksoy, E. E., Paus, F., & Asfour, T. (2017). Graph-based visual semantic perception for humanoid robots. In *2017 IEEE-RAS 17th International Conference on Humanoid Robotics (Humanoids)* (pp. 869–875). Birmingham: IEEE.
- Gu, S., Holly, E., Lillicrap, T., & Levine, S. (2017). Deep reinforcement learning for robotic manipulation with asynchronous off-policy updates. In *2017 IEEE International Conference on Robotics and Automation (ICRA)* (pp. 3389–3396). Singapore, Singapore: IEEE.
- Guan, Y., Wang, N., & Yang, C. (2021). Review of the techniques used in motor-cognitive human-robot skill transfer. *Cognitive Computation and Systems*, 3(3), 229–252. _eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1049/ccs2.12025>.
- Guenter, F., Hersch, M., Calinon, S., & Billard, A. (2007). Reinforcement Learning for Imitating Constrained Reaching Movements. *Advanced Robotics*, 21.
- Hajiloo, B., Anbarian, M., Esmaeili, H., & Mirzapour, M. (2020). The effects of fatigue on synergy of selected lower limb muscles during running. *Journal of Biomechanics*, 103, 109692.
- Hansen, N. (2016). The CMA Evolution Strategy: A Tutorial.
- Hansen, N., Müller, S. D., & Koumoutsakos, P. (2003). Reducing the Time Complexity of the Derandomized Evolution Strategy with Covariance Matrix Adaptation (CMA-ES). *Evolutionary Computation*, 11(1), 1–18. Conference Name: Evolutionary Computation.
- Hansen, N. & Ostermeier, A. (2001). Completely Derandomized Self-Adaptation in Evolution Strategies. *Evolutionary Computation*, 9(2), 159–195.

- Harnad, S. (1990). The symbol grounding problem. *Physica D: Nonlinear Phenomena*, 42(1), 335–346.
- Hase, K., Miyashita, K., Ok, S., & Arakawa, Y. (2003). Human gait simulation with a neuromusculoskeletal model and evolutionary computation. *Journal of Visualization and Computer Animation*, 14, 73–92.
- Hasselmo, M. E. (2011). *How We Remember: Brain Mechanisms of Episodic Memory*. The MIT Press.
- He, B., Si, Y., Wang, Z., & Zhou, Y. (2019). Hybrid CPG–FRI dynamic walking algorithm balancing agility and stability control of biped robot. *Autonomous Robots*, 43(7), 1855–1865.
- He, J. Z.-Y., Erickson, Z., Brown, D. S., Raghunathan, A., & Dragan, A. (2023). Learning representations that enable generalization in assistive tasks. In K. Liu, D. Kulic, & J. Ichnowski (Eds.), *Proceedings of The 6th Conference on Robot Learning*, volume 205 of *Proceedings of Machine Learning Research* (pp. 2105–2114).: PMLR.
- Heidrich-Meisner, V. & Igel, C. (2008). Evolution Strategies for Direct Policy Search. In G. Rudolph, T. Jansen, N. Beume, S. Lucas, & C. Poloni (Eds.), *Parallel Problem Solving from Nature – PPSN X*, volume 5199 (pp. 428–437). Berlin, Heidelberg: Springer Berlin Heidelberg. Series Title: Lecture Notes in Computer Science.
- Hereid, A., Hubicki, C. M., Cousineau, E. A., & Ames, A. D. (2018). Dynamic Humanoid Locomotion: A Scalable Formulation for HZD Gait Optimization. *IEEE Transactions on Robotics*, 34(2), 370–387.
- Hersch, M., Guenter, F., Calinon, S., & Billard, A. (2008). Dynamical System Modulation for Robot Learning via Kinesthetic Demonstrations. *IEEE Transactions on Robotics*, 24(6), 1463–1467.
- Higgins, I., Sonnerat, N., Matthey, L., Pal, A., Burgess, C. P., Bosnjak, M., Shanahan, M., Botvinick, M., Hassabis, D., & Lerchner, A. (2018). SCAN: Learning Hierarchical Compositional Visual Concepts.
- Hobbelen, D. G. E. & Wisse, M. (2007). Limit Cycle Walking. In M. Hackel (Ed.), *Humanoid Robots, Human-like Machines*. I-Tech Education and Publishing.
- Hock, H. S. (2015). Dynamic grouping motion: A method for determining perceptual organization for objects with connected surfaces. In J. Wagemans (Ed.), *The Oxford Handbook of Perceptual Organization*. Oxford University Press.
- Hof, A. L. (2008). The ‘extrapolated center of mass’ concept suggests a simple control of balance in walking. *Human Movement Science*, 27(1), 112–125.
- Hogan, N. & Sternad, D. (2013). Dynamic primitives in the control of locomotion. *Frontiers in Computational Neuroscience*, 7.
- Hogrel, J.-Y., Ledoux, I., & Duchêne, J. (2008). Reliability of muscle fibre conduction velocity distribution estimation from surface EMG. *Biomedical Signal Processing and Control*, 3(2), 118–125.

- Hooper, S. L. (2001). Central Pattern Generators. In eLS. John Wiley & Sons, Ltd. _eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1038/npg.els.0000032>.
- Hospedales, T. M., Antoniou, A., Micaelli, P., & Storkey, A. J. (2021). Meta-Learning in Neural Networks: A Survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, (pp. 1–1).
- Hsiao, K., Chitta, S., Ciocarlie, M., & Jones, E. G. (2010). Contact-reactive grasping of objects with partial shape information. In *2010 IEEE/RSJ International Conference on Intelligent Robots and Systems* (pp. 1228–1235). Taipei: IEEE.
- Huang, Y., Huang, K., & Tan, T. (2009). Computational primitives of visual perception. In *2009 16th IEEE International Conference on Image Processing (ICIP)* (pp. 1793–1796). Cairo, Egypt: IEEE.
- Huang, Y., Huang, K., Tan, T., & Tao, D. (2010). A Novel Visual Organization Based on Topological Perception. In D. Hutchison, T. Kanade, J. Kittler, J. M. Kleinberg, F. Mattern, J. C. Mitchell, M. Naor, O. Nierstrasz, C. Pandu Rangan, B. Steffen, M. Sudan, D. Terzopoulos, D. Tygar, M. Y. Vardi, G. Weikum, H. Zha, R.-i. Taniguchi, & S. Maybank (Eds.), *Computer Vision – ACCV 2009*, volume 5994 (pp. 180–189). Berlin, Heidelberg: Springer Berlin Heidelberg. Series Title: Lecture Notes in Computer Science.
- Huang, Y., Rozo, L., Silvério, J., & Caldwell, D. G. (2019). Kernelized movement primitives. *The International Journal of Robotics Research*, 38(7), 833–852.
- Hurmuzlu, Y. & Moskowitz, G. (1986). The role of impact in the stability of bipedal locomotion. *Dynamic and Stability of Systems*, 1.
- Ibarz, J., Tan, J., Finn, C., Kalakrishnan, M., Pastor, P., & Levine, S. (2021). How to train your robot with deep reinforcement learning: lessons we have learned. *The International Journal of Robotics Research*, 40(4-5), 698–721.
- Ijspeert, A., Nakanishi, J., & Schaal, S. (2002). Movement imitation with nonlinear dynamical systems in humanoid robots. In *Proceedings 2002 IEEE International Conference on Robotics and Automation (Cat. No.02CH37292)*, volume 2 (pp. 1398–1403 vol.2).
- Ijspeert, A. J. (2008). Central pattern generators for locomotion control in animals and robots: a review. *Neural networks*, 21(4), 642–653.
- Ijspeert, A. J. (2014). Biorobotics: Using robots to emulate and investigate agile locomotion. *Science*, 346(6206), 196–203.
- Ijspeert, A. J., Crespi, A., Ryczko, D., & Cabelguen, J.-M. (2007). From Swimming to Walking with a Salamander Robot Driven by a Spinal Cord Model. *Science*, 315(5817), 1416–1420. Publisher: American Association for the Advancement of Science.
- Ijspeert, A. J., Nakanishi, J., Hoffmann, H., Pastor, P., & Schaal, S. (2013). Dynamical movement primitives: learning attractor models for motor behaviors. *Neural Computation*, 25(2), 328–373.

- Ijspeert, A. J., Nakanishi, J., & Schaal, S. (2003). Learning Attractor Landscapes for Learning Motor Primitives.
- Il Hong Suh, Gi Hyun Lim, Wonil Hwang, Hyowon Suh, Jung-Hwa Choi, & Young-Tack Park (2007). Ontology-based multi-layered robot knowledge framework (OMRKF) for robot intelligence. In *2007 IEEE/RSJ International Conference on Intelligent Robots and Systems* (pp. 429–436). San Diego, CA, USA: IEEE.
- Inman, V. T. (1966). Human Locomotion. *Canadian Medical Association Journal*, 94(20), 1047–1054.
- Irani, P. & Ware, C. (2003). Diagramming information structures using 3D perceptual primitives. *ACM Transactions on Computer-Human Interaction*, 10(1), 1–19.
- Ivaldi, S., Nguyen, S. M., Lyubova, N., Droniou, A., Padois, V., Filliat, D., Oudeyer, P.-Y., & Sigaud, O. (2014). Object Learning Through Active Exploration. *IEEE Transactions on Autonomous Mental Development*, 6(1), 56–72.
- Ivanenko, Y. P., Cappellini, G., Dominici, N., Poppele, R. E., & Lacquaniti, F. (2005). Coordination of Locomotion with Voluntary Movements in Humans. *The Journal of Neuroscience*, 25(31), 7238–7253.
- Ivashko, D. G., Prilutsky, B. I., Markin, S. N., Chapin, J. K., & Rybak, I. A. (2003). Modeling the spinal cord neural circuitry controlling cat hindlimb movement during locomotion. *Neurocomputing*, 52-54, 621–629.
- Jang, G., Kim, J., Lee, S., & Choi, Y. (2016). EMG-Based Continuous Control Scheme With Simple Classifier for Electric-Powered Wheelchair. *IEEE Transactions on Industrial Electronics*, 63(6), 3695–3705.
- Jiang, Y., Malliaras, P., Chen, B., & Kulic, D. (2022). Real-time forecasting of exercise-induced fatigue from wearable sensors. *Computers in Biology and Medicine*, 148, 105905.
- Johnson, M., Shrewsbury, B., Bertrand, S., Wu, T., Duran, D., Floyd, M., Abeles, P., Stephen, D., Mertins, N., Lesman, A., Carff, J., Rifenburgh, W., Kaveti, P., Straatman, W., Smith, J., Griffioen, M., Layton, B., De Boer, T., Koolen, T., Neuhaus, P., & Pratt, J. (2015). Team IHMC's Lessons Learned from the DARPA Robotics Challenge Trials. *Journal of Field Robotics*, 32(2), 192–208.
- Kaiser, A., Ybanez Zepeda, J. A., & Boubekeur, T. (2019). A Survey of Simple Geometric Primitives Detection Methods for Captured 3D Data. *Computer Graphics Forum*, 38(1), 167–196.
- Kajita, S. & Espiau, B. (2008). Legged Robots. In *Handbook of Robotics*, volume Part B.
- Kalakrishnan, M., Righetti, L., Pastor, P., & Schaal, S. (2011). Learning force control policies for compliant manipulation. In *2011 IEEE/RSJ International Conference on Intelligent Robots and Systems* (pp. 4639–4644). San Francisco, CA: IEEE.

- Kalashnikov, D., Varley, J., Chebotar, Y., Swanson, B., Jonschkowski, R., Finn, C., Levine, S., & Hausman, K. (2021). MT-Opt: Continuous Multi-Task Robotic Reinforcement Learning at Scale. arXiv:2104.08212.
- Kamthe, S. & Deisenroth, M. (2018). Data-Efficient Reinforcement Learning with Probabilistic Model Predictive Control. In *Proceedings of the Twenty-First International Conference on Artificial Intelligence and Statistics* (pp. 1701–1710).: PMLR. ISSN: 2640-3498.
- Kaneko, K., Kanehiro, F., Kajita, S., Yokoyama, K., Akachi, K., Kawasaki, T., Ota, S., & Isozumi, T. (2002). Design of prototype humanoid robotics platform for HRP. In *IEEE/RSJ International Conference on Intelligent Robots and System*, volume 3 (pp. 2431–2436). Lausanne, Switzerland: IEEE.
- Karakasiliotis, K., Thandiackal, R., Melo, K., Horvat, T., Mahabadi, N. K., Tsitkov, S., Cabelguen, J. M., & Ijspeert, A. J. (2016). From cineradiography to biorobots: an approach for designing robots to emulate and study animal locomotion. *Journal of The Royal Society Interface*, 13(119), 20151089.
- Karlsson, M. K., Magnusson, H., von Schewelow, T., & Rosengren, B. E. (2013). Prevention of falls in the elderly—a review. *Osteoporosis international : a journal established as result of cooperation between the European Foundation for Osteoporosis and the National Osteoporosis Foundation of the USA*, 24(3), 747–762.
- Kaymak, İ., Uçar, A., & Güzeliş, C. (2023). Development of a New Robust Stable Walking Algorithm for a Humanoid Robot Using Deep Reinforcement Learning with Multi-Sensor Data Fusion. *Electronics*, 12(3), 568.
- Kendall, M. (1970). *Rank Correlation Methods*. Theory and applications of rank order-statistics. Griffin.
- Khansari-Zadeh, S. M. & Billard, A. (2011). Learning Stable Nonlinear Dynamical Systems With Gaussian Mixture Models. *IEEE Transactions on Robotics*, 27(5), 943–957. Conference Name: IEEE Transactions on Robotics.
- Kim, Y., Tagawa, Y., Obinata, G., & Hase, K. (2011). Robust control of CPG-based 3D neuromusculoskeletal walking model. *Biological Cybernetics*, 105(3), 269–282.
- Klatzky, R. & Reed, C. L. (2016). Haptic Exploration. In T. Prescott, E. Ahissar, & E. Izhikevich (Eds.), *Scholarpedia of Touch* (pp. 177–183). Paris: Atlantis Press.
- Klemetti, R., Moilanen, P., Avela, J., & Timonen, J. (2014). Effects of gait speed on stability of walking revealed by simulated response to tripping perturbation. *Gait & Posture*, 39(1), 534–539.
- Klink, P., Abdulsamad, H., Belousov, B., & Peters, J. (2020). Self-paced contextual reinforcement learning. 100, 513–529.
- Knowlton, B. J. & Schorn, J. M. (2024). Procedural and Motor Learning. In M. J. Kahana & A. D. Wagner (Eds.), *The Oxford Handbook of Human Memory, Two Volume Pack: Foundations and Applications*. Oxford University Press.

- Kober, J. (2014). *Learning motor skills: from algorithms to robot experiments*. PhD thesis.
- Kober, J., Mülling, K., Krömer, O., Lampert, C. H., Schölkopf, B., & Peters, J. (2010). Movement templates for learning of hitting and batting. In *2010 IEEE International Conference on Robotics and Automation* (pp. 853–858). ISSN: 1050-4729.
- Kober, J., Öztog, E., & Peters, J. (2011). Reinforcement learning to adjust robot movements to new situations. In *IJCAI Proceedings-International Joint Conference on Artificial Intelligence*, volume 22 (pp. 2650).
- Kober, J. & Peters, J. (2009). Learning motor primitives for robotics. In *2009 IEEE International Conference on Robotics and Automation* (pp. 2112–2118). Kobe: IEEE.
- Kober, J. & Peters, J. (2011). Policy search for motor primitives in robotics. *Machine Learning*, 84(1-2), 171–203.
- Kober, J., Wilhelm, A., Oztop, E., & Peters, J. (2012). Reinforcement learning to adjust parametrized motor primitives to new situations. *Autonomous Robots*, 33(4), 361–379.
- Koenig, N. & Howard, A. (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)*, volume 3 (pp. 2149–2154 vol.3).
- Kormushev, P., Calinon, S., & Caldwell, D. G. (2010). Robot motor skill coordination with EM-based Reinforcement Learning. In *2010 IEEE/RSJ International Conference on Intelligent Robots and Systems* (pp. 3232–3237). Taipei: IEEE.
- Kotseruba, I., Gonzalez, O. J. A., & Tsotsos, J. K. (2016). A Review of 40 Years of Cognitive Architecture Research: Focus on Perception, Attention, Learning and Applications.
- Kotseruba, I. & Tsotsos, J. K. (2018). 40 years of cognitive architectures: Core cognitive abilities and practical applications. *Artificial Intelligence Review*, 53(1), 17–94.
- Koutras, L. & Doulgeri, Z. (2020). Dynamic Movement Primitives for moving goals with temporal scaling adaptation. In *2020 IEEE International Conference on Robotics and Automation (ICRA)* (pp. 144–150). Paris, France: IEEE.
- Kragic, D. & Christensen, H. I. (2003). Robust Visual Servoing. *The International Journal of Robotics Research*, 22(10-11), 923–939.
- Kramberger, A., Gams, A., Nemec, B., Chrysostomou, D., Madsen, O., & Ude, A. (2017). Generalization of orientation trajectories and force-torque profiles for robotic assembly. *Robotics and Autonomous Systems*, 98, 333–346.

- Kroemer, O., Niekum, S., & Konidaris, G. (2020). A Review of Robot Learning for Manipulation: Challenges, Representations, and Algorithms. Number: arXiv:1907.03146.
- Krug, R. & Dimitrov, D. (2015). Model Predictive Motion Control based on Generalized Dynamical Movement Primitives. *Journal of Intelligent & Robotic Systems*, 77(1), 17–35.
- Kudithipudi, D., Aguilar-Simon, M., Babb, J., Bazhenov, M., Blackiston, D., Bongard, J., Brna, A. P., Chakravarthi Raja, S., Cheney, N., Clune, J., Daram, A., Fusi, S., Helfer, P., Kay, L., Ketz, N., Kira, Z., Kolouri, S., Krichmar, J. L., Kriegman, S., Levin, M., Madireddy, S., Manicka, S., Marjaninejad, A., McNaughton, B., Miikkulainen, R., Navratilova, Z., Pandit, T., Parker, A., Pilly, P. K., Risi, S., Sejnowski, T. J., Soltoggio, A., Soures, N., Tolia, A. S., Urbina-Meléndez, D., Valero-Cuevas, F. J., Van De Ven, G. M., Vogelstein, J. T., Wang, F., Weiss, R., Yanguas-Gil, A., Zou, X., & Siegelmann, H. (2022). Biological underpinnings for lifelong learning machines. *Nature Machine Intelligence*, 4(3), 196–210.
- Kulvicius, T., Ning, K., Tamosiunaite, M., & Wörgötter, F. (2012). Joining Movement Sequences: Modified Dynamic Movement Primitives for Robotics Applications Exemplified on Handwriting. *IEEE Transactions on Robotics*, 28(1), 145–157.
- Kumar, A. A. (2021). Semantic memory: A review of methods, models, and current challenges. *Psychonomic Bulletin & Review*, 28(1), 40–80.
- Kupcsik, A., Deisenroth, M., Peters, J., & Neumann, G. (2013). Data-Efficient Generalization of Robot Skills with Contextual Policy Search. *Proceedings of the AAAI Conference on Artificial Intelligence*, 27(1), 1401–1407.
- Kupcsik, A., Deisenroth, M. P., Peters, J., Loh, A. P., Vadakkepat, P., & Neumann, G. (2017). Model-based contextual policy search for data-efficient generalization of robot skills. *Artificial Intelligence*, 247, 415–439.
- Kyrdalen, I. L., Thingstad, P., Sandvik, L., & Ormstad, H. (2019). Associations between gait speed and well-known fall risk factors among community-dwelling older adults. *Physiotherapy Research International: The Journal for Researchers and Clinicians in Physical Therapy*, 24(1), e1743.
- Laird, J. E., Newell, A., & Rosenbloom, P. S. (1987). SOAR: An architecture for general intelligence. *Artificial Intelligence*, 33(1), 1–64.
- Langley, P. & Choi, D. (2006). A Unified Cognitive Architecture for Physical Agents. In *AAAI Conference on Artificial Intelligence*.
- Lee, D. D. & Seung, H. S. (1999). Learning the parts of objects by non-negative matrix factorization. *Nature*, 401(6755), 788–791.
- Lee, H., Kim, H., & Kim, H. J. (2018). Planning and Control for Collision-Free Cooperative Aerial Transportation. *IEEE Transactions on Automation Science and Engineering*, 15(1), 189–201. Conference Name: IEEE Transactions on Automation Science and Engineering.

- Lee, H.-M. (2019). Force direction and arm position affect contribution of clavicular and sternal parts of pectoralis major muscle during muscle strength testing. *Journal of Hand Therapy*, 32(1), 71–79.
- Lenat, D. B. (1995). CYC: a large-scale investment in knowledge infrastructure. *Commun. ACM*, 38(11), 33–38.
- Lenz, I., Lee, H., & Saxena, A. (2015). Deep learning for detecting robotic grasps. *The International Journal of Robotics Research*, 34(4-5), 705–724.
- Levine, S., Finn, C., Darrell, T., & Abbeel, P. (2016). End-to-End Training of Deep Visuomotor Policies. *Journal of Machine Learning Research*, 17(39), 1–40.
- Levine, S., Pastor, P., Krizhevsky, A., Ibarz, J., & Quillen, D. (2018). Learning hand-eye coordination for robotic grasping with deep learning and large-scale data collection. *The International Journal of Robotics Research*, 37(4-5), 421–436.
- Li, G., Jin, Z., Volpp, M., Otto, F., Lioutikov, R., & Neumann, G. (2023). ProDMP: A Unified Perspective on Dynamic and Probabilistic Movement Primitives. *IEEE Robotics and Automation Letters*, 8(4), 2325–2332. Conference Name: IEEE Robotics and Automation Letters.
- Li, M., Li, J., & Shu, M. (2020). Detection of Muscle Fatigue by Fusion of Agonist and Synergistic Muscle sEMG Signals. In *2020 IEEE 33rd International Symposium on Computer-Based Medical Systems (CBMS)* (pp. 95–98). ISSN: 2372-9198.
- Li, Y. (2018). Deep Reinforcement Learning: An Overview. Number: arXiv:1701.07274 arXiv:1701.07274.
- Li, Z., Zhao, T., Chen, F., Hu, Y., Su, C.-Y., & Fukuda, T. (2018). Reinforcement Learning of Manipulation and Grasping Using Dynamical Movement Primitives for a Humanoidlike Mobile Manipulator. *IEEE/ASME Transactions on Mechatronics*, 23(1), 121–131. Conference Name: IEEE/ASME Transactions on Mechatronics.
- Lioutikov, R., Neumann, G., Maeda, G., & Peters, J. (2017). Learning movement primitive libraries through probabilistic segmentation. *The International Journal of Robotics Research*, 36(8), 879–894. Publisher: SAGE Publications Ltd STM.
- Liu, G. L., Habib, M. K., Watanabe, K., & Izumi, K. (2008). Central pattern generators based on Matsuoka oscillators for the locomotion of biped robots. *Artificial Life and Robotics*, 12(1-2), 264–269.
- Liu, H. & Singh, P. (2004). ConceptNet - A Practical Commonsense Reasoning Tool-Kit. *BT Technology Journal*, 22(4), 211–226.
- Liu, J. Z., Brown, R. W., & Yue, G. H. (2002). A Dynamical Model of Muscle Activation, Fatigue, and Recovery. *Biophysical Journal*, 82(5), 2344–2359.

- Lockhart, T. E. (2013). Biomechanics of Human Gait – Slip and Fall Analysis. In J. A. Siegel, P. J. Saukko, & M. M. Houck (Eds.), *Encyclopedia of Forensic Sciences* (pp. 466–476). Waltham: Academic Press, second edition edition.
- Ma, W.-L., Kolathaya, S., Ambrose, E. R., Hubicki, C. M., & Ames, A. D. (2017). Bipedal Robotic Running with DURUS-2D: Bridging the Gap between Theory and Experiment. In *Proceedings of the 20th International Conference on Hybrid Systems: Computation and Control* (pp. 265–274). Pittsburgh Pennsylvania USA: ACM.
- Macedo, J., Santos, C., & Costa, L. (2014). Using cost-regularized kernel regression with a high number of samples. In *2014 IEEE International Conference on Autonomous Robot Systems and Competitions (ICARSC)* (pp. 199–204).: IEEE.
- Macedo, J. F. F. S. (2014). Towards the improvement of robot motion learning techniques. Master's thesis.
- Madl, T., Franklin, S., Chen, K., Montaldi, D., & Trapp, R. (2016). Towards real-world capable spatial memory in the LIDA cognitive architecture. *Biologically Inspired Cognitive Architectures*, 16, 87–104.
- Magyar, B., Tsiogkas, N., Brito, B., Patel, M., Lane, D., & Wang, S. (2019). Guided Stochastic Optimization for Motion Planning. *Frontiers in Robotics and AI*, 6. Publisher: Frontiers.
- Marigold, D. S. & Patla, A. E. (2002). Strategies for dynamic stability during locomotion on a slippery surface: Effects of prior experience and knowledge. *Journal of Neurophysiology*, 88(1), 339–353.
- Marsh, B., Sadka, A. H., & Bahai, H. (2022). A Critical Review of Deep Learning-Based Multi-Sensor Fusion Techniques. *Sensors*, 22(23), 9364.
- Martelli, D., Aprigliano, F., Tropea, P., Pasquini, G., Micera, S., & Monaco, V. (2017). Stability against backward balance loss: Age-related modifications following slip-like perturbations of multiple amplitudes. *Gait & Posture*, 53, 207–214.
- Martens, S., Carpenter, G. A., & Gaudiano, P. (1998). Neural sensor fusion for spatial visualization on a mobile robot. In P. S. Schenker & G. T. McKee (Eds.), *Photonics East (ISAM, VVDC, IEMB)* (pp. 100–111). Boston, MA.
- Martínez-Gómez, J., Marfil, R., Calderita, L. V., Bandera, J. P., Manso, L. J., Bandera, A., Romero-Garcés, A., & Bustos, P. (2014). Toward social cognition in robotics: Extracting and internalizing meaning from perception. In *XV Workshop of Physical Agents (WAF 2014), León, Spain* (pp. 93–104).
- Matheson, H. E. & Barsalou, L. W. (2018). Embodiment and Grounding in Cognitive Neuroscience. In J. T. Wixted (Ed.), *Stevens' Handbook of Experimental Psychology and Cognitive Neuroscience* (pp. 1–27). Wiley, 1. edition.

- Matos, V. & Santos, C. (2012). Central pattern generators with phase regulation for the control of humanoid locomotion. In *12th IEEE-RAS International Conference on Humanoid Robots (Humanoids)* (pp. 134–139).: IEEE.
- Matos, V. & Santos, C. P. (2014). Towards goal-directed biped locomotion: combining cpgs and motion primitives. *Robotics and Autonomous Systems*, 62(12), 1669–1690.
- Matos, V. E. d. S. (2014). *A Bio-inspired architecture for adaptive quadruped locomotion over irregular terrain*. PhD thesis.
- Matsubara, T., Hyon, S.-H., & Morimoto, J. (2011). Learning parametric dynamic movement primitives from multiple demonstrations. *Neural Networks*, 24(5), 493–500.
- Matsunaga, N., Okubo, Y., Isagawa, S., Niitsuma, J., Otsudo, T., Sawada, Y., & Akasaka, K. (2021). Muscle fatigue in the gluteus maximus changes muscle synergies during single-leg landing. *Journal of Bodywork and Movement Therapies*, 27, 493–499.
- Matuszek, C., Cabral, J., Witbrock, M., & DeOliveira, J. (2006). : (pp. 44–49).
- Mausfeld, R. (2003). ‘Colour’ As Part of the Format of Different Perceptual Primitives: The Dual Coding of Colour. In R. Mausfeld & D. Heyer (Eds.), *Colour Perception: Mind and the Physical World*. Oxford University Press.
- Mayeli, M. (2019). Neurophysiology of Visual Perception. In N. Rezaei & A. Saghzadeh (Eds.), *Biophysics and Neurophysiology of the Sixth Sense* (pp. 13–26). Cham: Springer International Publishing.
- McGeer, T. (1990). Passive Dynamic Walking. *The International Journal of Robotics Research*, 9(2), 62–82.
- Mei, Z., Chen, Y., Jiang, M., Wu, H., & Cheng, L. (2017). Mobile robots path planning based on dynamic movement primitives library. In *2017 36th Chinese Control Conference (CCC)* (pp. 6906–6911). Dalian, China: IEEE.
- Meier, F., Theodorou, E., Stulp, F., & Schaal, S. (2011). Movement segmentation using a primitive library. In *2011 IEEE/RSJ International Conference on Intelligent Robots and Systems* (pp. 3407–3412). San Francisco, CA: IEEE.
- Menager, D. H. & Choi, D. (2016). A Robust Implementation of Episodic Memory for a Cognitive Architecture. *Proceedings of the Annual Meeting of the Cognitive Science Society*, 38(0).
- Menciassi, A. & Laschi, C. (2014). Biorobotics. In *Robotics: Concepts, Methodologies, Tools, and Applications* (pp. 1613–1643). IGI Global.
- Messeri, C., Bicchi, A., Zanchettin, A. M., & Rocco, P. (2022). A Dynamic Task Allocation Strategy to Mitigate the Human Physical Fatigue in Collaborative Robotics. *IEEE Robotics and Automation Letters*, 7(2), 2178–2185. Conference Name: IEEE Robotics and Automation Letters.

- Meyer, A. J., Eskinazi, I., Jackson, J. N., Rao, A. V., Patten, C., & Fregly, B. J. (2016). Muscle Synergies Facilitate Computational Prediction of Subject-Specific Walking Motions. *Frontiers in Bioengineering and Biotechnology*, 4. Publisher: Frontiers.
- Mikołajczyk, T., Mikołajewska, E., Al-Shuka, H. F. N., Malinowski, T., Kłodowski, A., Pimenov, D. Y., Paczkowski, T., Hu, F., Giasin, K., Mikołajewski, D., & Macko, M. (2022). Recent Advances in Bipedal Walking Robots: Review of Gait, Drive, Sensors and Control Systems. *Sensors*, 22(12), 4440.
- Miller, B. D. & Clark, J. E. (2015). Towards highly-tuned mobility in multiple domains with a dynamical legged platform. *Bioinspiration & Biomimetics*, 10(4), 046001.
- Mioskowska, M., Stevenson, D., Onu, M., & Trkov, M. (2020). Compressed gas actuated knee assistive exoskeleton for slip-induced fall prevention during human walking. In *IEEE/ASME International Conference on Advanced Intelligent Mechatronics, AIM*, volume 2020-July (pp. 735–740).
- Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., Graves, A., Riedmiller, M., Fidjeland, A. K., Ostrovski, G., Petersen, S., Beattie, C., Sadik, A., Antonoglou, I., King, H., Kumaran, D., Wierstra, D., Legg, S., & Hassabis, D. (2015). Human-level control through deep reinforcement learning. *Nature*, 518(7540), 529–533.
- Mohan, V., Sandini, G., & Morasso, P. (2014). A Neural Framework for Organization and Flexible Utilization of Episodic Memory in Cumulatively Learning Baby Humanoids. *Neural Computation*, 26(12), 2692–2734.
- Monaco, V., Tropea, P., Aprigliano, F., Martelli, D., Parri, A., Cortese, M., Molino-Lova, R., Vitiello, N., & Micera, S. (2017). An ecologically-controlled exoskeleton can improve balance recovery after slippage. *Scientific Reports*, 7.
- Monjo, F. & Forestier, N. (2016). Muscle fatigue effects can be anticipated to reproduce a movement kinematics learned without fatigue. *Neuroscience*, 339, 100–108.
- Moro, F., Tsagarakis, N., & Caldwell, D. (2012). On the kinematic Motion Primitives (kMPs) – theory and application. *Frontiers in Neurobotics*, 6.
- Moyer, B. E., Redfern, M. S., & Cham, R. (2009). Biomechanics of trailing leg response to slipping - Evidence of interlimb and intralimb coordination. *Gait and Posture*, 29(4), 565–570.
- Mülling, K., Kober, J., Kroemer, O., & Peters, J. (2013). Learning to select and generalize striking movements in robot table tennis. *The International Journal of Robotics Research*, 32(3), 263–279.
- Mülling, K., Kober, J., & Peters, J. (2010). Learning table tennis with a mixture of motor primitives. In *Humanoid Robots (Humanoids), 2010 10th IEEE-RAS International Conference on* (pp. 411–416).: IEEE.
- Mussa-Ivaldi, F. A. (1999). Modular features of motor control and learning. *Current Opinion in Neurobiology*, 9(6), 713–717.

- Márquez-Figueroa, S., Shmaliy, Y. S., & Ibarra-Manzano, O. (2020). Optimal extraction of emg signal envelope and artifacts removal assuming colored measurement noise. *Biomedical Signal Processing and Control*, 57, 101679.
- Nachstedt, T., Tetzlaff, C., & Manoonpong, P. (2017). Fast Dynamical Coupling Enhances Frequency Adaptation of Oscillators for Robotic Locomotion Control. *Frontiers in Neurorobotics*, 11.
- Nadel, L. (1994). Multiple Memory Systems: What and Why, an Update. In D. L. Schacter & E. Tulving (Eds.), *Memory Systems 1994* (pp. 39–63). The MIT Press.
- Najarian, K. & Splinter, R. (2012). *Biomedical signal and image processing*. Boca Raton: CRC Press, 2nd ed edition. OCLC: 1317347438.
- Nakanishi, J., Morimoto, J., Endo, G., Cheng, G., Schaal, S., & Kawato, M. (2004). Learning from demonstration and adaptation of biped locomotion. *Robotics and Autonomous Systems*, 47(2-3), 79–91.
- Nandedkar, S. D. & Sanders, D. B. (1990). Measurement of the amplitude of the EMG envelope. *Muscle & Nerve*, 13(10), 933–938.
- Napolitano, A. (2020). Characterization of stochastic processes. In *Cyclostationary Processes and Time Series* (pp. 3–35). Elsevier.
- Nassour, J., Hénaff, P., Benouezdou, F., & Cheng, G. (2014). Multi-layered multi-pattern CPG for adaptive locomotion of humanoid robots. *Biological Cybernetics*, 108(3), 291–303.
- Nazifi, M. M., Beschorner, K. E., & Hur, P. (2017). Association between slip severity and muscle synergies of slipping. *Frontiers in Human Neuroscience*, 11.
- Neumann, G., Daniel, C., Paraschos, A., Kupcsik, A., & Peters, J. (2014). Learning modular policies for robotics. *Frontiers in computational neuroscience*, 8, 62.
- Nguyen, Q. & Sreenath, K. (2015). Optimal Robust Control for Bipedal Robots through Control Lyapunov Function based Quadratic Programs. In *Robotics: Science and Systems XI: Robotics: Science and Systems Foundation*.
- Niu, X., Tan, K., & Feng, L. (2024). Optimal Gait Control for a Tendon-driven Soft Quadruped Robot by Model-based Reinforcement Learning. arXiv:2406.07069.
- Nunes, P. F., Ostan, I., & Siqueira, A. A. G. (2020). Evaluation of Motor Primitive-Based Adaptive Control for Lower Limb Exoskeletons. *Frontiers in Robotics and AI*, 7, 575217.
- Nuxoll, A. M. & Laird, J. E. (2007). Extending cognitive architecture with episodic memory. In *Proceedings of the 22nd National Conference on Artificial Intelligence - Volume 2, AAAI'07* (pp. 1560–1565). Vancouver, British Columbia, Canada: AAAI Press.

- Ogihara, N. & Yamazaki, N. (2001). Generation of human bipedal locomotion by a bio-mimetic neuro-musculo-skeletal model. *Biological Cybernetics*, 84(1), 1–11.
- Oguz, O. S., Sari, O. C., Dinh, K. H., & Wollherr, D. (2017). Progressive stochastic motion planning for human-robot interaction. In *2017 26th IEEE International Symposium on Robot and Human Interactive Communication (RO-MAN)* (pp. 1194–1201). ISSN: 1944-9437.
- Okubo, Y., Sturnieks, D. L., Brodie, M. A., Duran, L., & Lord, S. R. (2019). Effect of Reactive Balance Training Involving Repeated Slips and Trips on Balance Recovery Among Older Adults: A Blinded Randomized Controlled Trial. *The journals of gerontology. Series A, Biological sciences and medical sciences*, 74(9), 1489–1496.
- Oliveira, A. S., Gizzi, L., Farina, D., & Kersting, U. G. (2014). Motor modules of human locomotion: influence of EMG averaging, concatenation, and number of step cycles. *Frontiers in Human Neuroscience*, 8.
- Oliveira, A. S., Gizzi, L., Ketabi, S., Farina, D., & Kersting, U. G. (2016). Modular Control of Treadmill vs Overground Running. *PLOS ONE*, 11(4), e0153307. Publisher: Public Library of Science.
- Olstad, B. (2015). *Muscle activation and kinematics in contemporary breaststroke swimming*. PhD Thesis.
- Orlovsky, G. N., Deliagina, T. G., & Grillner, S. (1999). *Neuronal control of locomotion: from mollusc to man*. New York: Oxford University Press.
- Ortega-Auriol, P. A., Besier, T. F., Byblow, W. D., & McMorland, A. J. C. (2018). Fatigue Influences the Recruitment, but Not Structure, of Muscle Synergies. *Frontiers in Human Neuroscience*, 12.
- O’Sullivan, L. & Gallwey, T. (2002). Upper-limb surface electro-myography at maximum supination and pronation torques: The effect of elbow and forearm angle. *Journal of electromyography and kinesiology : official journal of the International Society of Electrophysiological Kinesiology*, 12, 275–85.
- Owaki, D., Morikawa, L., & Ishiguro, A. (2012). Listen to body’s message: Quadruped robot that fully exploits physical interaction between legs. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems* (pp. 1950–1955). Vilamoura-Algarve, Portugal: IEEE.
- Pahič, R., Gams, A., Ude, A., & Morimoto, J. (2018). Deep Encoder-Decoder Networks for Mapping Raw Images to Dynamic Movement Primitives. In *2018 IEEE International Conference on Robotics and Automation (ICRA)* (pp. 1–6). Brisbane, QLD: IEEE.
- Pahič, R., Ridge, B., Gams, A., Morimoto, J., & Ude, A. (2020). Training of deep neural networks for the generation of dynamic movement primitives. *Neural Networks*, 127, 121–131.
- Palmer, S. E. (2015). Foreword. In J. Wagemans (Ed.), *The Oxford Handbook of Perceptual Organization*. Oxford University Press.

- Paraschos, A., Daniel, C., Peters, J. R., & Neumann, G. (2013). Probabilistic movement primitives. In *Advances in neural information processing systems* (pp. 2616–2624).
- Parijat, P. & Lockhart, T. E. (2008). Effects of lower extremity muscle fatigue on the outcomes of slip-induced falls. *Ergonomics*, 51(12), 1873–1884.
- Parijat, P. & Lockhart, T. E. (2012). Effects of moveable platform training in preventing slip-induced falls in older adults. *Annals of biomedical engineering*, 40(5), 1111–1121.
- Pastor, P., Hoffmann, H., Asfour, T., & Schaal, S. (2009). Learning and generalization of motor skills by learning from demonstration. In *2009 IEEE International Conference on Robotics and Automation* (pp. 763–768). Kobe: IEEE.
- Pastor, P., Kalakrishnan, M., Chitta, S., Theodorou, E., & Schaal, S. (2011a). Skill learning and task outcome prediction for manipulation. In *2011 IEEE International Conference on Robotics and Automation* (pp. 3828–3834). Shanghai, China: IEEE.
- Pastor, P., Kalakrishnan, M., Meier, F., Stulp, F., Buchli, J., Theodorou, E., & Schaal, S. (2013). From dynamic movement primitives to associative skill memories. *Robotics and Autonomous Systems*, 61(4), 351–361.
- Pastor, P., Kalakrishnan, M., Righetti, L., & Schaal, S. (2012). Towards associative skill memories. In *Humanoid Robots (Humanoids), 2012 12th IEEE-RAS International Conference on* (pp. 309–315).: IEEE.
- Pastor, P., Righetti, L., Kalakrishnan, M., & Schaal, S. (2011b). Online movement adaptation based on previous sensor experiences. In *2011 IEEE/RSJ International Conference on Intelligent Robots and Systems* (pp. 365–371). San Francisco, CA: IEEE.
- Paul, C., Bellotti, M., Jezernik, S., & Curt, A. (2005). Development of a human neuro-musculo-skeletal model for investigation of spinal cord injury. *Biological Cybernetics*, 93(3), 153–170.
- Pearson, K., Ekeberg, O., & Büschges, A. (2006). Assessing sensory function in locomotor systems using neuro-mechanical simulations. *Trends in Neurosciences*, 29(11), 625–631.
- Peña, G. G., Consoni, L. J., dos Santos, W. M., & Siqueira, A. A. G. (2019). Feasibility of an optimal EMG-driven adaptive impedance control applied to an active knee orthosis. *Robotics and Autonomous Systems*, 112, 98–108.
- Peng, X. B., Coumans, E., Zhang, T., Lee, T.-W., Tan, J., & Levine, S. (2020). Learning Agile Robotic Locomotion Skills by Imitating Animals. arXiv:2004.00784.
- Pervez, A., Mao, Y., & Lee, D. (2017). Learning deep movement primitives using convolutional neural networks. In *2017 IEEE-RAS 17th International Conference on Humanoid Robotics (Humanoids)* (pp. 191–197). Birmingham: IEEE.

- Peternel, L., Fang, C., Tsagarakis, N., & Ajoudani, A. (2019). A selective muscle fatigue management approach to ergonomic human-robot co-manipulation. *Robotics and Computer-Integrated Manufacturing*, 58, 69–79.
- Peternel, L., Noda, T., Petrič, T., Ude, A., Morimoto, J., & Babič, J. (2016). Adaptive Control of Exoskeleton Robots for Periodic Assistive Behaviours Based on EMG Feedback Minimisation. *PLOS ONE*, 11(2), e0148942. Publisher: Public Library of Science.
- Peters, J., Mülling, K., Kober, J., Nguyen-Tuong, D., & Krömer, O. (2011). Towards motor skill learning for robotics. In *Robotics Research* (pp. 469–482). Springer.
- Peters, J. & Schaal, S. (2008). Reinforcement learning of motor skills with policy gradients. *Neural Networks*, 21(4), 682–697.
- Petrič, T., Gams, A., Colasanto, L., Ijspeert, A. J., & Ude, A. (2018). Accelerated Sensorimotor Learning of Compliant Movement Primitives. *IEEE Transactions on Robotics*, 34(6), 1636–1642.
- Petrič, T., Gams, A., Zlajpah, L., & Ude, A. (2014). Online learning of task-specific dynamics for periodic tasks. In *2014 IEEE/RSJ International Conference on Intelligent Robots and Systems* (pp. 1790–1795). Chicago, IL, USA: IEEE.
- Phinyomark, A., Phukpattaranont, P., & Limsakul, C. (2012). Feature reduction and selection for EMG signal classification. *Expert Systems with Applications*, 39(8), 7420–7431.
- Plagemann, C., Stachniss, C., & Burgard, W. (2006). Efficient Failure Detection for Mobile Robots Using Mixed-Abstraction Particle Filters. In H. I. Christensen (Ed.), *European Robotics Symposium 2006*, volume 22 (pp. 93–107). Berlin/Heidelberg: Springer-Verlag. Series Title: Springer Tracts in Advanced Robotics.
- Poeppel, D. & Idsardi, W. (2011). Recognizing words from speech: The perception-action-memory loop. In G. Gaskell & P. Zwitserlood (Eds.), *Lexical Representation: A Multidisciplinary Approach* (pp. 171–196). De Gruyter Mouton.
- Potvin, J. R. & Fuglevand, A. J. (2017). A motor unit-based model of muscle fatigue. *PLOS Computational Biology*, 13(6).
- Poyil, A. T., Steuber, V., & Amirabdollahian, F. (2020). Influence of muscle fatigue on electromyogram–kinematic correlation during robot-assisted upper limb training. *Journal of Rehabilitation and Assistive Technologies Engineering*, 7.
- Pratt, J., Carff, J., Drakunov, S., & Goswami, A. (2006). Capture Point: A Step toward Humanoid Push Recovery. In *2006 6th IEEE-RAS International Conference on Humanoid Robots* (pp. 200–207). University of Genova, Genova, Italy: IEEE.

- Prochazka, A., Gillard, D., & Bennett, D. J. (1997). Positive force feedback control of muscles. *Journal of Neurophysiology*, 77(6), 3226–3236.
- Pugeault, N., Wörgötter, F., & Krüger, N. (2010). Visual primitives: Local, condensed, semantically rich visual descriptors and their applications in robotics. *International Journal of Humanoid Robotics*, 07(03), 379–405.
- Qiao, H., Chen, J., & Huang, X. (2022). A Survey of Brain-Inspired Intelligent Robots: Integration of Vision, Decision, Motion Control, and Musculoskeletal Systems. *IEEE Transactions on Cybernetics*, 52(10), 11267–11280. Conference Name: IEEE Transactions on Cybernetics.
- Quach, L., Galica, A. M., Jones, R. N., Procter-Gray, E., Manor, B., Hannan, M. T., & Lipsitz, L. A. (2011). The nonlinear relationship between gait speed and falls: The Maintenance of Balance, Independent Living, Intellect, and Zest in the Elderly of Boston Study. *Journal of the American Geriatrics Society*, 59(6), 1069–1073.
- QueiBer, J. F., Hammer, B., Ishihara, H., Asada, M., & Steil, J. J. (2018). Skill Memories for Parameterized Dynamic Action Primitives on the Pneumatically Driven Humanoid Robot Child Affetto. In *2018 Joint IEEE 8th International Conference on Development and Learning and Epigenetic Robotics (ICDL-EpiRob)* (pp. 39–45). Tokyo, Japan: IEEE.
- QueiBer, J. F. & Steil, J. J. (2018). Bootstrapping of Parameterized Skills Through Hybrid Optimization in Task and Policy Spaces. *Frontiers in Robotics and AI*, 5. Publisher: Frontiers Media SA.
- Raibert, M. H. (1986). *Legged robots that balance*. The MIT Press series in artificial intelligence. Cambridge, Mass.: MIT Press.
- Raman, N. & Maybank, S. J. (2016). Activity recognition using a supervised non-parametric hierarchical HMM. *Neurocomputing*, 199, 163–177.
- Ranganath, C. (2024). Episodic Memory. In M. J. Kahana & A. D. Wagner (Eds.), *The Oxford Handbook of Human Memory, Two Volume Pack: Foundations and Applications*. Oxford University Press.
- Ravichandar, H., Polydoros, A. S., Chernova, S., & Billard, A. (2020). Recent Advances in Robot Learning from Demonstration. *Annual Review of Control, Robotics, and Autonomous Systems*, 3(1), 297–330.
- Reaz, M. B. I., Hussain, M. S., & Mohd-Yasin, F. (2006). Techniques of EMG signal analysis: detection, processing, classification and applications. *Biological Procedures Online*, 8(1), 11–35.
- Reher, J. & Ames, A. D. (2021). Dynamic Walking: Toward Agile and Efficient Bipedal Robots. *Annual Review of Control, Robotics, and Autonomous Systems*, 4(1), 535–572.
- Reher, J., Cousineau, E. A., Hereid, A., Hubicki, C. M., & Ames, A. D. (2016). Realizing dynamic and efficient bipedal locomotion on the humanoid robot DURUS. In *2016 IEEE International Conference on Robotics and Automation (ICRA)* (pp. 1794–1801). Stockholm, Sweden: IEEE.

- Reher, J., Ma, W.-L., & Ames, A. D. (2019). Dynamic Walking with Compliance on a Cassie Bipedal Robot. In *2019 18th European Control Conference (ECC)* (pp. 2589–2595). Naples, Italy: IEEE.
- Reher, J. P., Hereid, A., Kolathaya, S., Hubicki, C. M., & Ames, A. D. (2020). Algorithmic Foundations of Realizing Multi-Contact Locomotion on the Humanoid Robot DURUS. In K. Goldberg, P. Abbeel, K. Bekris, & L. Miller (Eds.), *Algorithmic Foundations of Robotics XII*, volume 13 (pp. 400–415). Cham: Springer International Publishing. Series Title: Springer Proceedings in Advanced Robotics.
- Renjewski, D. & Seyfarth, A. (2012). Robots in human biomechanics—a study on ankle push-off in walking. *Bioinspiration & Biomimetics*, 7(3), 036005. Publisher: IOP Publishing.
- Rizzi, A. & Koditschek, D. (1994). Further progress in robot juggling: solvable mirror laws. In *Proceedings of the 1994 IEEE International Conference on Robotics and Automation* (pp. 2935–2940 vol.4).
- Robotics, S. (2019). Softbank robotics documentation, nao: User guide. Accessed: 2019-03-14.
- ROBOTIS (2019). Robotis e-manual: Darwin-op. Accessed: 2019-03-14.
- Romero-Garcés, A., Calderita, L., Martínez-Gómez, J., Bandera-Rubio, J. P., Marfil, R., Manso, L. J., Bustos, P., & Bandera-Rubio, A. J. (2015). The cognitive architecture of a robotic salesman. In *Conferencia de La Asociación Española Para La Inteligencia Artificial (CAEPIA'15)* Albacete, Spain.
- Rosenbaum, D. A., Cohen, R. G., Jax, S. A., Weiss, D. J., & Van Der Wel, R. (2007). The problem of serial order in behavior: Lashley's legacy. *Human Movement Science*, 26(4), 525–554.
- Rothfuss, J., Ferreira, F., Aksoy, E. E., Zhou, Y., & Asfour, T. (2018). Deep Episodic Memory: Encoding, Recalling, and Predicting Episodic Experiences for Robot Action Execution. *IEEE Robotics and Automation Letters*, 3(4), 4007–4014.
- Rozo, L., Calinon, S., Caldwell, D. G., Jimenez, P., & Torras, C. (2016). Learning Physical Collaborative Robot Behaviors From Human Demonstrations. *IEEE Transactions on Robotics*, 32(3), 513–527.
- Rückert, E. & d'Avella, A. (2013). Learned parametrized dynamic movement primitives with shared synergies for controlling robotic and musculoskeletal systems. *Frontiers in computational neuroscience*, 7, 138.
- Ruesch, J., Lopes, M., Bernardino, A., Hornstein, J., Santos-Victor, J., & Pfeifer, R. (2008). Multimodal saliency-based bottom-up attention a framework for the humanoid robot iCub. In *2008 IEEE International Conference on Robotics and Automation* (pp. 962–967).
- Rückert, E. A., Neumann, G., Toussaint, M., & Maass, W. (2013). Learned graphical models for probabilistic planning provide a new class of movement primitives. *Frontiers in Computational Neuroscience*, 6.
- Rückstieß, T., Sehnke, F., Schaul, T., Wierstra, D., Sun, Y., & Schmidhuber, J. (2010). Exploring Parameter Space in Reinforcement Learning. *Paladyn, Journal of Behavioral Robotics*, 1(1), 14–24.

- Saito, A., Tomita, A., Ando, R., Watanabe, K., & Akima, H. (2018). Muscle synergies are consistent across level and uphill treadmill running. *Scientific Reports*, 8.
- Samant, R., Behera, L., & Pandey, G. (2016). Adaptive learning of dynamic movement primitives through demonstration. In *2016 International Joint Conference on Neural Networks (IJCNN)* (pp. 1068–1075). Vancouver, BC, Canada: IEEE.
- Sardain, P. & Bessonnet, G. (2004). Forces acting on a biped robot. Center of pressure-zero moment point. *IEEE Transactions on Systems, Man, and Cybernetics - Part A: Systems and Humans*, 34(5), 630–637. Conference Name: IEEE Transactions on Systems, Man, and Cybernetics - Part A: Systems and Humans.
- Saveriano, M., Abu-Dakka, F. J., Kramberger, A., & Peternel, L. (2021). Dynamic Movement Primitives in Robotics: A Tutorial Survey. Publisher: arXiv Version Number: 1.
- Schaal, S. (1999). Is imitation learning the route to humanoid robots? *Trends in Cognitive Sciences*, 3(6), 233–242.
- Schaal, S., Atkeson, C. G., & Sternad, D. (1996). One-Handed Juggling: A Dynamical Approach to a Rhythmic Movement Task. *Journal of Motor Behavior*, 28(2), 165–183.
- Schaal, S., Mohajerian, P., & Ijspeert, A. (2007). Dynamics systems vs. optimal control—a unifying view. *Progress in Brain Research*, 165, 425–445.
- Schaal, S., Peters, J., Nakanishi, J., & Ijspeert, A. (2005). Learning movement primitives. In *Robotics research. the eleventh international symposium* (pp. 561–572).: Springer.
- Scherzinger, S., Becker, P., Roennau, A., & Dillmann, R. (2023). Motion Macro Programming on Assistive Robotic Manipulators: Three Skill Types for Everyday Tasks. In *2023 20th International Conference on Ubiquitous Robots (UR)* (pp. 792–799).
- Schmidt, R. A., Lee, T. D., Winstein, C. J., Wulf, G., & Zelaznik, H. N. (2018). The Learning Process. In *Motor Control and Learning: A Behavioral Emphasis, 6th Edition*, Motor Control and Learning: A Behavioral Emphasis, 5th Ed (pp. 529–563). Champaign, IL, US: Human Kinetics, sixth edition edition.
- Sella, G. E. (2021). Agonism and antagonism of the muscles of the elbow joint: An semg approach. *Biofeedback*, 49(4), 89–94.
- Semini, C., Barasuol, V., Boaventura, T., Frigerio, M., & Buchli, J. (2016). Is Active Impedance the Key to a Breakthrough for Legged Robots? In M. Inaba & P. Corke (Eds.), *Robotics Research: The 16th International Symposium ISRR*, Springer Tracts in Advanced Robotics (pp. 3–19). Cham: Springer International Publishing.

- Seok, S., Wang, A., Chuah, M. Y., Hyun, D. J., Lee, J., Otten, D. M., Lang, J. H., & Kim, S. (2015). Design Principles for Energy-Efficient Legged Locomotion and Implementation on the MIT Cheetah Robot. *IEEE/ASME Transactions on Mechatronics*, 20(3), 1117–1129. Conference Name: IEEE/ASME Transactions on Mechatronics.
- Seyde, T., Shrivastava, A., Engelsberger, J., Bertrand, S., Pratt, J., & Griffin, R. J. (2018). Inclusion of Angular Momentum During Planning for Capture Point Based Walking. In *2018 IEEE International Conference on Robotics and Automation (ICRA)* (pp. 1791–1798). Brisbane, QLD: IEEE.
- Shams, L. & von der Malsburg, C. (2002). Acquisition of visual shape primitives. *Vision Research*, 42(17), 2105–2122.
- Shankar, T., Tulsiani, S., Pinto, L., & Gupta, A. (2020). Discovering motor programs by recomposing demonstrations.
- Shiavi, R. (1990). Quantitative representation of electromyographic patterns generated during human locomotion. *IEEE Engineering in Medicine and Biology Magazine*, 9(1), 58–60.
- Shourijeh, M. S., Flaxman, T. E., & Benoit, D. L. (2016). An approach for improving repeatability and reliability of non-negative matrix factorization for muscle synergy analysis. *Journal of Electromyography and Kinesiology*, 26, 36–43.
- Sigaud, O. & Stulp, F. (2019). Policy search in continuous action domains: An overview. *Neural Networks*, 113, 28–40.
- Silver, D., Huang, A., Maddison, C. J., Guez, A., Sifre, L., van den Driessche, G., Schrittwieser, J., Antonoglou, I., Panneershelvam, V., Lanctot, M., Dieleman, S., Grewe, D., Nham, J., Kalchbrenner, N., Sutskever, I., Lillicrap, T., Leach, M., Kavukcuoglu, K., Graepel, T., & Hassabis, D. (2016). Mastering the game of Go with deep neural networks and tree search. *Nature*, 529(7587), 484–489.
- Silver, D., Singh, S., Precup, D., & Sutton, R. S. (2021). Reward is enough. *Artificial Intelligence*, 299.
- Singh, A., Yang, L., & Levine, S. (2017). GPLAC: Generalizing Vision-Based Robotic Skills Using Weakly Labeled Images. In *2017 IEEE International Conference on Computer Vision (ICCV)* (pp. 5852–5861). Venice: IEEE.
- Singh, R. E., Iqbal, K., White, G., & Hutchinson, T. E. (2018). A Systematic Review on Muscle Synergies: From Building Blocks of Motor Behavior to a Neurorehabilitation Tool. *Applied Bionics and Biomechanics*, 2018, 3615368.
- Singh, T. & Latash, M. L. (2011). Effects of muscle fatigue on multi-muscle synergies. *Experimental brain research*, 214(3), 335–350.
- Smaldone, F. M., Scianca, N., Lanari, L., & Oriolo, G. (2022). From Walking to Running: 3D Humanoid Gait Generation via MPC. *Frontiers in Robotics and AI*, 9. Publisher: Frontiers Media S.A.

- Smale, K. B., Shourijeh, M. S., & Benoit, D. L. (2016). Use of muscle synergies and wavelet transforms to identify fatigue during squatting. *Journal of Electromyography and Kinesiology*, 28, 158–166.
- Smith, S. W. (1999). *The scientist and engineer's guide to digital signal processing*. San Diego (Calif.): California Technical Pub., 2nd edition edition. OCLC: 493473234.
- Sousa, A. & Tavares, J. (2012). Surface electromyographic amplitude normalization methods: A review. In *Electromyography: New Developments, Procedures and Applications* (pp. 85–102). Journal Abbreviation: Electromyography: New Developments, Procedures and Applications.
- Spence, C. (2015). Cross-modal perceptual organization. In J. Wagemans (Ed.), *The Oxford Handbook of Perceptual Organization*. Oxford University Press.
- Sprowitz, A., Tuleu, A., Vespignani, M., Ajallooeian, M., Badri, E., & Ijspeert, A. J. (2013). Towards dynamic trot gait locomotion: Design, control, and experiments with Cheetah-cub, a compliant quadruped robot. *The International Journal of Robotics Research*, 32(8), 932–950. Publisher: SAGE Publications Ltd STM.
- Squire, L. R. (1994). Declarative and Nondeclarative Memory: Multiple Brain Systems Supporting Learning and Memory. In D. L. Schacter & E. Tulving (Eds.), *Memory Systems 1994* (pp. 200–228). The MIT Press.
- Sreenath, K., Park, H.-W., Poulakakis, I., & Grizzle, J. (2013). Embedding active force control within the compliant hybrid zero dynamics to achieve stable, fast running on MABEL. *The International Journal of Robotics Research*, 32(3), 324–345.
- Sridhar, S., Khamaj, A., & Asthana, M. K. (2023). Cognitive neuroscience perspective on memory: Overview and summary. *Frontiers in Human Neuroscience*, 17.
- Stegeman, D. & 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, 8–12.
- Stickgold, R. (2005). Sleep-dependent memory consolidation. *Nature*, 437(7063), 1272–1278. Number: 7063 Publisher: Nature Publishing Group.
- Stiles, N. R. B. & Shimojo, S. (2015a). Auditory Sensory Substitution is Intuitive and Automatic with Texture Stimuli. *Scientific Reports*, 5(1), 15628.
- Stiles, N. R. B. & Shimojo, S. (2015b). Sensory substitution: A new perceptual experience. In J. Wagemans (Ed.), *The Oxford Handbook of Perceptual Organization*. Oxford University Press.
- Strang, A. J. & Berg, W. P. (2007). Fatigue-induced adaptive changes of anticipatory postural adjustments. *Experimental Brain Research*, 178(1), 49–61.

- Stulp, F. & Sigaud, O. (2012). Policy Improvement Methods: Between Black-Box Optimization and Episodic Reinforcement Learning.
- Stulp, F. & Sigaud, O. (2013). Robot Skill Learning: From Reinforcement Learning to Evolution Strategies. *Paladyn, Journal of Behavioral Robotics*, 4(1), 49–61. Publisher: De Gruyter Open Access Section: Paladyn.
- Stulp, F., Theodorou, E. A., & Schaal, S. (2012). Reinforcement Learning With Sequences of Motion Primitives for Robust Manipulation. *IEEE Transactions on Robotics*, 28(6), 1360–1370. Conference Name: IEEE Transactions on Robotics.
- Sun, R. (2007). The importance of cognitive architectures: An analysis based on CLARION. *Journal of Experimental & Theoretical Artificial Intelligence*, 19(2), 159–193.
- Sünderhauf, N., Brock, O., Scheirer, W., Hadsell, R., Fox, D., Leitner, J., Upcroft, B., Abbeel, P., Burgard, W., Milford, M., & Corke, P. (2018). The limits and potentials of deep learning for robotics. *The International Journal of Robotics Research*, 37(4-5), 405–420.
- Sutherland, D. (1994). Kinematics of normal human walking. *Human walking*.
- Sutton, R. (2019). The bitter lesson. <http://www.incompleteideas.net/IncIdeas/BitterLesson.html>. Accessed: 2024-06-14.
- Sörnmo, L. & Laguna, P. (2005). The Electromyogram. In *Bioelectrical Signal Processing in Cardiac and Neurological Applications* (pp. 337–410). Elsevier.
- Søgaard, K., Gandevia, S. C., Todd, G., Petersen, N. T., & Taylor, J. L. (2006). The effect of sustained low-intensity contractions on supraspinal fatigue in human elbow flexor muscles. *The Journal of Physiology*, 573(2), 511–523.
- Takenaka, T., Matsumoto, T., & Yoshiike, T. (2009). Real time motion generation and control for biped robot -1st report: Walking gait pattern generation. In *2009 IEEE/RSJ International Conference on Intelligent Robots and Systems* (pp. 1084–1091). St. Louis, MO, USA: IEEE.
- Tallent, J., Woodhead, A., Frazer, A. K., Hill, J., Kidgell, D. J., & Howatson, G. (2021). Corticospinal and spinal adaptations to motor skill and resistance training: Potential mechanisms and implications for motor rehabilitation and athletic development. *European Journal of Applied Physiology*, 121(3), 707–719.
- Taylor, J. L., Amann, M., Duchateau, J., Meeusen, R., & Rice, C. L. (2016). Neural Contributions to Muscle Fatigue: From the Brain to the Muscle and Back Again. *Medicine & Science in Sports & Exercise*, 48(11), 2294–2306.
- Teixeira, C., Costa, L., & Santos, C. (2014). Biped locomotion - Improvement and adaptation. In *2014 IEEE International Conference on Autonomous Robot Systems and Competitions (ICARSC)* (pp. 110–115).

- Tenorth, M. & Beetz, M. (2009). KNOWROB — knowledge processing for autonomous personal robots. In *2009 IEEE/RSJ International Conference on Intelligent Robots and Systems* (pp. 4261–4266). ISSN: 2153-0866.
- Thomas, S. J., Castillo, G. C., Topley, M., & Paul, R. W. (2023). The effects of fatigue on muscle synergies in the shoulders of baseball players. *Sports Health*, 15(2), 282–289. PMID: 35492023.
- Thrun, S. (1995). Is learning the n-th thing any easier than learning the first? In *Proceedings of the 8th International Conference on Neural Information Processing Systems, NIPS'95* (pp. 640–646). Cambridge, MA, USA: MIT Press.
- Trkov, M., Wu, S., Chen, K., Yi, J., Liu, T., & Zhao, Q. (2017). Design of a Robotic Knee Assistive Device (ROKAD) for Slip-Induced Fall Prevention during Walking. *IFAC-PapersOnLine*, 50(1), 9802–9807.
- Tulving, E. (1984). Précis of Elements of episodic memory. *Behavioral and Brain Sciences*, 7(2), 223–238.
- Ude, A., Gams, A., Asfour, T., & Morimoto, J. (2010). Task-Specific Generalization of Discrete and Periodic Dynamic Movement Primitives. *IEEE Transactions on Robotics*, 26(5), 800–815.
- Ude, A., Nemec, B., Petrič, T., & Morimoto, J. (2014). Orientation in Cartesian space dynamic movement primitives. In *2014 IEEE International Conference on Robotics and Automation (ICRA)* (pp. 2997–3004). Hong Kong, China: IEEE.
- USBJI (2015). *United States Bone and Joint Initiative: The Burden of Musculoskeletal Diseases in the United States (BMUS)*. Rosemont, IL, 3. edition.
- Van der Noot, N. & Barrea, A. (2014). Zero-moment point on a bipedal robot under bio-inspired walking control. In *Electrotechnical Conference (MELECON), 2014 17th IEEE Mediterranean* (pp. 85–90).: IEEE.
- Varnet, L., Knoblauch, K., Serniclaes, W., Meunier, F., & Hoen, M. (2015). A Psychophysical Imaging Method Evidencing Auditory Cue Extraction during Speech Perception: A Group Analysis of Auditory Classification Images. *PLOS ONE*, 10(3).
- Vernon, D. (2022). Cognitive Architectures. In A. Cangelosi & M. Asada (Eds.), *Cognitive Robotics*. The MIT Press.
- Verrelst, B., Ham, R. V., Vanderborght, B., Daerden, F., Lefeber, D., & Vermeulen, J. (2005). The Pneumatic Biped “Lucy” Actuated with Pleated Pneumatic Artificial Muscles. *Autonomous Robots*, 18(2), 201–213.
- Verschure, P. F. M. J. & Althaus, P. (2003). A real-world rational agent: Unifying old and new AI. *Cognitive Science*, 27(4), 561–590.
- Vukobratović, M. & Stepanenko, J. (1972). On the stability of anthropomorphic systems. *Mathematical biosciences*, 15(1-2), 1–37.

- Vukobratović, M. & Borovac, B. (2004). Zero-moment point — thirty five years of its life. *International Journal of Humanoid Robotics*, 01(01), 157–173.
- Wang, C., Xie, G., Wang, L., & Cao, M. (2011). CPG-Based Locomotion Control of a Robotic Fish: Using Linear Oscillators and Reducing Control Parameters via PSO. *International journal of innovative computing information and control*, 7(7B), 4237–4249.
- Wang, L., Wang, Y., Ma, A., Ma, G., Ye, Y., Li, R., & Lu, T. (2018). A Comparative Study of EMG Indices in Muscle Fatigue Evaluation Based on Grey Relational Analysis during All-Out Cycling Exercise. *BioMed Research International*, 2018, 1–8.
- Wang, Y. & Wang, H. (2020). Gait Learning for Lower Limb Rehabilitation Exoskeleton using Dynamical Movement Primitives based on General Regression Neural Networks. In *2020 Chinese Automation Congress (CAC)* (pp. 7046–7051). Shanghai, China: IEEE.
- Watanabe, K., Tajima, A., & Izumi, K. (2008). Locomotion pattern generation of semi-looper type robots using central pattern generators based on van der Pol oscillators. In *2008 6th IEEE International Conference on Industrial Informatics* (pp. 377–382). Daejeon, South Korea: IEEE.
- Weavil, J. C. & Amann, M. (2018). Corticospinal excitability during fatiguing whole body exercise. In *Progress in Brain Research*, volume 240 (pp. 219–246). Elsevier.
- Weitschat, R. & Aschemann, H. (2018). Safe and Efficient Human–Robot Collaboration Part II: Optimal Generalized Human-in-the-Loop Real-Time Motion Generation. *IEEE Robotics and Automation Letters*, 3(4), 3781–3788. Conference Name: IEEE Robotics and Automation Letters.
- Werner, A., Lampariello, R., & Ott, C. (2014). Trajectory optimization for walking robots with series elastic actuators. In *53rd IEEE Conference on Decision and Control* (pp. 2964–2970). ISSN: 0191-2216.
- Westervelt, E., Grizzle, J., & Koditschek, D. (2003). Hybrid zero dynamics of planar biped walkers. *IEEE Transactions on Automatic Control*, 48(1), 42–56.
- Wierzbicka, M., Wiegner, A., & Shahani, B. (1986). Role of agonist and antagonist muscles in fast arm movements in man. *Experimental Brain Research*, 63(2).
- Willingham, D. B. (1998). A neuropsychological theory of motor skill learning. *Psychological Review*, 105(3), 558–584.
- Wilson, M. (2002). Six views of embodied cognition. *Psychonomic Bulletin & Review*, 9(4), 625–636.
- Wolpert, D. M., Diedrichsen, J., & Flanagan, J. R. (2011). Principles of sensorimotor learning. *Nature Reviews Neuroscience*, 12(12), 739–751. Number: 12 Publisher: Nature Publishing Group.

- Wu, G., Siegler, S., Allard, P., Kirtley, C., Leardini, A., Rosenbaum, D., Whittle, M., D'Lima, D. D., Cristofolini, L., Witte, H., Schmid, O., & Stokes, I. (2002). ISB recommendation on definitions of joint coordinate system of various joints for the reporting of human joint motion—part I: ankle, hip, and spine. *Journal of Biomechanics*, 35(4), 543–548.
- Yamane, K., Revfi, M., & Asfour, T. (2013). Synthesizing object receiving motions of humanoid robots with human motion database. In *2013 IEEE International Conference on Robotics and Automation* (pp. 1629–1636). ISSN: 1050-4729.
- Yamane, K., Yamaguchi, Y., & Nakamura, Y. (2011). Human motion database with a binary tree and node transition graphs. *Autonomous Robots*, 30(1), 87–98.
- Yang, C., Zeng, C., Cong, Y., Wang, N., & Wang, M. (2019). A Learning Framework of Adaptive Manipulative Skills From Human to Robot. *IEEE Transactions on Industrial Informatics*, 15(2), 1153–1161.
- Yang, C.-Y., Gamborino, E., Fu, L.-C., & Chang, Y.-L. (2022). A Brain-Inspired Self-Organizing Episodic Memory Model for a Memory Assistance Robot. *IEEE Transactions on Cognitive and Developmental Systems*, 14(2), 617–628.
- Yang, F., Anderson, F. C., & Pai, Y.-C. (2008). Predicted threshold against backward balance loss following a slip in gait. *Journal of biomechanics*, 41(9), 1823–1831.
- Yaserifar, M., fallah mohammadi, Z., Hosseinienejad, E., Esmaili Paeen Afrakoti, I., Meijer, K., & Boonstra, T. (2021). *The effects of long-term exercise training on the neural control of walking*.
- Yu, J., Wu, Z., Wang, M., & Tan, M. (2016). CPG Network Optimization for a Biomimetic Robotic Fish via PSO. *IEEE Transactions on Neural Networks and Learning Systems*, 27(9), 1962–1968. Conference Name: IEEE Transactions on Neural Networks and Learning Systems.
- Yun, Y., Dancausse, S., Esmatloo, P., Serrato, A., Merring, C. A., Agarwal, P., & Deshpande, A. D. (2017). Maestro: An EMG-driven assistive hand exoskeleton for spinal cord injury patients. In *2017 IEEE International Conference on Robotics and Automation (ICRA)* (pp. 2904–2910).
- Zaier, R. & Kanda, S. (2007). Piecewise-Linear Pattern Generator and Reflex System for Humanoid Robots. In *Proceedings 2007 IEEE International Conference on Robotics and Automation* (pp. 2188–2195). ISSN: 1050-4729.
- Zandvoort, C. S., Daffertshofer, A., & Dominici, N. (2022). Cortical contributions to locomotor primitives in toddlers and adults. *iScience*, 25(10).
- Zeng, C., Li, S., Fang, B., Chen, Z., & Zhang, J. (2022). Generalization of Robot Force-Relevant Skills Through Adapting Compliant Profiles. *IEEE Robotics and Automation Letters*, 7(2), 1055–1062.
- Zhang, L.-Q., Shiavi, R., Hunt, M., & Chen, J.-J. (1991). Clustering analysis and pattern discrimination of EMG linear envelopes. *IEEE Transactions on Biomedical Engineering*, 38(8), 777–784.

- Zhang, X., Wang, S., Ma, S., Liu, S., & Gao, W. (2013). Entropy of primitive: A top-down methodology for evaluating the perceptual visual information. In *2013 Visual Communications and Image Processing (VCIP)* (pp. 1–6). Kuching, Malaysia: IEEE.
- Zhang, Y., Chen, S., Cao, W., Guo, P., Gao, D., Wang, M., Zhou, J., & Wang, T. (2021). MFFNet: Multi-dimensional feature fusion network based on attention mechanism for semg analysis to detect muscle fatigue. *Expert Systems with Applications*, 185.
- Zhou, C., Wang, B., Zhu, Q., & Wu, J. (2016). An online gait generator for quadruped walking using motor primitives. *International Journal of Advanced Robotic Systems*, 13(6). Publisher: SAGE Publications.
- Zhou, X. & Bi, S. (2012). A survey of bio-inspired compliant legged robot designs. *Bioinspiration & Biomimetics*, 7(4), 041001. Publisher: IOP Publishing.
- Zhou, Y., Gao, J., & Asfour, T. (2019). Learning Via-Point Movement Primitives with Inter- and Extrapolation Capabilities. In *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)* (pp. 4301–4308). Macau, China: IEEE.
- Zlatoff, N., Tellez, B., & Baskurt, A. (2008). Combining local belief from low-level primitives for perceptual grouping. *Pattern Recognition*, 41(4), 1215–1229.
- Zmigrod, S. & Hommel, B. (2013). Feature Integration across Multimodal Perception and Action: A Review. *Multisensory Research*, 26(1-2), 143–157.
- Åhsberg, E., Garnberale, F., & Kjellberg, A. (1997). Perceived quality of fatigue during different occupational tasks Development of a questionnaire. *International Journal of Industrial Ergonomics*, 20(2), 121–135.
- Öberg, T., Sandsjö, L., & Kadefors, R. (1994). Subjective and objective evaluation of shoulder muscle fatigue. *Ergonomics*, 37(8), 1323–1333.
- Žabkar, J. & Leonardis, A. (2016). Motor memory: Representation, learning and consolidation. *Biologically Inspired Cognitive Architectures*, 16, 64–74.