

Politecnico di Torino

Master's Degree in Biomedical Engineering

Academic Year 2025/2026

Graduation Session: September 2026

A Markerless-Driven Musculoskeletal Model of the Hand and Wrist: sEMG Validation, EMG-Informed Feasibility and Mechanical Characterisation

Master's Degree Thesis carried out at the Translational Neural Engineering Laboratory, École Polytechnique Fédérale de Lausanne





Supervisors

Prof. Danilo Demarchi

Prof. Silvestro Micera

Dr. Vincent Mendez

Candidate

Michele Gurra

A Markerless-Driven Musculoskeletal Model of the Hand and Wrist: sEMG Validation, EMG-Informed Feasibility and Mechanical Characterisation

Michele Gurra

Abstract

Musculoskeletal models estimate muscle forces non-invasively, but they all rest on a redundancy: many muscles act on few degrees of freedom, so infinitely many activation patterns reproduce the same motion. The criterion that selects one of them has been tested against physiological measurements elsewhere in the body, and barely at the hand. The methods that bring electromyography into the estimation have taken shape on the lower limb. The inverse-dynamics procedure used here had been demonstrated on this hand model on synthetic trajectories, where it reproduces the imposed motion accurately, but its activations had never been confronted with a measurement.

This thesis supplies that term of comparison through a reproducible pipeline driven by markerless kinematics alone. Joint angles for four wrist movements and two middle-finger movements, recorded from seven participants, were reconstructed from two cameras and used to drive an existing MuJoCo-based inverse-dynamics solver on the MyoHand model distributed with MyoSuite, while six surface channels recorded four wrist muscles and the third-finger flexor and extensor compartments. The unconstrained estimates were first compared against the recorded envelopes. Recorded activity was then introduced into the same solver as a soft constraint, posed as a feasibility question rather than a claim of accuracy: once the recording is an input it can no longer validate the solution, and what remains testable is whether recorded activity and recorded motion can be satisfied at once, and at what cost. A leave-one-muscle-out check followed, and the mechanics of the model were characterised through joint torque.

The unconstrained estimates reproduced the leading muscle of the task in shape in four tasks out of six, with correlations between 0.61 and 0.86, while the coefficient of determination was negative in thirty-five of thirty-six muscle–task pairs: the timing of an estimated activation carries information, its amplitude does not. In the feasibility test, agreement with the reference rose systematically, as its construction guarantees; the informative result is its price. At the wrist the movement was reproduced as before, the cost falling on joints merely held still. At the finger the movement itself was spent: at the distal interphalangeal joint the error grew from 11.3 to 17.6 degrees in flexion and from 3.4 to 15.0 degrees in extension, with the coefficient of determination falling from -0.43 to -1.18 and from -0.02 to -11.4, the residual in extension exceeding the excursion asked of the joint. Withholding one recorded muscle from the reference left the movement untouched, but that muscle was recovered only partially: the wrist extensors were reconstructed whatever their

role, whereas the flexors were recovered less and less as their involvement grew, failing in the movements they lead.

The characterisation accounts for the asymmetry: the six electrodes carry roughly 78 to 94 per cent of the torque on the degrees of freedom the task moves at the wrist, and about half of it at the finger, with little difference between the two solutions; the rest is produced by muscles no electrode observes.

The two requirements can therefore be met simultaneously at the wrist, and not at the finger. What such a model can achieve is set by how much of the mechanics the electrodes reach rather than by the solver; whether it can be made quantitatively reliable at the finger depends on how the muscles it does not observe are accounted for.

Table of Contents

1 Introduction **1**

1.1 Context and motivation 1
1.2 Problem statement 2
1.3 Objectives and research questions 2
1.4 Contributions 3
1.5 Thesis outline 4

2 State of the art **5**

2.1 Motion capture and kinematics estimation 5
2.2 Electromyography 6
2.3 Musculoskeletal modelling 8
2.4 Inverse dynamics and the muscle redundancy problem 10
2.5 Integrating EMG into MSK models 11
2.6 Positioning of this work 13

3 Materials and Methods **16**

3.1 From recordings to muscle activation estimates 18
3.1.1 Participants and experimental protocol 18
3.1.2 Data acquisition 21
3.1.3 Signal preprocessing 25
3.1.4 EMG envelopes and MVC normalisation 26
3.1.5 Detrending and repetition segmentation 28
3.1.6 Group averages of muscle activity 32
3.1.7 Kinematic inputs for the MSK model 35
3.1.8 Musculoskeletal model 40
3.1.9 Inverse-dynamics muscle-activation estimation (MjID) 42
3.1.10 Subject-specific parameter optimisation 44
3.2 Validation against sEMG 46
3.2.1 Reproduction of the imposed trajectories 46
3.2.2 Modelled activations and measured envelopes 48
3.2.3 Time normalisation and alignment 48
3.2.4 Metrics 49
3.2.5 Aggregation and statistics 51
3.3 EMG-informed MjID: a feasibility test 52

3.3.1 Measured envelopes as an input . . . 52
3.3.2 Construction of the control reference . . . 53
3.3.3 Supplying the reference to the solver . . . 54
3.3.4 Re-optimisation of the parameters . . . 54
3.3.5 Comparisons and metrics . . . 55
3.3.6 Robustness to a missing constraint: the LOMO test . . . 57
3.4 Mechanical characterisation of the model . . . 58
3.4.1 Joint torque as the measure of what a muscle does . . . 58
3.4.2 Joint torque read from the simulation . . . 59
3.4.3 Anatomical capacity of the model . . . 59
3.4.4 Partition between movement and stabilisation . . . 60
3.4.5 Share of the torque carried by the recorded muscles . . . 60
3.4.6 Co-contraction . . . 61

4 Results 63
4.1 The recorded signals . . . 63
4.1.1 Signal quality . . . 63
4.1.2 Effect of the detrending . . . 66
4.1.3 Group averages of the envelopes . . . 67
4.2 Validation against sEMG . . . 69
4.2.1 Reproduction of the imposed trajectories . . . 69
4.2.2 Agreement in shape . . . 72
4.2.3 Agreement in amplitude . . . 74
4.3 EMG-informed MjID: a feasibility test . . . 75
4.3.1 Comparison of the configurations . . . 75
4.3.2 Agreement with the control reference . . . 77
4.3.3 Cost of the agreement . . . 79
4.3.4 Robustness check: the LOMO test . . . 81
4.4 Mechanical characterisation of the model . . . 83
4.4.1 Anatomical capacity . . . 84
4.4.2 Movement and stabilisation . . . 85
4.4.3 Agonists, antagonists and the recorded muscles . . . 87
4.4.4 Co-contraction . . . 91

5 Discussion 94
5.1 Validation against sEMG: timing without amplitude . . . 94
5.2 EMG-informed MjID: agreement, cost and redundancy . . . 96
5.3 Mechanical characterisation: the coverage of the six electrodes . . . 97
5.4 Methodological choices and limitations . . . 99
5.5 Implications . . . 100

6 Conclusions and future work 101
6.1 Conclusions . . . 101
6.2 Future work . . . 102

A Reproducibility material **103**

B Supporting data **112**

Bibliography **119**

List of Figures

2.1 Surface electromyography records the summed activity of several motor units . . . 7
2.2 Hill-type musculotendon model . . . 9
2.3 Muscle redundancy . . . 10
2.4 Approaches to the muscle redundancy problem . . . 12

3.1 Overview of the pipeline . . . 17
3.2 Temporal structure of the acquisition session . . . 19
3.3 Placement of the six sEMG probes on the forearm . . . 22
3.4 Acquisition setup, seen from the side and from above . . . 23
3.5 The MyoHand MSK model . . . 40
3.6 The three stages of the MjID method . . . 42

4.1 Effect of the detrending on a kinematic channel . . . 66
4.2 Group averages of the sEMG envelopes . . . 67
4.3 Amplitude in the active phase and differences between muscles . . . 68
4.4 Estimated activations and measured envelopes . . . 73
4.5 Comparison of the configurations . . . 76
4.6 Effect of the constraint, muscle by muscle . . . 78
4.7 Cost of the constraint on the tracking . . . 79
4.8 Cost on the driven finger joints . . . 80
4.9 Overview of the LOMO test . . . 81
4.10 Recovery of the withheld muscle against its involvement . . . 83
4.11 Anatomical capacity of the model . . . 85
4.12 Movement and stabilisation . . . 86
4.13 Torque over the cycle . . . 88
4.14 Share on the driven DOF . . . 89
4.15 Share in wrist stabilisation . . . 90
4.16 Share in finger stabilisation . . . 91
4.17 Co-contraction . . . 91

B.1 A trial whose repetitions cannot be located on kinematics . . . 114
B.2 Agreement between repetitions and between their averages . . . 115
B.3 Controls of the whole model, wrist flexion . . . 116
B.4 Controls of the whole model, wrist extension . . . 116

B.5 Controls of the whole model, radial deviation 117
B.6 Controls of the whole model, ulnar deviation 117
B.7 Controls of the whole model, finger flexion 118
B.8 Controls of the whole model, finger extension 118

List of Tables

2.1 Main technologies for measuring hand and wrist kinematics 6

3.1 Voluntary contractions used to locate each muscle 22

4.1 SNR ratio of the sEMG envelopes 64
4.2 SNR ratio of the kinematic channels 65
4.3 SNR ratio of the kinematic channels in fist closure 66
4.4 Reproduction of the imposed trajectories 70
4.5 Reproduction of the imposed trajectories, joint by joint 71
4.6 Correlation between the estimated activations and the measured envelopes 72
4.7 Coefficient of determination and gain between the estimated activations and the measured envelopes 74
4.8 Agonist and antagonist groups 92

A.1 Normalisation reference per channel 104
A.2 Segmentation settings per participant and task 106
A.3 Corrections applied to the resting posture 107
A.4 Joint ranges of the MyoHand model 108
A.5 Search bounds of the parameter optimisation 109
A.6 Settings of the differential evolution algorithm 109
A.7 Optimised parameter values 109
A.8 Optimised parameter values, EMG-informed configuration 111

B.1 Number of repetitions per participant and task 112

Acronyms

MSK	musculoskeletal.
EMG	electromyography.
sEMG	surface electromyography.
HD-sEMG	high-density surface electromyography.
IMU	inertial measurement unit.
MVC	maximum voluntary contraction.
CEDE	Consensus for Experimental Design in Electromyography.
MjID	MuJoCo-based inverse dynamics.
RMSE	root-mean-square error.
nRMSE	normalised root-mean-square error.
RMS	root-mean-square.
SNR	signal-to-noise ratio.
ROM	range of motion.
IQR	interquartile range.
DOF	degree of freedom.
LOMO	leave-one-muscle-out.
FDS	flexor digitorum superficialis.
FCR	flexor carpi radialis.
FCU	flexor carpi ulnaris.
EDC	extensor digitorum communis.
ECRL	extensor carpi radialis longus.
ECU	extensor carpi ulnaris.
FDP	flexor digitorum profundus.
EPL	extensor pollicis longus.
MCP	metacarpophalangeal.
PIP	proximal interphalangeal.
DIP	distal interphalangeal.

Chapter 1

Introduction

1.1 Context and motivation

Understanding how individual forearm and hand muscles are recruited during movement is a long-standing goal in biomechanics, with direct relevance to several fields. In clinical assessment and rehabilitation, the pattern of muscle recruitment reflects motor control and its impairment, and is used to characterise neuromuscular function and to monitor recovery [1]. In prosthetics and human–machine interfaces, myoelectric control decodes the user's motor intent from muscle signals, so that a better understanding of muscle activity translates into more reliable control [2]. In ergonomics and injury prevention, knowing which muscles and tendons bear the load during repetitive manual tasks helps to explain and mitigate musculoskeletal disorders [3]. In all of these settings the quantity of interest is not the movement itself, which can be observed, but the underlying muscular activity that produces it.

The hand is a particularly important and challenging target. It is central to most activities of daily living and a frequent focus of rehabilitation and prosthetic design, yet it is also one of the most mechanically complex effectors of the human body: a large number of degrees of freedom (DOFs) are actuated by many muscles, comprising both extrinsic muscles located in the forearm and intrinsic muscles within the hand itself, several of which are small or deep [4]. This combination of functional importance and mechanical complexity makes the hand both a compelling application and a stringent test case for any method that attempts to recover muscle activity from movement.

Measuring that activity directly, however, is difficult: surface electromyography (sEMG) can only be placed over a limited number of superficial muscles, while the deep and intrinsic muscles of the hand remain largely inaccessible to non-invasive recording [5]. Movement, on the other hand, has become increasingly easy to capture. Markerless motion capture from ordinary video now provides movement kinematics without markers or dedicated laboratory equipment, considerably lowering the barrier to acquiring movement data outside specialised settings [6]. This asymmetry, abundant kinematics and scarce muscle measurements, motivates an indirect route: if muscle activity can be inferred from the observed movement through modelling,

recruitment could be estimated even for muscles that cannot be recorded.

1.2 Problem statement

The sEMG signal is an indirect and partial observation of muscle activity. It is restricted to superficial muscles, is affected by crosstalk between neighbouring sources, and is sensitive to electrode placement and to the normalisation procedure, so that its amplitude does not correspond one-to-one to the activation of an individual muscle: a single electrode may capture contributions from adjacent muscles [7], and its envelope is related to activation only through the excitation–activation dynamics rather than being that activation [8]. For the hand these limitations are especially severe, since many intrinsic muscles cannot be reached at all by surface electrodes, and invasive fine-wire recordings are impractical in most settings [9]. Alone, sEMG therefore cannot provide a complete picture of hand muscle recruitment.

Musculoskeletal (MSK) modelling combined with inverse dynamics offers a complementary, model-based route from an observed movement to the underlying joint moments [10], from which the activations of the individual modelled muscles can in turn be estimated, including those that cannot be recorded. Because the MSK system is redundant, however, this inverse problem is underdetermined: many combinations of muscle activations reproduce the same movement, and the solution is usually selected by minimising a cost criterion such as overall muscle effort [11]. The chosen solution reproduces the motion but need not coincide with the recruitment that the subject actually used, and may be physiologically implausible; the resulting estimates therefore need to be validated against independent measurements.

This is where the two sources of information become complementary. Kinematics and sEMG are independent measurements, one dense but silent about muscles, the other sparse but directly muscular, and reconciling them could steer the estimated recruitment towards physiological plausibility, particularly for the muscles that are not monitored. The open question, and the one most relevant to the project above, is whether sEMG can be introduced as an additional constraint on the inverse-dynamics solution so that the estimate remains dynamically consistent with the measured movement while becoming more neurophysiologically plausible, without the two requirements coming into conflict [12]. Two further questions accompany it, and its answer cannot be interpreted without them: whether what a single recorded muscle contributes is peculiar to it or already implicit in the other five, and how the model distributes effort over the muscles that no electrode reaches.

1.3 Objectives and research questions

The general objective of this thesis is to build and validate a reproducible pipeline that estimates forearm and hand muscle activations from markerless kinematics, and to establish how the estimation behaves once experimental sEMG is introduced into it as a soft constraint. This objective is articulated into four research questions (RQs):

- **RQ1.** Can forearm and hand muscle activations be estimated from markerless kinematics through inverse dynamics of a MuJoCo-based MSK model?
- **RQ2.** To what extent do the estimated activations agree with experimentally recorded sEMG, in terms of temporal correlation and error?
- **RQ3.** How does information extracted from sEMG, introduced as a soft constraint, affect the inverse-dynamics solution: at what cost to the tracking of the measured kinematics, and how much of the information carried by one recorded channel is redundant with the other five?
- **RQ4.** How does the model distribute mechanical effort across the muscles for which no measurement exists, and how much of the joint torque required by the tasks do the six recorded muscles account for?

1.4 Contributions

The contributions of this work are the following:

- an end-to-end and reproducible pipeline that goes from markerless video (MediaPipe) and sEMG to estimated muscle activations;
- a validation against sEMG of hand muscle-activation estimates obtained by MuJoCo-based inverse dynamics applied to a detailed MSK model of the hand, in a modelling ecosystem where the hand has no shared reference model and where the corresponding validations have been produced with other software [13];
- a feasibility framework for EMG-constrained inverse dynamics, in which information from sEMG is used as a reference for the control without altering the solver formulation;
- a robustness check of that framework, in which each recorded muscle is withheld in turn from the constraint and the solution is compared with the recording of the muscle it was not shown;
- a mechanical characterisation of the model, quantifying how the joint torque the tasks require divides between moving and stabilising, and how much of it the six recorded muscles account for.

These contributions rest on a number of supporting choices (the synergy-based reconstruction of the distal finger joints, a C^1-continuous soft-clipping of the kinematic inputs, and the spectral-matching alignment of repetitions), which are described where they are used. The settings and the per-participant values of the pipeline are collected in Appendix A.

1.5 Thesis outline

The remainder of this thesis is organised as follows. Chapter 2 reviews the relevant background on markerless motion capture, sEMG, MSK modelling, inverse dynamics, and EMG-informed methods, and positions the present work. Chapter 3 presents the materials and methods: the experimental setup and the signal-processing pipeline, the MSK inverse-dynamics model and its subject-specific optimisation, the validation against sEMG, the EMG-constrained feasibility formulation, and the mechanical characterisation of the model. Chapter 4 reports the results for the baseline estimation, the sEMG validation, the feasibility analysis with its robustness check, and the mechanical characterisation (addressing RQ1–RQ4). Chapter 5 discusses the findings across the study, and Chapter 6 draws conclusions and outlines future work. Two appendices collect the supporting material: Appendix A gathers the software, the settings and the per-participant values behind the processing steps, and Appendix B the data behind the results, from the repetition counts to the controls of every actuator of the model.

Chapter 2

State of the art

2.1 Motion capture and kinematics estimation

The starting point of any model-based estimate of muscle activity is the movement itself, described as the time course of the joint angles. In the pipeline proposed in this thesis the kinematics are the main experimental input that drives the MSK model, so the way in which they are acquired, and their accuracy, directly bound the quality of the downstream estimates. Several technologies are available to measure hand and wrist kinematics, each trading accuracy against intrusiveness, cost, and practicality (Table 2.1).

Marker-based optical motion capture is generally regarded as the gold standard for movement measurement [14]. Reflective markers tracked by an array of infrared cameras yield accurate three-dimensional trajectories, but the approach is confined to instrumented laboratories, remains comparatively expensive, and is poorly suited to the hand: the segments are small and closely spaced, markers are difficult to place without impeding the movement, soft-tissue artefact is significant, and frequent self-occlusion of the fingers degrades tracking [15].

Inertial measurement units (IMUs) offer a wearable, line-of-sight-free alternative that can be used in ambulatory conditions [10]. They estimate segment orientation from accelerometer, gyroscope and magnetometer data, but are affected by integration drift and by magnetic disturbances, which require sensor-fusion algorithms and careful calibration; moreover, mounting the many small sensors required to resolve the individual segments of the hand is impractical.

Instrumented (data) gloves capture hand posture by measuring the finger joint angles directly, through sensors embedded in a worn glove. A range of sensing technologies is used, including resistive bend sensors and strain gauges, optical-fibre sensors, and Hall-effect (electromagnetic) sensors [16]. Several aspects are relevant for the present work: many gloves are designed around the fingers and do not, on their own, capture wrist and forearm motion (although some systems add dedicated sensors for this purpose); the abduction angles are influenced by the flexion angles [17], which limits their accuracy; and wearing an instrumented glove adds hardware to the hand that may affect natural movement.

Table 2.1: Main technologies for measuring hand and wrist kinematics, along the trade-off axes discussed in the text.

Technology	Accuracy	Intrusive.	Cost	Hand suitability
Marker-based optical	High	High	High	Poor: occlusion, soft-tissue artefact
IMUs	Moderate	Moderate	Moderate	Poor: too many sensors needed
Data gloves	Moderate	Moderate	Moderate	Good on fingers, limited on wrist
Markerless video	Moderate/low	Low	Low	Moderate, but non-interfering

Markerless motion capture has emerged as a low-barrier alternative that extracts kinematics from ordinary video. Deep-learning pose estimators such as OpenPose [18] and MediaPipe [19] locate anatomical landmarks in the image and, from one or more camera views, reconstruct two- or three-dimensional joint coordinates without markers or dedicated equipment. This considerably lowers the cost and setup burden of acquiring movement data and enables recordings outside specialised laboratories. Accuracy, however, remains a concern: while temporospatial measures are generally comparable to marker-based systems, the estimated joint centres and joint angles are not yet sufficiently accurate for demanding clinical applications, in part because the underlying pose-estimation networks were not designed or trained for biomechanical use [6]. For the hand specifically the problem is even harder, owing to the large number of DOFs, the small size of the segments, and pervasive self-occlusion during grasping [20].

2.2 Electromyography

Electromyography (EMG) records the electrical activity produced by contracting muscles; the present work uses its non-invasive surface form, sEMG. The signal originates at the level of the motor unit, the functional ensemble of a motoneuron and the muscle fibres it innervates. When a motor unit is recruited, an action potential propagates along the sarcolemma of each of its fibres, and the summation of these fibre potentials produces a motor unit action potential; the signal detected at the skin is the spatio-temporal summation of the action potentials of all the active motor units, attenuated and low-pass filtered by the intervening tissue [21]. Surface sEMG is therefore, from the outset, an interference signal that reflects the combined activity of many sources rather than the activation of a single muscle.

Before it can be used, the raw signal is conditioned through a standard chain. A band-pass filter removes motion artefact at low frequencies and noise at high frequencies, and a notch filter suppresses power-line interference; to obtain an amplitude estimate, the so-called linear envelope, the signal is rectified and low-pass filtered [22]; finally, it is commonly normalised to the amplitude recorded during a maximum voluntary contraction (MVC), so that values can be compared across muscles and subjects [23].

Several limitations constrain what sEMG can reveal. Crosstalk, the detection by an electrode of activity generated by neighbouring muscles, blurs the attribution

of the signal to a specific muscle [7]. Only superficial muscles are accessible, while deep muscles cannot be reached from the skin [5]. Amplitude cancellation, the partial cancellation of the positive and negative phases of superimposed action potentials, makes the recorded amplitude a non-linear underestimate of the underlying activity [24]. Both effects are illustrated in Fig. 2.1. The signal is further sensitive to electrode size, spacing and placement and to the normalisation procedure and, more fundamentally, its envelope is related to muscle activation only through the excitation–activation dynamics, so that it is not a one-to-one measure of activation [8].

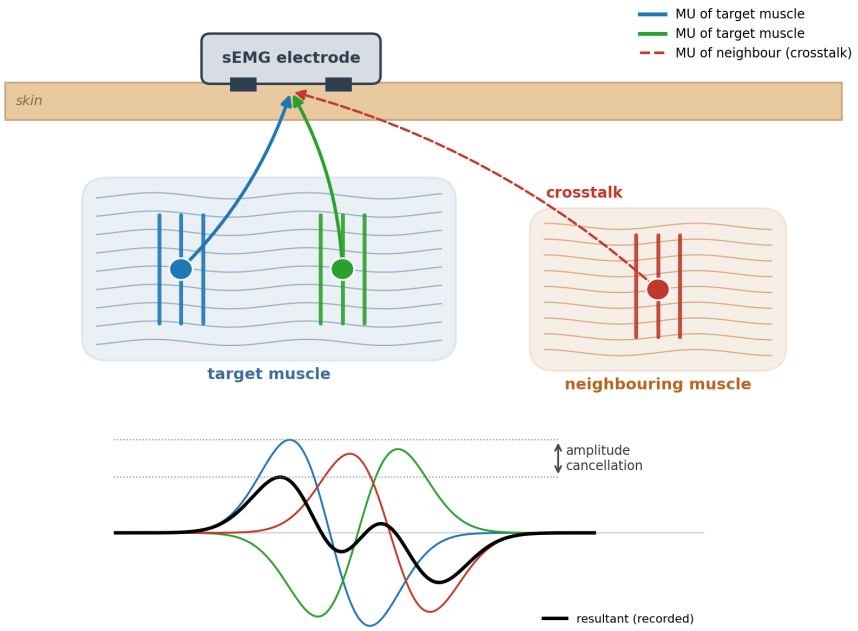


Figure 2.1: Surface sEMG records the summed activity of several motor units: the electrode picks up motor units of the target muscle and, through crosstalk, of neighbouring muscles, and overlapping motor unit action potentials partially cancel (amplitude cancellation).

Two directions are commonly pursued to mitigate these limitations. High-density surface electromyography (HD-sEMG), which uses grids of closely spaced electrodes, improves spatial selectivity, helps to identify and reduce crosstalk, and enables the decomposition of the signal into the activity of individual motor units [25]. Worn as wearable garments, such dense electrode grids in the form of HD-sEMG sleeves provide multi-muscle forearm recordings with up to a hundred and fifty channels [26]. In parallel, muscle-synergy analysis extracts a small number of coordinated activation patterns from multi-muscle recordings, providing a compact description of recruitment and a way to relate a limited set of recorded channels to broader muscle coordination [27].

When the muscles of interest are deep or too small to be resolved at the surface, intramuscular recordings with fine-wire or needle electrodes provide selective access [5]. This selectivity comes at a cost: intramuscular EMG is invasive, requires skilled placement, causes discomfort and is impractical for routine or large-scale use, as

documented for the intrinsic muscles of the hand [9].

Taken together, sEMG provides only a partial, largely superficial and indirect picture of muscle recruitment: it cannot observe all the muscles involved in a movement, and even for those it does record the relationship to activation is imperfect. Despite these shortcomings, it remains the best available compromise for measuring muscle activity in practice: unlike intramuscular recordings it is non-invasive, painless and quick to apply, which makes it the only option that can realistically be used on healthy volunteers and patients outside a clinical procedure.

2.3 Musculoskeletal modelling

Computational biomechanics studies the mechanics of the human body by means of numerical models, replacing quantities that are difficult or impossible to measure in vivo with quantities that can be simulated [28]. Within this field, MSK modelling represents the body as a mechanical system of bones, joints and muscles, and makes it possible to compute internal variables, such as muscle forces, joint loads and muscle activations, from external observations of the movement. Such models are used to simulate surgical or rehabilitative interventions, to assess loading in ergonomics and sport, to design prostheses and exoskeletons [28], and, more recently, as environments for learning motor control [29]. What they share is the ability to provide access to muscle-level quantities that experimental measurement alone cannot deliver, which is also the role they play in this work.

At the core of any MSK model is a description of how muscles generate force. The most widely used is the Hill-type model, which represents the musculotendon unit as a small number of lumped elements: an active contractile element, whose force depends on the muscle's activation and on its instantaneous length and shortening velocity through the force–length and force–velocity relations, placed in parallel and in series with passive elastic elements that capture the stiffness of the surrounding tissue and of the tendon (Fig. 2.2) [8]. The neural input is not mechanical force directly: the excitation delivered to the muscle is transformed into activation through first-order excitation–activation dynamics, and only then into force, so that activation is an internal state of the model instead of a directly measured quantity [8].

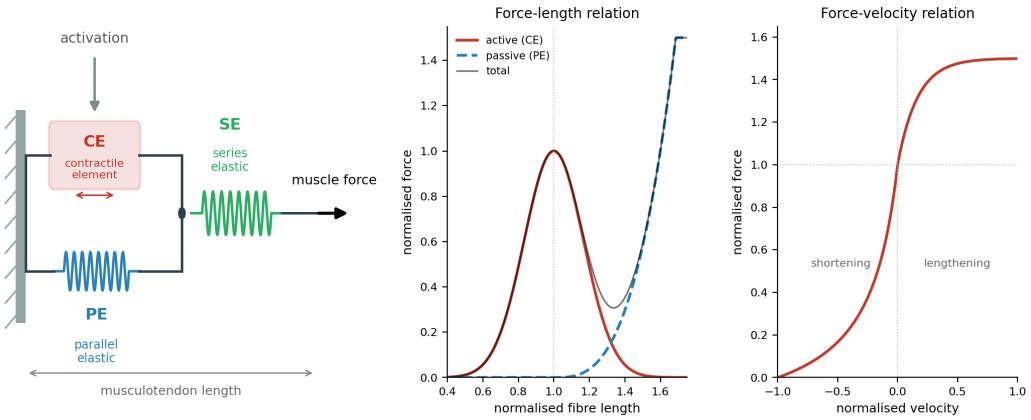


Figure 2.2: Hill-type musculotendon model: an active contractile element (CE) in parallel with a passive elastic element (PE) and in series with the elastic element representing the tendon (SE).

A complete MSK model embeds these actuators in a skeletal structure. The bones are represented as rigid bodies connected by joints that define the DOFs of the system, and each muscle is routed along an anatomical path that fixes its line of action and hence its moment arm about every joint it spans. Given the joint angles, such a model can be used in two complementary directions: in the forward direction, muscle activations produce movement; in the inverse direction, an observed movement is used to infer the underlying muscle forces and activations.

Several software platforms implement MSK models of this kind. OpenSim is the de facto standard in biomechanics, providing Hill-type actuators together with tools for inverse kinematics, inverse dynamics, static optimisation and computed muscle control, and a large library of shared, validated models [13]. AnyBody offers a comparable modelling environment based on inverse-dynamics analysis [30]. More recently, general-purpose physics engines have been adopted for musculoskeletal simulation: MuJoCo provides fast and differentiable rigid-body dynamics with contact [31], and the MyoSuite framework builds on it to offer musculoskeletal models designed for simulation and control, including reinforcement-learning applications [29].

The hand is among the most demanding structures to model. Its many DOFs are actuated by a large number of muscles, divided into extrinsic muscles located in the forearm, whose long tendons cross the wrist to reach the fingers, and intrinsic muscles within the hand itself [4]. The tendons are routed through sheaths and retinacula that make their paths, and therefore their moment arms, configuration-dependent: tendon excursion does not vary linearly with joint rotation, so the moment arm of a given tendon changes along the range of motion (ROM) [32]. Many muscles, moreover, act across several joints at once. The result is a highly redundant system in which the mapping from muscle forces to joint moments is high-dimensional and strongly coupled, which is precisely what makes activation estimation for the hand difficult.

Despite the availability of such models, MSK modelling of the hand remains markedly more fragmented than modelling of the lower limb. Whereas gait and

the large joints of the leg are served by established models, validation studies and inverse-dynamics pipelines, hand models differ widely in the kinematic structure they adopt, in the muscles they include and in the anatomical data they are built on, and no single formulation has become a shared reference [33].

2.4 Inverse dynamics and the muscle redundancy problem

Once a movement has been measured and a MSK model is available, the first step towards muscle-level quantities is inverse dynamics. From the joint angles and their time derivatives, together with the inertial properties of the segments, inverse dynamics computes the net moment that must act about each joint to produce the observed motion [10]. These net joint moments are a purely mechanical quantity: they specify the total effect required at each joint but say nothing, on their own, about how the individual muscles share the load.

Recovering the individual muscle contributions is where the difficulty lies. Because a joint is typically spanned by more muscles than it has DOFs, the system is muscularly redundant: infinitely many combinations of muscle forces produce exactly the same net joint moments (Fig. 2.3) [11]. The inverse problem of distributing the moments among the muscles is therefore underdetermined and cannot be solved from the mechanics alone; an additional criterion is needed to select one solution.

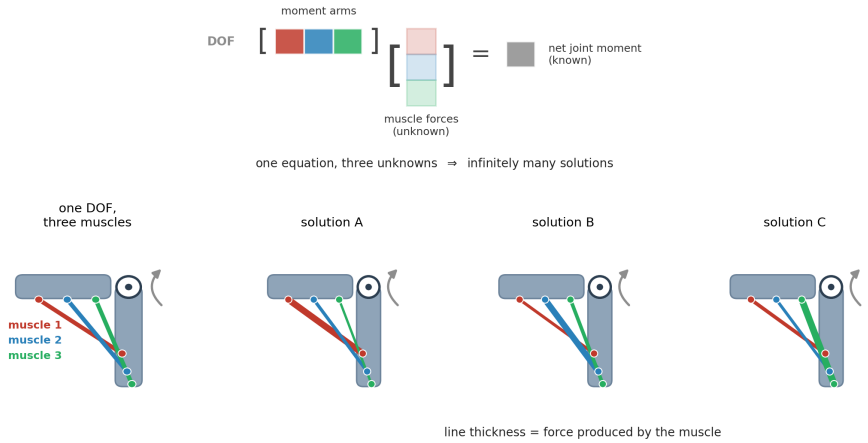


Figure 2.3: Muscle redundancy. A joint with a single DOF is spanned by several muscles, so that the net joint moment required to produce the observed motion can be obtained by many different distributions of force among them.

The most common approach is static optimisation, which resolves the redundancy at each instant by choosing the muscle forces that reproduce the joint moments while minimising a cost function. A widely used criterion is the sum of squared, or higher-power, muscle stresses or activations, motivated by the idea that the central nervous system tends to minimise muscular effort or fatigue [11]. This yields a unique and motivated solution at low computational cost, which is why it underlies most

muscle-force estimation pipelines.

The choice of cost function is not neutral, however. When several criteria are compared on the same movement and the same model, and the resulting activations are validated against recorded EMG, the agreement differs appreciably from one criterion to another, in particular in how the load is shared among synergistic muscles [34]. The most appropriate criterion is not universally agreed upon, and the estimate inherits whatever assumption is built into it.

Alternatives to instantaneous static optimisation exist. Computed muscle control and forward dynamic optimisation account for muscle activation dynamics and for the time history of the movement, tracking the desired kinematics with a dynamically consistent set of activations instead of solving each instant in isolation [35]. Because they embed activation and musculotendon dynamics, these methods respect the fact that activation cannot change instantaneously and that muscle force depends on the history of the movement, at the price of a higher computational cost and a greater sensitivity to modelling detail.

For the hand specifically, model-based estimation of muscle forces and activations has been attempted in only a limited number of studies, often restricted to individual fingers or to reduced sets of muscles, and rarely validated against experimental muscle activity [36].

Common to all these methods is that they select a single solution to an underdetermined problem on the basis of an assumed criterion: the resulting estimate reproduces the measured movement, but does not necessarily coincide with the recruitment actually adopted by the subject. Independent information on muscle activity is therefore needed, either to validate the estimate or to constrain it, which is the subject of the following section.

2.5 Integrating EMG into MSK models

The redundancy of the MSK system can be addressed in two fundamentally different ways. The first, described in the previous section, keeps the movement as the only experimental input and resolves the ambiguity by assuming a criterion that the central nervous system is supposed to optimise. The second replaces that assumption with a measurement: if the neural drive to the muscles is recorded, the distribution of the load no longer has to be inferred, because it is observed. This is the principle of EMG-driven modelling, and the methods that combine the two, using the measured muscle activity without discarding the mechanical requirement of reproducing the movement, occupy the ground between them (Fig. 2.4).

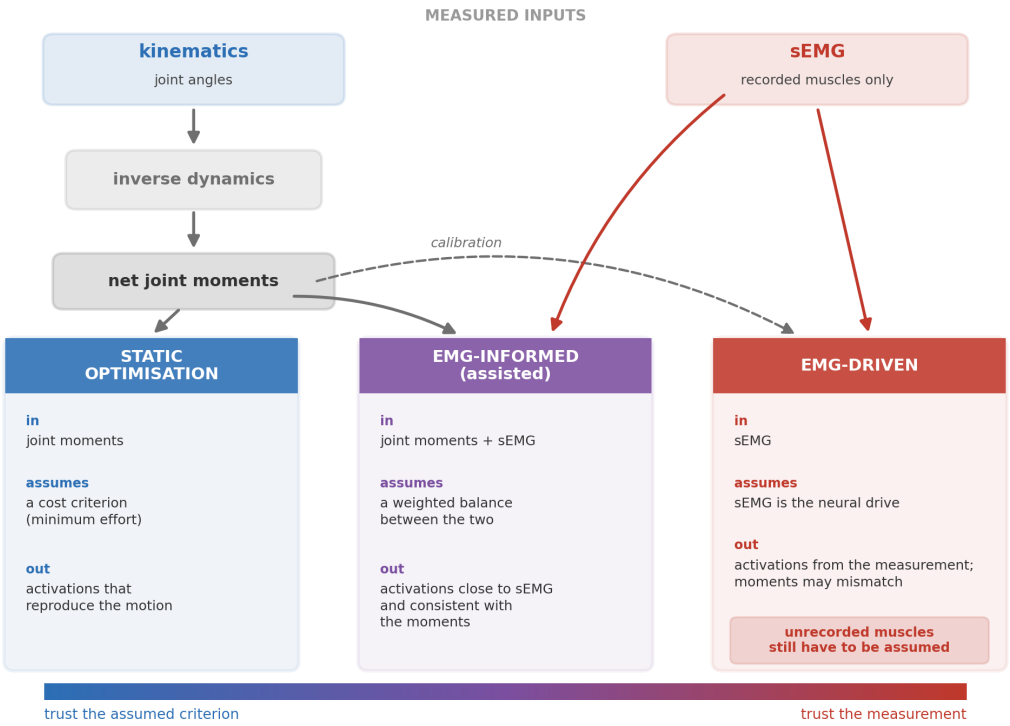


Figure 2.4: Approaches to the muscle redundancy problem, ordered by how much they rely on an assumed criterion instead of the measured muscle activity. For each approach the inputs, the underlying assumption and the resulting estimate are indicated.

In an EMG-driven model the processed sEMG envelopes are taken as the excitations of the musculotendon units, and are converted into muscle activations through the excitation–activation dynamics. In parallel, the measured movement provides the length and contraction velocity of each musculotendon unit. Feeding activations and musculotendon kinematics into a Hill-type model yields the muscle forces, which, multiplied by their moment arms, give the joint moments that the recorded muscle activity would produce [37]. The computation therefore runs forward, from neural input to mechanical output, and a subject-specific calibration step adjusts the parameters that cannot be measured, such as maximum isometric forces, optimal fibre lengths and tendon slack lengths, so that the predicted moments match those obtained from inverse dynamics [37, 38]. The appeal of this formulation is that the recruitment strategy is measured instead of being assumed, so that subject- and task-specific features, in particular co-contraction of antagonists, are retained instead of being penalised away by a cost function [12].

Two limitations, however, restrict how far this can be taken. The first is coverage: the neural drive is measured only for the muscles that carry an electrode, and the activity of all the others still has to be assumed, so that redundancy is removed only in part. This is a marginal issue when a few superficial muscles dominate the joint of interest, but a severe one for the hand, where most muscles are deep, small or intrinsic and cannot be recorded at the surface at all. The second is consistency: because of measurement noise, electrode placement, normalisation and the muscles

that are missing, the moments predicted from the recorded EMG generally do not reproduce exactly those required by inverse dynamics, and the residual mismatch has to be resolved somehow [38].

EMG-informed methods address precisely this mismatch by relaxing the requirement that the measured excitations must be followed exactly. The estimation is posed as an optimisation in which the excitations are allowed to deviate from the recorded envelopes within bounds, and the objective balances two competing terms: staying close to the measured EMG and reproducing the joint moments from inverse dynamics. The CEINMS toolbox implements this family of solutions within a single framework, making it possible to compare purely EMG-driven, EMG-assisted and static-optimisation solutions on the same model and the same movement [12]. By tuning the relative weight of the two terms, these methods span the whole range between full trust in the measurement and full trust in the mechanical criterion, and they have been shown to preserve physiologically plausible recruitment, including co-contraction, better than optimisation alone [12].

A separate line of work addresses the coverage problem rather than the selection problem, and can therefore be combined with either of the approaches above. Synergy extrapolation exploits the low-dimensional structure of muscle coordination: a small number of synergies is extracted from the channels that are recorded, and the excitations of the muscles that are not instrumented are reconstructed from those same synergies, with the synergy weights of the missing muscles estimated as part of the calibration [38]. The approach makes it possible to drive a model with more muscles than there are electrodes, at the price of assuming that the unrecorded muscles share the coordination structure of the recorded ones.

The reference works of this body of literature have been developed on the lower limb, and on gait in particular, within the OpenSim ecosystem, where the muscles that matter most are superficial and accessible and where marker-based motion capture provides accurate kinematics [37, 12]. Applications to the hand are far less common, and they inherit the difficulties already described: many muscles, most of them unrecordable, and kinematics that are harder to measure. Moreover, the emphasis is generally placed on obtaining a single best estimate of the activations, whereas the prior question of whether measured kinematics and measured sEMG can be satisfied simultaneously by the model at all, and at what cost in tracking accuracy, has received little attention.

2.6 Positioning of this work

The preceding sections have followed a single thread. Movement can now be captured with little instrumentation, but only approximately; muscle activity can be measured directly, but only for a few superficial muscles and only as an indirect image of activation; MSK models can supply the muscles that measurement cannot reach, but only by resolving a redundancy that the mechanics alone leaves open; and the methods that bring EMG into the estimation have taken shape on the lower limb,

where none of these difficulties is as severe as it is in the hand. The present work sits at the intersection of these four threads, and this section states the choices it makes and the gap it addresses.

Markerless motion capture was chosen as the source of kinematics. The decision reflects a deliberate trade-off: markerless video requires minimal instrumentation, does not interfere with the forearm sEMG electrodes, and involves little cost. The price of these advantages is a lower kinematic accuracy than marker-based motion capture or instrumented gloves, a limitation that is acknowledged and discussed in the following chapters.

Muscle activity is recorded with conventional per-muscle surface electrodes rather than with an HD-sEMG sleeve, and it is used as such, without reducing it to synergies or feeding it to a learned decoder. This choice follows from the aim of the work: the estimated activation of a modelled muscle is to be compared with the activity recorded from that same muscle, one muscle at a time, so that agreement and disagreement remain interpretable in physiological terms instead of being absorbed into a low-dimensional representation. The sEMG signal remains a partial and indirect measurement and it is treated here as the reference against which the model estimates are assessed, as the only alternative that reaches the deep muscles is intramuscular recording, which is invasive and impractical outside a clinical procedure.

The MSK model adopted in this work is MyoHand, distributed with MyoSuite and built on the MuJoCo physics engine, which provides a detailed representation of the wrist and hand with musculotendon actuators spanning both the extrinsic and the intrinsic muscles [29]. Adopting it places this work outside the OpenSim ecosystem in which most of the literature reviewed above has been produced, and allows the muscle activations to be estimated by a MuJoCo-based inverse-dynamics procedure, described in the next chapters, together with the previous work from which it originates.

Against this background, the gap addressed here can be stated explicitly. MSK modelling of the hand has not converged on a shared reference model in the way that gait and the lower limb have; validations of hand muscle-activation estimates against experimental sEMG are rare; and the integration of EMG into the estimation has been developed on the leg, within the OpenSim ecosystem, and with the aim of producing a single best estimate. The question of whether measured kinematics and measured sEMG can be satisfied *simultaneously* by the same model, and at what cost in tracking accuracy when they cannot, has received little attention. What such a model does with the muscles that no electrode reaches (how it distributes the effort among them, and how much of the torque the movement requires the recorded ones account for) has not been characterised either. It is these questions, on the hand and with a MSK model driven by markerless kinematics, that the present work addresses.

These four concerns map directly onto the research questions stated in Chapter 1. Estimating forearm and hand muscle activations from markerless kinematics through inverse dynamics of the MuJoCo-based model corresponds to RQ1; assessing the agreement between those estimates and recorded sEMG, muscle by muscle,

corresponds to RQ2; establishing how the inverse-dynamics solution behaves once information extracted from sEMG is introduced as a soft constraint, whether the two requirements can be met at once, at what price in kinematic tracking, and how much of what one recorded channel carries is redundant with the others, corresponds to RQ3; and characterising how the model distributes mechanical effort over the muscles it cannot observe corresponds to RQ4.

Chapter 3

Materials and Methods

This chapter describes how the recordings were made and how they were turned into an estimate of muscle activity obtained from movement alone; how that estimate is compared with the muscle activity actually measured; how the same measurement is afterwards introduced as an input to the estimation itself; and how the mechanical means by which the model produces the movement are characterised. Fig. 3.1 summarises the whole pipeline and indicates where each step is described.

The four sections that follow correspond to the four research questions of Sec. 1.3: Sec. 3.1 builds the estimate from movement alone (RQ1), Sec. 3.2 compares it with the recorded sEMG (RQ2), Sec. 3.3 introduces the sEMG envelopes as a soft constraint on the estimate (RQ3), and Sec. 3.4 characterises the mechanical means the model has at its disposal (RQ4).

The two modalities are recorded together and processed in parallel up to the point at which the individual repetitions are located; only then do they separate, the kinematics becoming the input of the musculoskeletal model and the sEMG the quantity against which the model is judged. Nothing measured on the muscles enters the estimation of Sec. 3.1, not even through the parameters, which are fitted on the kinematics alone. The same envelopes reappear on the right of the figure in the opposite role, no longer as a reference but as a constraint on the estimation itself, which is the subject of Sec. 3.3.

The whole pipeline is written in Python and runs on the MSK model distributed with MyoSuite; the libraries used at every step, the two landmark models on which the kinematic extraction relies and the machines on which the computations were run are listed in Appendix A, under *Software and hardware*.

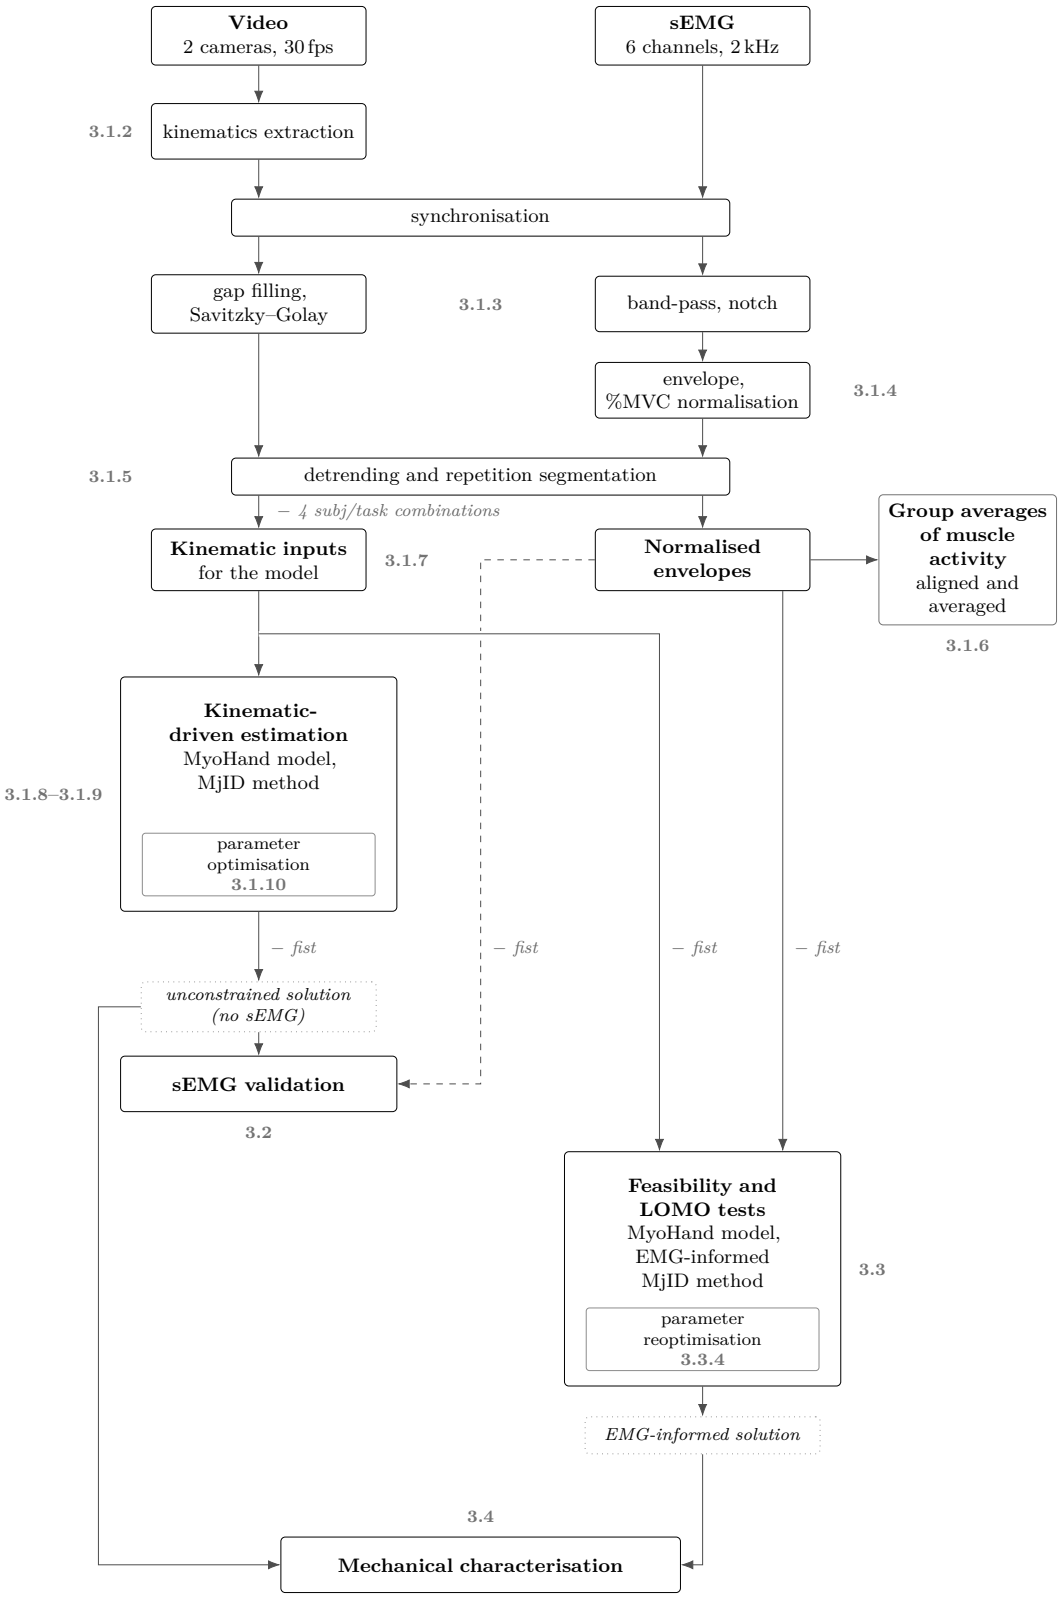


Figure 3.1: Overview of the pipeline described in this chapter; the numbers indicate the section in which each step is described. Solid arrows mark a signal used to produce something, dashed arrows a signal used as a term of comparison.

3.1 From recordings to muscle activation estimates

This section follows the recordings from the laboratory to the estimate of muscle activity that the rest of the chapter uses. Sec. 3.1.1 describes who took part and how a session was organised, and Sec. 3.1.2 the instruments, the placement of the electrodes and of the cameras, and the extraction of joint angles from the video. The five subsections that follow bring those recordings into a usable form: the two streams are merged and filtered (Sec. 3.1.3), the sEMG is reduced to a normalised envelope (Sec. 3.1.4), the filtered recordings are detrended and individual repetitions are located (Sec. 3.1.5) and then aligned in time so that they can be averaged to extract the muscle activity patterns for every analysed movement (Sec. 3.1.6), and the measured angles are turned into the input the model requires at every instant (Sec. 3.1.7).

The last three subsections concern the muscle activation estimation itself. Sec. 3.1.8 presents the MSK model, Sec. 3.1.9 the inverse-dynamics method that reconstructs muscle activity from an imposed movement, and Sec. 3.1.10 the way in which the parameters of the method are fitted to each participant.

3.1.1 Participants and experimental protocol

Seven healthy adults, four men and three women, took part in the study (age 23.1±0.9 years, range 22–24). They were volunteers recruited at the Translational Neural Engineering Laboratory (EPFL, Campus Biotech, Geneva), where all recordings took place, with no known neuromuscular disorder or recent injury of the upper limb; the study is concerned with muscle activity in healthy individuals, so no clinical population was involved. All participants had given written informed consent for sEMG and video recordings within the laboratory research protocol approved by the Regional Research Ethics Committee of Geneva (CCER, ref. 2024-00574). Six participants were right-handed and one was left-handed; recordings were made on the dominant forearm and hand, and the kinematics of the left-handed participant were mirrored onto the right-hand structure at the extraction stage (Sec. 3.1.2), since the model used throughout is a right-hand model (Sec. 3.1.8). No anthropometric measurements were collected, so the model is used with its default segment dimensions and is not scaled to the individual participant (Sec. 3.1.8). The first participant was the experimenter, and that session also served to finalise the protocol: it is the only one containing three baseline recordings instead of one, a difference that is accounted for in the processing and has no effect on the analysis.

Each participant completed a single session, organised in three consecutive blocks (Fig. 3.2). The first block was a sixty-second baseline recording, with the participant seated and the hand at rest; it provides the reference for the resting condition of both the sEMG signals and the kinematics. The second block was the MVC calibration: one maximal isometric contraction of ten seconds was performed for each of the six recorded muscles, against manual resistance, with a single attempt per muscle and

a rest of roughly thirty to sixty seconds between contractions. These contractions were intended to provide the reference amplitudes against which the sEMG signals are later normalised, following the standardised isometric MVC normalisation of the Consensus for Experimental Design in Electromyography (CEDE) matrix [23]. That reference prescribes taking the maximum of the values obtained and recommends between three and five contractions producing consistent values, separated by rest periods of one to two minutes. A single attempt per muscle, with a shorter rest, is therefore fewer than is recommended, and the reference amplitude may fall short of the true maximum. The third block contained the seven motor tasks.

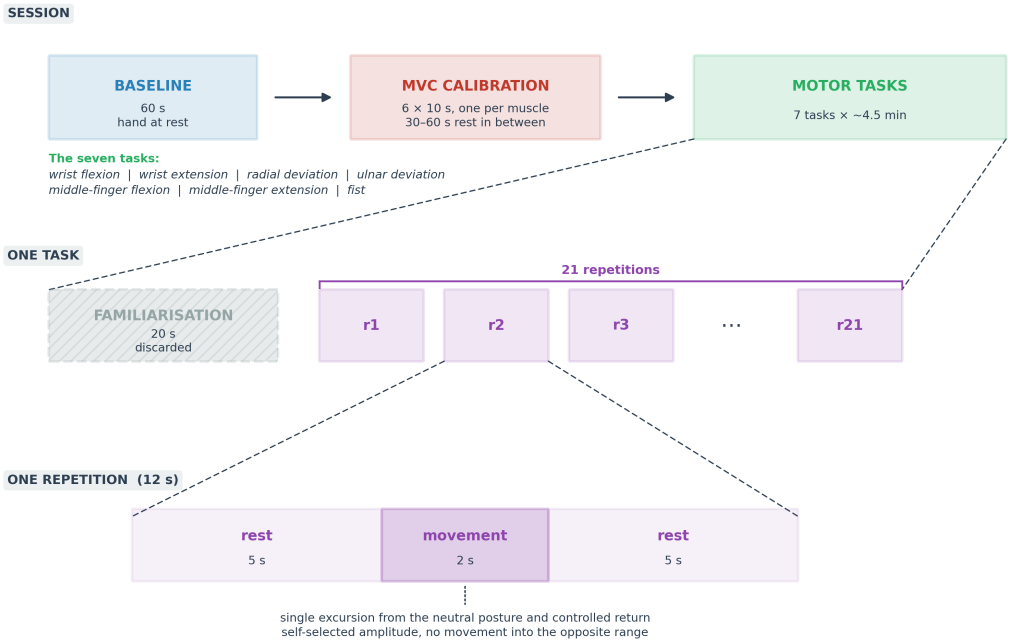


Figure 3.2: Temporal structure of the acquisition session, at three nested levels: the session, one motor task, and one repetition. Durations are given in the labels; the blocks are not drawn to scale.

Throughout the session the participant sat comfortably at a table, with the forearm resting on two supports and the hand protruding beyond the edge of the table, free to move under gravity. The forearm was in mid-position, neither pronated nor supinated. The non-dominant arm was kept still and away from the working area, either resting on the thigh or hanging at the side of the body. In this configuration the hand at rest does not lie exactly in the anatomical neutral posture: under its own weight the wrist settles into a slight flexion and a slight ulnar deviation, and the fingers rest slightly flexed. This resting posture is the reference configuration of the whole study: it is the posture held during the baseline recording and, nominally, during the rest intervals between repetitions, and it is the origin from which every joint excursion is measured. It is also what motivates the constant wrist offset applied when the measured kinematics are mapped onto the model (Sec. 3.1.7). The posture held during the inter-repetition rest is nominally the same as during the baseline, but the recorded signals drift slowly over a trial; this drift is removed at the detrending

stage (Sec. 3.1.5).

The tasks span the main DOFs of the wrist and, separately, the flexion and extension of a single finger and of all the fingers: wrist flexion, wrist extension, wrist radial deviation, wrist ulnar deviation, middle-finger flexion, middle-finger extension, and fist closure. They were performed in this order by every participant. Fist closure differs from the other six tasks in engaging all the fingers at once. The distinction is not between a single and multiple DOFs (the middle-finger tasks are themselves multi-articular, involving the metacarpophalangeal (MCP), proximal interphalangeal (PIP), and distal interphalangeal (DIP) joints), but between one digit and the whole hand. As portrayed in Fig. 3.2, each task was recorded as a single continuous trial with a fixed temporal structure: an initial twenty-second period of free movement, during which the participant became familiar with the gesture and the markerless tracking settled, followed by twenty-one repetitions of the movement cycle. Each repetition consisted of five seconds at rest, a two-second movement, and five further seconds at rest, for a total duration of about four and a half minutes per task. The intervals were announced by an interval-timer application, so that the participant did not have to keep count and the same cadence was imposed on every trial and every participant. About one minute of rest separated one task from the next. The amplitude of the movement was not imposed and was chosen by the participant. The initial twenty seconds are excluded from all subsequent analyses.

Every repetition was a single excursion away from the resting posture followed by a controlled return to it, without continuing into the opposite range: in the wrist flexion task, for instance, the wrist flexed and then came back to rest, but did not go on to extend. The same holds for the two deviation tasks, for the two finger tasks, and for fist closure, in which the fingers closed from their resting semi-flexed position and reopened only as far as that position, without extending further. Flexion and extension are therefore two distinct and unidirectional tasks rather than the two halves of a single oscillation. This defines the movement cycle on which segmentation and time normalisation operate (Sec. 3.1.5), and it also means that the return phase is a controlled return to the resting posture rather than a movement actively driven in the opposite direction.

The electrode–skin interface, and with it the amplitude of the recorded signal, changes over time and, above all, whenever electrodes are removed and applied again [22]. Recording all the repetitions of a task within a single continuous trial, instead of recording them as separate takes, removes that source of variability between repetitions, and it is what allows the segmentation to operate on a homogeneous signal. The twenty-one repetitions are those performed; the number retained after segmentation and quality pruning (Sec. 3.1.5) is smaller, varies across trials and subjects, and is reported with the corresponding results.

3.1.2 Data acquisition

Muscle activity was recorded with a Cometa WaveX wireless sEMG system (Cometa srl, Bareggio, Italy), using six probes of two form factors (two MiniX and four PicoX) which share the same electrical specifications: a sampling rate of 2 kHz per channel, 16-bit resolution, an input impedance above 100 MΩ, a common-mode rejection ratio above 100 dB, and a baseline noise below 3 μV. Each probe applies a fixed analogue band-pass between 10 and 500 Hz, so the signal is already band-limited at acquisition; the digital filtering described in Sec. 3.1.3 operates within that band. All six channels are sampled synchronously by the same receiver. The two form factors differ only in how the electrodes are attached: the PicoX probes carry their two contacts at a fixed spacing set by the sensor body, whereas the MiniX probes connect through flexible leads, so their interelectrode distance is chosen at the time of placement. The interelectrode distance was therefore not identical across the six channels; it was approximately 25 mm, slightly above the 20 mm recommended by SENIAM [39]. Disposable pre-gelled Ag/AgCl electrodes of 24 mm diameter (Kendall H124SG) were used throughout.

The six recorded muscles were three wrist and finger flexors and extensors: flexor digitorum superficialis (FDS), flexor carpi radialis (FCR), flexor carpi ulnaris (FCU), extensor digitorum communis (EDC), extensor carpi radialis longus (ECRL), and extensor carpi ulnaris (ECU). What restricts the choice is accessibility: these six muscles are superficial and can be reached by surface electrodes, whereas the deeper flexors and the intrinsic hand muscles cannot. Being represented in the MSK model (Sec. 3.1.8) is not what selects them, since the model contains deep and intrinsic muscles as well; it is the condition that makes the rest of the work possible, because every recorded channel can then be put beside the activation estimated for the same muscle. For the two multi-tendon muscles, FDS and EDC, the channel was matched to the modelled compartment serving the third digit, consistently with the finger tasks of the protocol. Without this one-to-one correspondence the muscle-by-muscle comparison of Sec. 3.2 could not be made. The forearm muscles are, however, closely packed and partly overlapping, so the recorded signals are affected by crosstalk from neighbouring muscles [7]. This is an intrinsic limitation of surface recording on the forearm and its consequences for the comparison are taken up in the discussion.

Before electrode application, the skin over each site was cleaned with an alcohol-based disinfectant solution; no abrasion was performed and electrode–skin impedance was not measured, which is a departure from the more thorough preparation recommended for sEMG detection [22]. Each site was then found by palpating the muscle belly while the participant repeated a contraction that isolates the target muscle as far as it can be isolated voluntarily; the manoeuvre used for each muscle is given in Tab. 3.1. The probes were aligned with the estimated fibre direction. The general procedure follows the SENIAM recommendations [39], but those recommendations do not cover the forearm muscles recorded here, so the individual sites were established by palpation. The resulting layout is shown in Fig. 3.3. The same experimenter

placed the electrodes in every session, so the procedure was applied consistently; the exact position of each site nonetheless varied slightly from one participant to the next, since it was found on the participant.

Table 3.1: Voluntary contractions used to locate each muscle by palpation before applying the electrodes.

Muscle	Test contraction
FDS	flexion of the middle finger at the MCP joint
FCR	wrist flexion towards the thumb (radial)
FCU	wrist flexion towards the little finger (ulnar)
EDC	extension of the middle finger at the MCP joint
ECRL	wrist extension towards the thumb (radial)
ECU	wrist extension towards the little finger (ulnar)

Each contraction was repeated several times while the site was palpated. In every case the participant was asked to keep the joints not involved in the manoeuvre as relaxed as possible and not to move them actively.

Movement was recorded with two off-the-shelf USB webcams (Logitech C525 and Logitech StreamCam), both acquiring at 1280×720 pixels and 30 frames per second, under the ambient lighting of the laboratory, which was not controlled. No scientific camera system was used, and the two cameras were neither genlocked to each other nor hardware-triggered by the sEMG system: keeping the instrumentation minimal and inexpensive is part of the rationale for markerless capture stated in Sec. 2.6,

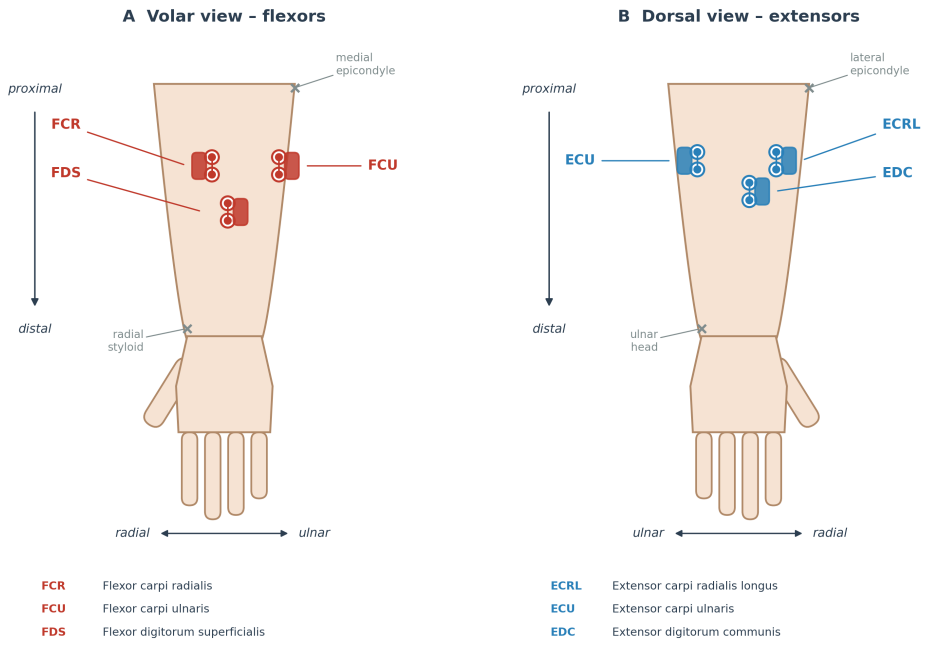


Figure 3.3: Placement of the six sEMG probes on the forearm: (A) volar aspect, flexors; (B) dorsal aspect, extensors. The layout is drawn on a right forearm; the same sites were used on the left forearm of the left-handed participant.

and the accuracy that is given up in exchange is discussed there. The first camera was fixed to a laboratory cabinet at a height of roughly 2 m above the floor and approximately 50 cm horizontally from the hand, and looked down on it from above; the second stood on a tripod at approximately 70–80 cm above the floor and about 1.5 m from the hand, and viewed the volar aspect of the hand from the medial side. The arrangement is shown in Fig. 3.4.

Two views were needed because a single one cannot resolve all the recorded DOFs: a movement that develops out of the image plane of a camera is poorly resolved by that camera. The two views could not be assigned freely, however. The wrist angles are defined between the forearm and the palm and therefore require the body landmarks, which were estimated from the overhead view only, as described below; both wrist flexion–extension and radial–ulnar deviation were consequently taken from that view, even though the deviation develops out of its image plane. That this was nevertheless the better of the two options was confirmed in practice, the deviation being tracked more reliably from above than from the medial view, but it remains a limitation and is returned to in the discussion. The finger joint angles were taken from the medial view, which sees the digits from the side. For the finger tasks the tripod camera was moved closer to the hand, so as to gain spatial resolution on the digits. Both cameras recorded continuously in every trial, so each task has a recording from both views even though only one of the two contributes the joint angles used downstream; which view is used for which task, and how the two are combined, is described in Sec. 3.1.3. Both cameras (except for the tripod camera on the finger tasks) framed the whole upper body rather than the hand alone: this is required by the pose estimation described next.

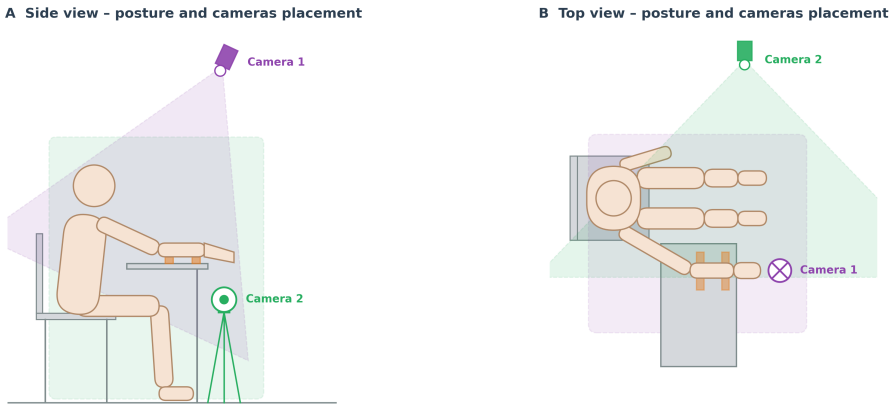


Figure 3.4: Acquisition setup, seen from the side (A) and from above (B): seated posture, forearm resting on two foam supports with the hand protruding beyond the table edge, and the placement of the two cameras with their fields of view. The drawing shows a right-handed participant; the arrangement was mirrored for the left-handed one. The figure is not drawn to scale.

Kinematics were extracted from the video offline, without markers, using the MediaPipe framework [19]. The Hand Landmarker model was run on both views to obtain the hand landmarks, and the heavy Pose Landmarker model was run on the

overhead view to obtain the body landmarks. Both models were used in their float16 form and in video mode, with a detection, presence and tracking confidence threshold of 0.5; the two model files are kept with the code, so repeating the extraction does not depend on what the distribution endpoint currently serves, but it does depend on the version of the MediaPipe package used to run it. The wrist angles cannot be obtained from the hand landmarks alone, since they are defined between the forearm and the palm: the forearm axis was therefore taken from the pose landmarks and the palm direction from the hand landmarks, the two sets being expressed in a common global orientation. This is why the body, and not only the hand, has to be in frame, and why the wrist angles are available from the overhead view alone. The finger joint angles were computed from the hand landmarks alone. The sign conventions of the resulting joint angles are stated in Sec. 3.1.7, together with the transformations that bring them onto the model. The extraction code derives from an in-house implementation developed within the laboratory.

Two aspects of the kinematics extraction required adjustments. The first concerns handedness: for the left-handed participant the mediolateral component of the wrist angles was negated, so that the radial and ulnar directions keep the same meaning on both arms, and the landmarks were relabelled onto the right-hand structure at this stage, so that all participants enter the rest of the pipeline in the same form. The second concerns the quality of the video. In the medial view the hand was occasionally poorly lit, and for two participants a gamma correction was applied to some of their trials before landmark detection; the overhead view never required it. The portion of each video to be processed was delimited by hand from the session annotations, with the exception of the ulnar deviation task, for which, for some participants, the annotated trimming was not applied and the whole recording was processed. Where the tracking remained too unreliable to yield a usable joint-angle trajectory, the corresponding participant–task combination was excluded from the analysis altogether; how that judgement was made is described in Sec. 3.1.5.

Both the sEMG samples and the video frames were acquired by a single in-house recording application, which drives all the devices and stamps every sample and every frame with a time read from the same host clock. The alignment between the two modalities is therefore obtained in software, at recording time, rather than by an electrical synchronisation of the devices; no external trigger, visual cue or post-hoc realignment was used. Because the webcams are free-running, their frames are not equally spaced in time: the effective rate was 29.8–29.9 frames per second rather than exactly 30, and the timestamp of each individual frame, not a nominal frame rate, is what carries the kinematics onto the common time axis. Every recording is accompanied by an automatic integrity report produced by the same application, which compares the nominal and the effective sampling rates and the instants at which each stream starts and stops. Across the recordings the effective sEMG rate was 2000.0 Hz, and the streams started and stopped within less than one second of each other.

3.1.3 Signal preprocessing

The two streams were first brought together. For each trial the sEMG samples and the extracted joint angles were written into a single file, each retaining its own timestamps: nothing was resampled at this stage, and the two signals remain on the common time axis established at recording time (Sec. 3.1.2). Which view feeds the merged file depends on the task. For the four wrist tasks the overhead view is used, and only the two wrist angles are retained; for the finger tasks and for the baseline the medial view is used, and the fifteen finger joint angles are retained. Each trial therefore carries the angles of the joints that the task actually involves and not the others: there is no wrist trajectory in the finger trials, and no finger trajectory in the wrist trials. The MVC recordings contain no kinematics and are carried through this stage as sEMG alone.

The sEMG channels were then band-limited and cleared of power-line interference, as is standard practice for surface recordings [22]. A fourth-order Butterworth band-pass between 20 and 450 Hz was applied first, followed by a notch at the power-line frequency of 50 Hz and at its harmonics up to 450 Hz, the upper edge of the band-pass. Both filters were applied forward and backward, so that no phase distortion is introduced: this matters here because the quantity of interest is the timing of muscle activity relative to the movement, and a filter that delayed one channel with respect to another would corrupt exactly what Sec. 3.2 sets out to measure. The digital cut-off frequencies lie inside the analogue band of the probes, so the two filters act in series. Baseline, MVC and task recordings all received the same treatment. No rectification was applied at this stage; the envelope is extracted in Sec. 3.1.4.

The kinematic traces carry the high-frequency jitter typical of frame-by-frame landmark estimation, and in addition a defect that the sEMG does not have: in some frames no estimate exists at all, because the detection returned nothing. These gaps have to be closed before any filter can be applied, and were filled by linear interpolation between the neighbouring valid samples; a channel that was empty for a whole trial was set to zero. The traces were then smoothed with a Savitzky–Golay filter of second order over a window of 21 samples, which at about 30 frames per second corresponds to roughly 0.7 s. The filter fits a low-order polynomial to the samples inside a sliding window by least squares and retains its value at the centre of the window; compared with a moving average, or with a low-pass filter of comparable bandwidth, it removes the jitter while preserving the height and the shape of the peaks, which here carry the amplitude of the movement. Being symmetric, it introduces no phase shift either.

From this point on the joint angles are treated as a signal uniformly sampled at 30 Hz. This is an approximation, since the effective frame rate is 29.8–29.9 frames per second, and it is used only where a duration has to be converted into a number of samples, an operation in which an error of half a percent amounts to a few tens of milliseconds. Whenever an operation involves both modalities at once (evaluating the sEMG at the instants of the video frames, for instance, or cutting one signal

over an interval defined by the other), the timestamps are used and the nominal rate plays no part.

3.1.4 EMG envelopes and MVC normalisation

The interference signal produced by the filtering of Sec. 3.1.3 oscillates around zero and cannot be compared directly with a muscle activation, which is a non-negative quantity varying slowly over the movement. It was therefore reduced to its envelope: each channel was full-wave rectified and then low-pass filtered with a fourth-order Butterworth filter at 2 Hz, applied forward and backward so that no delay is introduced between the envelope and the movement that produced it. Baseline, MVC and task recordings were all rectified and smoothed in the same way.

There is no single conventional cut-off for a linear envelope, and the value follows from what the envelope is meant to represent. Cut-offs of tens of hertz preserve the timing of individual bursts, whereas work in which the envelope has to stand for the slow course of muscle activation uses a few hertz: the EMG-driven musculoskeletal model that established the approach obtains its envelope with a low-pass cut-off of 6 Hz [37], and values of 4 to 6 Hz are common in that literature. The value adopted here, 2 Hz, is lower still.

What motivates it is what the envelope has to show: not the individual bursts of the interference signal, but the slow rise and fall of activity over a movement that lasts two seconds. The waveforms recorded across participants and repetitions were markedly variable in their fine structure, and a higher cut-off preserved that variability without making the underlying course of the activity any easier to read.

A cut-off this low has a price: part of what it suppresses is variation the model would have reproduced, and where that happens the suppression appears as a disagreement between the two. That price was accepted here, because the movements last about two seconds and the activity of interest varies far more slowly than the cut-off.

Each envelope was then divided by a reference amplitude, one per muscle and per participant, so that the result is dimensionless and can be placed beside a modelled activation, which is bounded between zero and one. Two candidate references were computed. The first comes from the MVC calibration block: within the annotated contraction interval the envelope reaches a peak, and the reference is the mean of the envelope over a window of 0.5 s centred on that peak, rather than the instantaneous maximum, so that a single sample cannot set the scale. The second is the largest envelope value reached by that muscle in any of the motor tasks of the same session. The reference finally used is the larger of the two.

Taking the larger of the two guarantees that the normalised envelope never exceeds one, which would otherwise happen whenever a muscle is more active during a task than during its calibration contraction, and would make the comparison with a bounded activation meaningless. It happened often: in 25 of the 42 participant–muscle combinations the reference comes from the tasks and not from the calibration, and in

the most extreme cases the task maximum is close to three times the calibration value. This is consistent with the single MVC attempt per muscle discussed in Sec. 3.1.1.

The two candidates are not the same statistic (the calibration value is a mean over half a second, the task value an instantaneous maximum), so the selection is in principle biased in favour of the second. That asymmetry is not, however, what drives the outcome. Recomputing both candidates as instantaneous maxima changes which of the two prevails in one channel out of 42, and recomputing both as means over 0.5 s changes it in six; in either case the tasks still prevail in the large majority of channels. The muscles therefore exceed their calibration contractions for a substantive reason and not because of the way the two quantities were summarised. The per-channel values of both candidates, and of the reference that would have been selected under each of the two uniform choices, are reported in Table A.1.

The task maximum was searched over every task recorded in the session, including the fist closure trial, which is afterwards withdrawn from the model-based analysis (Sec. 4.2.1).

The consequence changes what the normalised amplitude means. For the channels in which the calibration prevailed, the unit is a maximum voluntary contraction, and the procedure is the standardised isometric MVC normalisation of the CEDE matrix [23]. For the majority in which the task maximum prevailed, the unit is instead the peak activity observed in that session. This is closest to the approach the same matrix describes as normalisation to the peak amplitude in the task, with one difference: the peak is taken across all the tasks recorded in the session and not within the task being analysed. The normalised envelope is therefore a fraction of the largest activity recorded for that muscle, and only in part a fraction of its maximum voluntary contraction. Since the reference is a single constant per channel, this affects the amplitude of the normalised signal but not its shape or its timing; which of the comparisons of Sec. 3.2 are sensitive to it is taken up there. For brevity, normalised amplitudes are expressed throughout this work as a percentage and denoted %MVC; the notation is a convention, and the denominator it refers to is always the reference amplitude defined in this section, that is, the calibration value or the session peak, whichever is larger.

The normalised envelopes are used in two different roles in what follows, and the distinction matters because the two are separated by the excitation-to-activation dynamics of Sec. 3.1.8. In Sec. 3.2 they serve as a term of comparison for the muscle activations the model produces. In Sec. 3.3 they are used instead as a reference for the muscle *controls*, which are the input to those dynamics rather than their output; what that substitution assumes is discussed there.

The same normalisation was applied to every task trial and to the baseline recordings. The latter were also restricted to the interval annotated during the session as the one in which the participant was effectively still, so that the settling at the beginning of the recording and its tail are excluded; these annotated intervals last about a minute. From this point on the sEMG of each trial (baseline, calibration contraction or motor task), for every participant, is represented by six non-negative

envelopes, one per recorded muscle, expressed as a fraction of their reference and sampled at the original 2 kHz.

3.1.5 Detrending and repetition segmentation

Throughout this chapter, a *trial* is the continuous recording of one task by one participant, and a *repetition* is one execution of the movement within it. Up to this point each task is a single trial of about four and a half minutes, containing twenty-one repetitions separated by rest intervals. Two operations are needed before the repetitions can be treated as comparable observations: the slow drift that affects both signals over such a trial has to be removed, and the boundaries of each repetition have to be located.

The two operations depend on each other. A drift estimated over a window that contains a movement is biased by that movement, and a boundary located on a drifting signal moves with the drift. The procedure therefore works in two passes: the drift is first estimated in a way that does not require knowing where the repetitions are, the boundaries are located on the signal so detrended, and the drift is then estimated a second time, now using those boundaries to keep the movements out of the estimate. The two passes use different methods, for the reasons given below.

Both signals are detrended, not only the kinematics. A drift is not necessarily an artefact: a slow change of posture over four minutes may be real, and if it is real it is sustained by muscle activity, which appears as a slow change in the envelopes as well. Removing it from one signal and leaving it in the other would break the correspondence between them, since the same slow variation would be erased on one side and preserved on the other. Removing it from both keeps the two describing the same thing, at the price of discarding whatever genuinely slow component the movement may have had, a component that the repetition-based analysis of Sec. 3.1.6 could not have used in any case.

Trial boundaries

The first twenty seconds of familiarisation are discarded, as anticipated in Sec. 3.1.1. The start of the analysed portion is then placed five seconds before the first repetition, whose time was annotated during the session, so that every trial begins with a stretch of rest of the same length, and the last three seconds are discarded as well. Both values were adjusted for individual trials in which the annotation did not correspond to the recording. For the first participant, whose session precedes the final annotation format, the boundaries are not derived from the annotations at all: the first 13.0 s of the recording and its last 5.0 s are removed.

The trimmed kinematic interval then defines the interval of the sEMG. The envelopes are cut between the timestamp of the first retained kinematic frame and that of the last, rather than by counting a number of samples at the nominal sampling rate. This matters over a recording of four and a half minutes: had the window been obtained by counting samples, every small discrepancy between the nominal

and the effective sampling rate would have accumulated along the trial, and the two modalities would have ended it displaced with respect to each other. Anchoring the window to measured timestamps keeps them aligned from beginning to end.

Locating the repetitions

One kinematic channel is chosen as the reference for each task (the wrist angle in the plane of the movement for the four wrist tasks, the MCP angle of the middle finger for the two finger tasks and for fist closure) and low-pass filtered at 1 Hz, again forward and backward.

The drift of this reference is estimated with an asymmetric least-squares baseline [40]. The estimate must follow the resting level of the signal and ignore the excursions, and the asymmetry is what decides on which side of the signal it is kept: when the movement raises the reference angle, the fit is constrained to run along the lower edge of the signal, where the rest lies, and when the movement lowers it, along the upper edge. Subtracting this estimate leaves a signal whose resting level is constant and whose excursions are preserved.

A repetition is then detected as a continuous interval in which the detrended reference exceeds a threshold on the side of the movement. The threshold takes one of two forms, depending on the trial: either a fixed angle in degrees, or a multiple k of the standard deviation of the detrended signal, measured from its resting level. Three rules refine the raw intervals so obtained: gaps shorter than one second are closed, so that a momentary crossing back through the threshold in the middle of a movement does not split one repetition into two; intervals shorter than 300 ms are discarded as spurious; and once an interval is accepted, no other is accepted until eight seconds later, which is shorter than the twelve-second cycle imposed by the protocol and longer than any single movement. Each accepted interval is finally extended by a margin of the order of a second on each side, so that the segment contains the beginning of the rise and the end of the return, and the instant of maximum excursion is recorded as the peak of the repetition.

Per-participant and per-task tuning

The parameters of this procedure (the form and value of the threshold, the two margins, and the trial boundaries described above) are not the same for every trial. They were set for each combination of participant and task, by inspecting the reference signal together with the detected boundaries, and they had to be: the movement amplitude was chosen by the participant and not imposed (Sec. 3.1.1), so the excursion that separates movement from rest differs by an order of magnitude across participants, and a single threshold rule cannot serve all of them. In one trial, a middle-finger extension, the reference channel had to be reflected about its resting level, the movement having been recorded with the opposite sign to the one expected for that task. In addition, individual repetitions were listed for exclusion where the detected boundaries were visibly wrong, most often at the beginning of

a trial: thirteen repetitions across the whole dataset. The complete set of values is reported in Table A.2 of Appendix A.

This tuning is a manual, supervised step, and it is the least automated part of the pipeline. It operates on the boundaries of the repetitions and not on their content: no criterion applied here can change the shape of a signal within a repetition, only which portion of the trial is called a repetition and which repetitions are kept.

Four combinations segmented on the sEMG envelope

In four combinations of participant and task the reference kinematic channel did not permit the repetitions to be located at all. The excursion of the recorded angle during the movement was not distinguishable from the fluctuation of the same angle at rest, so that the individual repetitions did not emerge from the continuous signal and no choice of threshold separated one condition from the other; Fig. B.1 shows one of these channels beside a channel from the same task in which the repetitions are evident, on the same scale. For these four, the segmentation was performed on the sEMG envelope of the muscle that is the agonist of the task (ECU for the three cases of ulnar deviation, EDC for the one case of middle-finger extension) which retains the timing of the repetitions even when the recorded angle does not. The rule that locates the repetitions is the one described above, applied to the envelope instead of the angle: a threshold set at a multiple of the standard deviation above the initial rest in two of the four combinations, and a fixed level, 20 and 30 %MVC, in the other two, the values being listed with the rest in Table A.2.

These four combinations are precisely the four that are excluded from the model-based analysis. The exclusion follows from the same observation that forced the change of segmentation reference: if the kinematics of a trial are not usable to say when the movement happened, they are not usable as the input from which joint angles are imposed on the model (Sec. 3.1.7). Their sEMG is unaffected by this and the trials still contribute to the group averages of muscle activity (Sec. 3.1.6), which is why they were segmented at all rather than discarded outright. Every combination for which the kinematics allowed segmentation is retained; no other criterion was applied.

Detrending of the exported signals

The asymmetric baseline described above serves only to locate the repetitions, and it could not serve for more: weighting the samples above it far less than those below makes it an estimate of the lower edge of the fluctuation rather than of its centre, which is exactly what keeps it from being drawn upwards by the movements, and exactly what disqualifies it as a quantity to be subtracted from the amplitudes that are carried forward. Once the boundaries are known, the signals that are carried forward are therefore detrended with an estimate that uses them explicitly. For every kinematic channel, not only the reference one, the trend is a moving average over three seconds, replaced by a linear interpolation across each repetition, from

two seconds before its onset to two seconds after its offset; the movement therefore contributes nothing to the estimate of its own trend. Excluding the repetitions from the estimate in this way is what allows a plain symmetric average to be used where a robust but biased estimator would otherwise be needed. The detrended signal is the difference between the filtered signal and this trend, offset back to the mean level of the initial rest so that the resting posture is preserved.

The sEMG envelopes are detrended in the same spirit but with a statistic suited to a non-negative signal. An envelope spends most of the trial close to its quiet level and departs from it in bursts during the movements, so a moving average would be pulled upwards by the bursts and would not represent the quiet level at all. The trend is therefore taken as a low quantile of the envelope over a moving window of eight seconds, the fifteenth percentile: low enough to remain below the bursts, high enough not to follow the minima of the noise. Here the quantity to be removed is the quiet level itself and not the centre of a fluctuation, so a low quantile is the appropriate estimate and not a compromise. The exact value is a choice rather than a critical threshold. As for the kinematics, the trend is interpolated across each repetition. From the result the mean envelope level measured on that participant's baseline recording is subtracted, so that a channel which is not active sits at zero rather than at its own resting amplitude, and the outcome is clipped at zero, since a negative activity has no meaning. This clipping is not neutral: it removes the negative half of whatever residual fluctuation remains around the resting level, and it is the reason the normalised envelopes of an inactive muscle are small positive numbers.

At the end of this stage each trial is represented by its detrended kinematics, its detrended and normalised envelopes, and the onset, peak and offset times of its repetitions. Of the twenty-one repetitions performed, twenty on average were retained, with a minimum of eleven in two trials; the counts for every participant and task are reported in Table B.1.

SNR

Both quantities the repetitions carry, the envelope and the joint angle, are summarised by their signal-to-noise ratio (SNR). For one participant, one task and one channel it is computed repetition by repetition as

$$\mathrm{SNR}_i = 20 \log_{10} \frac{A_i}{A_n} \quad [\mathrm{dB}], \tag{3.1}$$

where A_i is the amplitude of the signal during the i-th repetition and A_n the amplitude of the background of the same recording. The background is the set of samples that belong to no repetition, that is everything outside a guard of half a second around every onset and offset, and A_n is estimated on all of it at once rather than repetition by repetition: the spread of the ratio across repetitions then reflects the contraction alone. Below, N_R is the number of samples of the repetition under consideration and N_B the number of background samples, and the bar restricts the

expression to the samples of the i-th repetition or to those of the background, B.

For the sEMG both amplitudes are root-mean-square (RMS) levels of the normalised envelope e_t:

$$A_i = \sqrt{\frac{1}{N_R}\sum_t e_t^2\Big|_i}, \qquad A_n = \sqrt{\frac{1}{N_B}\sum_t e_t^2\Big|_B}. \tag{3.2}$$

Neither is centred, and this is deliberate: after the baseline subtraction and the clipping described above, a null envelope means no activity, so zero is not an arbitrary origin but the ideal quiet state, and the distance from it is precisely the quantity of interest. Both are taken before the detrending rather than after it, because that detrending clips at zero and would empty the quiet stretches of a channel the task does not recruit, making A_n collapse.

For the kinematics neither choice carries over, because the angle θ_t has no such origin: its value at rest is a posture, fixed by the convention of the measurement, and only its variation is signal. The background amplitude is therefore the RMS of the angle about its own mean over the background, $\bar{\theta}_B$, which is its standard deviation there, and the amplitude during a repetition is an excursion:

$$A_i = (P_{95} - P_5)\big|_i, \qquad A_n = \sqrt{\frac{1}{N_B}\sum_t (\theta_t - \bar{\theta}_B)^2\Big|_B}, \tag{3.3}$$

with P_{95} and P_5 the ninety-fifth and the fifth percentile of the angle. Percentiles are used rather than the plain range because the range rests on two single samples, and the markerless estimate does produce isolated frames far from the trajectory; a spike of that kind would enter the numerator, so that the ratio would improve on exactly the repetitions that were tracked worst. Both quantities are computed on the detrended angle, which is the one used downstream and which the clipping does not affect.

The ratio is reported at one stage of the processing only, the one described here.

3.1.6 Group averages of muscle activity

The repetitions of a task are nominally identical, but they are not identical in practice: they differ in duration, because the participant follows an announced cadence and they differ in the instant at which the movement actually begins within the announced interval. Averaging them as they are would blur precisely what the average is meant to show, since a fixed feature occurring at slightly different times in different repetitions is spread out. The repetitions are therefore aligned in time and brought to a common duration before being averaged.

Everything described in this subsection is applied to all seven tasks and to every participant–task combination for which repetitions could be located, including the four of Sec. 3.1.5 that had to be segmented on the sEMG envelope. The restrictions introduced later concern the use of the kinematics as an input to the model and leave the sEMG envelopes untouched, so there is no reason to withhold any recording from

the description of muscle activity.

Temporal alignment of the repetitions

Each repetition is extracted from the detrended envelopes over an interval that extends half a second before its onset and half a second after its offset, so that the segment contains the quiet level on both sides of the movement and the alignment has something to work on outside the burst itself.

Alignment requires a reference waveform, and one muscle per task is designated to provide it: the muscle that is the principal agonist of that movement, that is, FCR for wrist flexion, ECRL for wrist extension and radial deviation, FCU for ulnar deviation, FDS for middle-finger flexion and EDC for middle-finger extension; fist closure uses FDS, as middle-finger flexion does, the two tasks sharing the muscle that leads them. The choice matters because the alignment can only be as good as the timing information contained in the signal it is computed on, and the agonist is the channel in which the movement produces the clearest and most reproducible burst. A first template is formed as the mean of the unaligned segments of that muscle.

Each repetition is then aligned to the template in two stages. A cross-correlation provides a coarse estimate of the delay, to the nearest sample, and this estimate initialises a refinement carried out in the frequency domain, in which the delay that minimises the least-squares distance between the two waveforms is found by a Newton iteration on the phase of their Fourier transforms. The refinement returns a sub-sample delay, which matters here because the envelopes vary slowly and the information about their relative timing is not carried by individual samples. Spectral matching of this kind was introduced in electromyography to estimate the delay between signals detected at different points along the muscle fibres, from which the conduction velocity is obtained [41]; the estimator is used here for a different purpose, to remove the timing jitter between repetitions of the same movement, but the quantity it recovers is the same.

The delay estimated on the agonist is then applied to all six channels of that repetition. This is the point of the procedure: the repetitions are brought into register with one another, but the relative timing of the six muscles within a repetition is left exactly as recorded. An average built this way can still be read for the order in which muscles become active, whereas one built by aligning each muscle to its own template could not.

Time normalisation and the per-participant average

The aligned segments still differ in duration. Each is therefore resampled by linear interpolation onto a common number of samples, chosen as the mean duration of the repetitions of that trial plus the two half-second margins. Linear interpolation is equivalent to a convolution with a triangular kernel, whose frequency response attenuates the signal increasingly with frequency [42]; this is why it is avoided when the waveform carries information at frequencies close to the sampling rate. That

condition does not arise here, because the envelopes have already been low-pass filtered at 2 Hz (Sec. 3.1.4) and contain essentially nothing above that, so the choice is without consequence for these signals. The same argument covers the two further resamplings these signals undergo: the one onto the common group grid described below, and the time normalisation applied to both modalities in Sec. 3.2.3. The time axis is then expressed as a percentage of the movement, with 0 % at the onset and 100 % at the offset; the margins occupy the negative part of the axis and the part beyond 100 %, so that the rest preceding and following the movement remains visible and is not compressed into the cycle. Averaging the resampled repetitions gives, for each participant, task and muscle, a mean waveform and its standard deviation across repetitions.

Group average

The mean waveforms of the individual participants are not yet on a common axis, because the fraction of the window occupied by the movement depends on the mean duration of that participant's repetitions. Each is therefore resampled once more onto a fixed grid spanning from −15 % to 115 % of the movement, and the group average is the mean across participants of these curves, shown with the standard deviation across participants as its dispersion. The standard error of the mean is not used, for the reason given in Sec. 3.2.5: it measures the precision with which the mean is known and not the spread of the participants, which is what the band is meant to convey. Seven participants contribute to every task, that is, every participant–task combination of the dataset.

Amplitude of the active phase and group statistics

For the comparisons between muscles, each participant, task and muscle is summarised by a single amplitude: the mean of the participant's mean envelope over the interval that begins at the movement onset and ends at the mean position of the kinematic peak, that is, over the phase in which the movement is being produced, rather than the one in which the limb returns to rest. For the four combinations segmented on the sEMG envelope this position is not available, since no kinematic peak was identified, and it is fixed at the midpoint of the movement by convention; the amplitude of those combinations is therefore measured over a window defined by a convention and they are nevertheless retained in the comparisons that follow.

Within each task the six muscles are compared with a Friedman test on the seven participants, the repeated measures being the six muscles. Where the test is significant, the pairs of muscles are compared post hoc with Wilcoxon signed-rank tests, and the resulting p-values are corrected for multiple comparisons with the Benjamini–Hochberg procedure; the outcome is summarised as a compact letter display, in which two muscles sharing a letter are not distinguishable at the chosen level. The unit of analysis is therefore the participant, not the repetition: the variability across repetitions has already been absorbed into the per-participant

mean, and the test asks whether the ordering of the muscles is consistent across participants, not whether it holds within any one of them.

Purpose of the group averages

These averages are a description of measured muscle activity and they serve two distinct purposes in what follows. The first is to establish, within this dataset, which muscles are recruited and in what order during each of the tasks. The comparisons of Sec. 3.2 are made recording by recording and not on these averages, but they are read against this picture: a correlation or an amplitude computed on a muscle that does almost nothing during a task does not carry the same meaning as the same quantity computed on the agonist of that task, and the group averages are what establishes which is which. The second is to check that the recordings behave as the literature would lead one to expect.

That check is possible for the wrist, where the activity of the forearm muscles during flexion, extension and the two deviations has been described repeatedly [43], but it is only partial: those studies are largely isometric or performed against a controlled force, whereas the movements recorded here are free, unloaded, and of an amplitude chosen by the participant, so the comparison can only concern which muscles are involved and not the amplitudes at which they act. For the isolated flexion and extension of the middle finger a comparable reference is considerably less available, and the group averages of those two tasks are therefore closer to a description of new material than to a confirmation of known behaviour.

3.1.7 Kinematic inputs for the MSK model

Everything described so far concerns measured signals. This subsection describes how those signals are turned into an input for the model, which is a different kind of object: the model requires a value for every DOF of the hand at every instant, including the ones the participant was not asked to move and the ones the markerless tracking does not measure reliably. Supplying that complete configuration requires conventions and, in places, reconstruction. They are set out here in full, because the quality of everything that follows depends on them and because some of them are assumptions, and not measurements.

Sign conventions

The convention was verified directly on the recorded data rather than taken from the documentation of the extraction code, which is incorrect on this point. For the wrist, a positive *wrist_flexion* angle is a flexion and a positive *wrist_deviation* angle is a radial deviation: in the flexion trials the first quantity is positive and in the extension trials it is negative, and likewise for radial and ulnar deviation. The two wrist angles therefore already carry the sign the model expects and are used as they are, with no inversion.

For the finger joints an inversion is required. The tracking returns the angle between adjacent segments, so that an extended finger gives a value near 180° and flexion decreases it. Each finger angle is therefore converted as 180° minus the measured value, which places an extended joint at zero and increases with flexion, as the model expects. This conversion is systematic and applies to every finger channel of every trial; it should not be confused with the single reflected trial of Sec. 3.1.5, which was an isolated case of a signal recorded with the wrong sign and not a property of the convention.

The resting configuration

Every joint angle is expressed relative to the resting posture, which Sec. 3.1.1 defines as the reference configuration of the study, and the resting posture itself has to be given a numerical value in the model. For the wrist that value is not taken from the measurement, because the measurement is not trustworthy on those two angles at rest: with the hand motionless the tracking returned resting orientations that were not compatible with the observed posture, in some recordings deviations of several tens of degrees on a wrist that was plainly close to neutral. A flexion of $+5°$ and an ulnar deviation of $-7°$ are imposed instead, chosen as a plausible description of the posture the hand takes under its own weight in the seated configuration of Sec. 3.1.1.

The measured resting angle of each trial is subtracted from the signal and these two constants are added in its place, so that every trial starts from the same wrist configuration regardless of how the wrist happened to be reconstructed in that recording. Declaring the resting measurement unreliable and then using the same channels during the movement would otherwise look contradictory, but the two uses are not the same quantity. The resting posture is an absolute orientation, which is what a markerless reconstruction estimates worst, because it depends on recovering the three-dimensional attitude of foreshortened segments from a camera view. What drives the model is instead the excursion of the angle away from the rest of that same trial: the measured resting value is subtracted and replaced, so any constant error on the absolute orientation cancels in the difference and never reaches the model. What is required of the measurement, and what it can deliver, is that it track how much the wrist moved, not where it was.

The price is two assumptions: that the resting wrist posture is the same for every participant, and that its value is the one adopted here. Since the excursions are measured from this configuration and not from an absolute reference, the choice shifts the operating point of the wrist joints in the model and therefore the moment arms and the passive forces at which the movement begins; it does not affect the amplitude of the excursions themselves.

For the finger joints the resting posture is measured, from the baseline recording of that participant (Sec. 3.1.1), averaged over the annotated quiet interval. Two corrections are applied to it. The DIP joint angles are capped at 60°, because at rest the tracking returns distal flexion values that the hand plainly did not hold,

an artefact of occlusion and perspective on the smallest and most foreshortened segments; and all values are constrained within the joint limits of the model. For one participant the whole set of measured resting finger angles was replaced by a plausible set, the baseline recording having caught the hand in a clenched posture instead of a relaxed one. Both corrections are acknowledged departures from measurement and are listed in Table A.3.

A natural question is why the resting finger posture is corrected in this way rather than reconstructed from the MCP angle, as the moving joints are in the next paragraphs. The reason is that the reconstruction describes how the joints of a finger move together during flexion, and at rest there is no movement to couple. Measurement bears this out: across the seven participants, and even after the cap has been applied, the resting PIP angle of the middle finger is a median of 1.6 times the resting MCP angle and the resting DIP angle 2.3 times it, whereas the proportions used during the movement are 1.1 and 0.8. Imposing the movement proportions on the resting posture would therefore not correct it but distort it, flattening a hand that rests with markedly flexed distal joints into one that does not.

Joints not involved in the task

In the wrist tasks the fingers are not asked to move, and in the two middle-finger tasks neither the wrist nor the other digits are; in fist closure all the digits move and it is the wrist that is held still. Those joints are not left free and they are not driven by their own measurement either: they are held fixed at the participant's resting value for the whole repetition.

This is a deliberate trade-off and not a complete description of what the hand did. Those joints do move a little, and it would be surprising if they did not. Flexing or extending the wrist passively displaces the finger tendons and produces a corresponding movement of the digits, the tenodesis effect, so the fingers cannot be truly still during a wrist task [44]. Moving one finger in isolation is equally difficult for a different reason: the extrinsic flexors and extensors are multi-tendoned and their compartments are not independently controllable, so a middle-finger movement is normally accompanied by a smaller movement of the adjacent fingers [45].

The residual movement is therefore real, but it is small, and it falls on the joints that the markerless reconstruction recovers worst. Driving the model with those channels would mean imposing a movement whose measured part is mostly noise, and an inverse-dynamics procedure asked to reproduce noise answers with joint torques, and therefore with muscle activity, that has no real cause. Holding those joints fixed discards the residual movement rather than fabricating it, which is the lesser of the two errors.

Declaring a joint motionless tells the model that no force is needed to move it, whereas in the participant that joint did move a little and its movement was produced by muscle activity. The activity associated with the residual movement is therefore absent from the estimate, which is expected to fall slightly below the true one, and to

do so most in the tasks where the coupling between the fixed joints and the moving one is strongest.

Reconstruction of the distal finger joints

The joints that do move are handled differently, and this is the most consequential assumption of the subsection. Of the three joints of a moving finger, only the MCP angle is taken from the measurement. The PIP and DIP angles are not used as measured: they are reconstructed from the MCP angle by fixed proportions, 1.1 for the proximal and 0.8 for the distal joint, applied instant by instant. The same construction is used for all four fingers in the fist closure trials, and for the thumb, whose MCP and interphalangeal angles are obtained from the carpometacarpal one with proportions 1.0 and 0.6.

The reason for reconstructing these trajectories is that the measured interphalangeal angles are not usable, and this can be shown directly on the middle-finger flexion trials. In four of the seven participants the measured distal angle extends while the finger flexes, by an average of 9 to 50 degrees depending on the participant, and in the remaining three it flexes by 6 to 21 degrees. A slight extension of the distal joint while the MCP joint flexes is physiologically possible, but not of that magnitude, and not with the sign of the movement changing from one participant to another during the same task. The proximal angle is affected in a different way: its measured excursion reaches 114 and 123 degrees in two participants, beyond the ROM of that joint. Neither channel can be given to the model as an observation.

Both interphalangeal joints are reconstructed, and not only the distal one, even though the proximal angle is measured considerably better. The two joints are coupled to each other by the same tendinous apparatus, and it is that coupling which makes the reconstruction credible in the first place. Keeping the measured proximal angle and deriving only the distal one from the MCP angle would leave the ratio between the two joints as an incidental by-product of two quantities drawn from different sources, and would allow the pair to take combinations that a finger cannot physically assume. Deriving both from the same source guarantees that their mutual ratio is the physiological one and confines the questionable assumption to a single place, the link with the MCP angle. The price is that a usable measurement of the proximal joint is discarded; the choice is a trade-off between using all the available information and preserving the internal consistency of the finger chain, resolved in favour of the second.

The proportions themselves were chosen for this work and not taken from a published source; they are stated here as a modelling choice so that they can be judged as such. They are not arbitrary. The interphalangeal joints of a flexing finger are kinematically coupled through the tendinous and ligamentous apparatus [46], and the ratio between the two reconstructed joints implied by the values adopted here, $0.8/1.1 \simeq 0.73$, is close to the value of 0.75 reported for the middle finger during fist closure [47], the middle finger being the digit that moves in the two isolated

finger tasks. Only that ratio is supported by the literature: the proportion linking the reconstructed joints to the MCP angle is not, and it is on weaker ground, since the coupling between the MCP and PIP joints is markedly more variable than the coupling between the two interphalangeal joints, both across individuals and across tasks, so that no single value can represent it.

For every finger task, two of the three joint angles that drive the model are therefore not observations but a deterministic function of the third, with proportions fixed a priori and identical for every participant. No result concerning those joints can be read as a measurement of the participant's finger, and the inter-individual variability of finger coordination is, by construction, absent from the model input.

The scale on which this applies is not the same in every task. In the two middle-finger tasks a single digit drives the model, and two reconstructed angles accompany the one that is measured. In fist closure four fingers move together, so eight interphalangeal angles are derived at once, from a single pair of proportions applied to four digits whose coupling is not necessarily the same, with two further angles derived for the thumb. Nothing in the recording can be used to check the outcome, since the measured interphalangeal angles are the very quantity the reconstruction was introduced to replace. Whether a trajectory built in this way can be reproduced by the model is therefore a question the recording cannot answer and the simulation can, and it is taken up in Sec. 4.2.1.

Joint limits and repetition selection

Each angle is finally constrained within the ROM of the corresponding model joint. The constraint is a smooth algebraic saturation that leaves the signal untouched over most of the range and bends it progressively as it approaches a limit. A hard clip would introduce a corner in the target trajectory, and the controller of Sec. 3.1.9 converts the target trajectory into commanded accelerations: a corner there becomes an abrupt change in the commanded acceleration and hence a step in the estimated muscle controls, at an instant chosen by the clipping rather than by the movement itself. The smooth saturation avoids introducing that artefact.

Two automatic selections are then applied to the repetitions that survived Sec. 3.1.5, both aimed at keeping out of the model those repetitions that contain an artefact or an abrupt spurious movement, which the inverse dynamics would answer with implausibly large torques and muscle activity.

The first examines the shape of the excursion on the joint that drives the task. A clean repetition rises once to a maximum and returns, whereas an artefact or a second, unintended movement appears as a further prominent peak within the same segment; a repetition containing more than one such peak is rejected. This check has a threshold that must be declared, because it is not a tolerance but an exemption: it is applied only where the peak excursion of the repetition reaches ten degrees, and it is skipped entirely below that. In trials whose excursions are small the criterion is therefore inactive and no repetition can be rejected by it. The reason is that below

that amplitude the distinction between a second movement and a fluctuation of the resting level cannot be made reliably, so applying the rule would discard repetitions on the basis of noise; the consequence is that the protection this rule offers is not uniform across trials.

The second is a rejection of outliers, applied channel by channel. The repetitions of a trial are resampled to a common length and averaged, and each repetition is compared with that average through its RMS difference from it; a repetition whose difference exceeds the mean by more than three standard deviations, on any channel, is rejected. Unlike the tuning of Sec. 3.1.5, both criteria are automatic and identical for every participant and task. Because they are applied here and not at segmentation, the number of repetitions that reach the model is smaller than the number reported in Sec. 3.1.5; both counts are given in Table B.1.

Each accepted repetition is then resampled to one hundred points spanning the movement cycle, which is the form in which it is given to the model.

Recordings excluded from the model

One exclusion takes effect here: everything described in Sec. 3.1.6 was applied to the complete set of recordings. It concerns the four participant–task combinations of Sec. 3.1.5, whose kinematics did not allow the repetitions to be located and are for the same reason unusable as the input from which joint angles are imposed on the model.

3.1.8 Musculoskeletal model

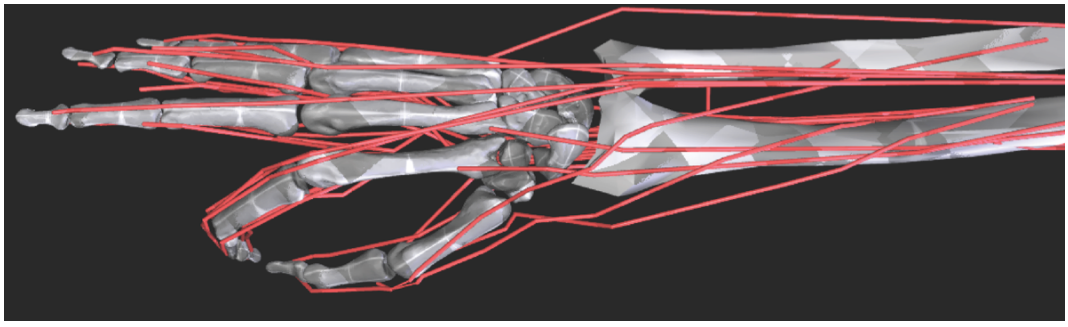


Figure 3.5: The MyoHand model: the skeletal segments of the forearm and hand, and the paths of the muscle actuators that act on them. Render from the MyoSuite documentation [29].

The model used throughout is MyoHand, one of the musculoskeletal models distributed with MyoSuite [29] and simulated by the MuJoCo physics engine [31]. It represents a right forearm and hand with twenty-three DOFs and thirty-nine muscle actuators (Fig. 3.5). Its geometry and its inertial properties are used exactly as distributed: no segment length, mass or moment arm was modified. Its joint ranges are not. The limits of eight groups of joints were widened to physiological values taken from the clinical literature, because the ranges shipped with the model

are narrower than the excursion the participants actually produced, and a target configuration that falls outside the admitted range cannot be imposed on the model at all. Which joints were widened, from which value to which, and on which reference, is reported in Table A.4. No anthropometric measurements were collected, so the model cannot be scaled to the individual and every participant is represented by the same body; and the recordings of the one left-handed participant are mirrored onto the right-hand structure, since no left-hand version of the model exists.

The actuators are Hill-type muscles. The force a muscle exerts is the sum of an active component, which depends on its activation and is modulated by the length and the shortening velocity of the fibres, and a passive component, which depends on length alone; the torque produced at the joints is obtained from these forces through the moment arms, which vary with the configuration of the hand [8].

Between the control signal and the force there is a first-order filter that represents the excitation-to-activation dynamics. Denoting by *ctrl* the control applied to a muscle and by *act* its activation, both bounded between zero and one, the activation evolves as

$$\dot{act} = \frac{ctrl - act}{\tau(ctrl, act)}, \tag{3.4}$$

where the time constant is not a single number but depends on whether the muscle is activating or relaxing and on how activated it already is:

$$\tau_A = \tau_{act}\,(0.5 + 1.5\,act), \qquad \tau_D = \frac{\tau_{deact}}{0.5 + 1.5\,act}. \tag{3.5}$$

The two nominal constants are $\tau_{act} = 10\,\text{ms}$ and $\tau_{deact} = 40\,\text{ms}$, and they are identical for all thirty-nine actuators of the model. The effective time constants are not: as the activation ranges from zero to one, the activation constant varies between 5 and 20 ms and the deactivation constant between 80 and 20 ms, which correspond to corner frequencies of roughly 8–32 Hz and 2–8 Hz. Eq. 3.4 is a property of the muscle model of the physics engine and not of the estimation method described next.

The model is redundant: thirty-nine actuators act on twenty-three DOFs, so infinitely many combinations of muscle activity produce the same movement, and any procedure that recovers activations from kinematics must add a criterion to choose among them. The clearest illustration is a hand closed softly and a hand closed tightly, which are the same posture with very different muscle activity; kinematics alone cannot distinguish them. This is the point at which the sEMG becomes relevant, first as a term of comparison (Sec. 3.2) and then as a soft constraint (Sec. 3.3).

The six muscles recorded in this study exist in the model as individual actuators, and can therefore be compared one to one: FCR, FCU, ECRL and ECU are single actuators, while the two multi-tendon muscles are represented by separate compartments, of which the ones serving the middle finger are used here, since the middle finger is the digit that moves in the finger tasks and one of the four that move in fist closure. The correspondence is one of names: a surface electrode records the summed activity of the motor units within its detection volume, including crosstalk from neighbouring muscles (Sec. 3.1.2), whereas the model actuator is an idealised

single muscle. The comparison between the two is the subject of Sec. 3.2, and this difference is one of the reasons it cannot be expected to be exact.

3.1.9 Inverse-dynamics muscle-activation estimation (MjID)

The estimation of muscle activity from the measured kinematics is performed with MuJoCo-based inverse dynamics (MjID), a method developed for the MuJoCo physics engine [48]. The present work is a direct continuation of that one, developed by Andrea Prestia at the Translational Neural Engineering laboratory of EPFL: the method is taken as it stands, without modification, and what is added here is the term of comparison it was missing, for the reason set out at the end of this subsection.

MjID is a method and not a model. Its author states that it is agnostic to the musculoskeletal model it is applied to, and demonstrates it on a hand, a leg and an elbow model; it is used here with the MyoHand model of Sec. 3.1.8. The problem it solves is the inverse of the one a physics engine normally solves. In forward dynamics, forces are the input and the resulting motion is the output; in inverse dynamics the movement is given and the forces that would produce it are sought. For a muscle-actuated system this is underdetermined, for the reason recalled in Sec. 3.1.8 and introduced in Sec. 2.4.

The three stages

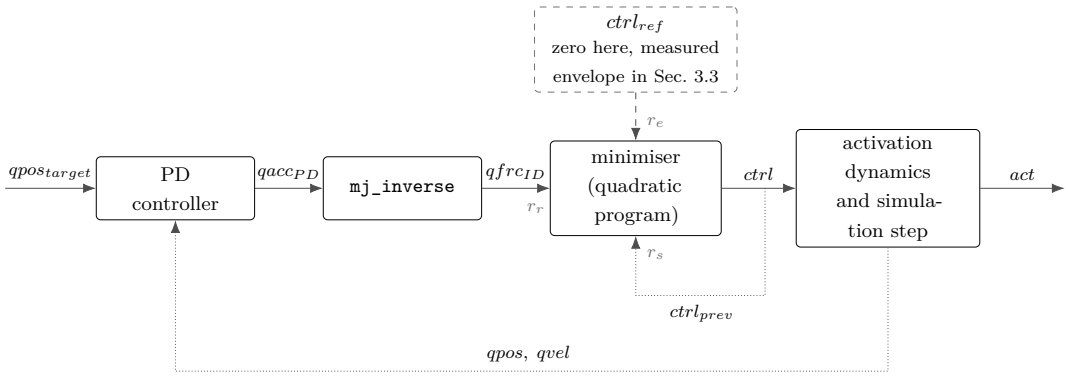


Figure 3.6: The three stages of MjID and the quantities that connect them, adapted from [48]. Solid arrows follow one simulation step, dotted arrows carry quantities from the step before, and the dashed block is the control reference.

At each simulation step the method proceeds in three stages, shown together in Fig. 3.6. First, a proportional–derivative controller compares the current configuration of the model with the target configuration, one frame of the kinematics prepared in Sec. 3.1.7, and produces the joint accelerations that would close the gap,

$$qacc_{PD} = k_p \odot (qpos_{target} - qpos) - k_v \odot qvel, \tag{3.6}$$

where $\odot$ is the element-wise product and k_p and k_v are gains defined per DOF.

The controller is placed upstream of the inverse dynamics because the torques

that result from it are stable, which is the reason its author gives for the choice [48]. It has a second consequence that matters when the target trajectory comes from a measurement rather than from a synthetic reference, as it does here: the accelerations are produced by the controller from the position error and from the model's own velocity, so the measured trajectory is never differentiated twice. A second derivative of a measured signal would be dominated by noise, whereas the accelerations obtained this way are bounded and smooth by construction.

Second, the joint torques required to realise those accelerations are computed by the inverse-dynamics routine of the physics engine. The computation accounts for the inertial properties of the segments, for gravity, and for the coupling between DOFs, that is, for the fact that accelerating one joint also produces torques at the others, since the segments of the hand are mechanically connected and their motions are not independent. This gives a target torque $qfrc_{ID}$ at every joint.

Third, a combination of muscle controls is sought whose joint torques reproduce that target. The torque produced by the actuators is

$$qfrc_{actuator}(ctrl) = MA \cdot (gain \odot act(ctrl) + bias), \tag{3.7}$$

with MA the matrix of moment arms, *gain* the active force, dependent on fibre length and velocity, *bias* the passive force, and *act* the activations obtained from the controls through the dynamics of Eq. 3.4. Because the system is redundant, this stage is posed as a minimisation with three terms,

$$\min_{ctrl} \|qfrc_{actuator}(ctrl) - qfrc_{ID}\|_{r_r}^2 + \|ctrl - ctrl_{ref}\|_{r_e}^2 + \|ctrl - ctrl_{prev}\|_{r_s}^2. \tag{3.8}$$

The three terms of the minimisation problem

The first term is replication: it asks that the muscles produce the torques the movement requires. Its weight r_r is defined per DOF rather than as a single number, so that the accuracy demanded of the model can be set separately for different joints; how that freedom is used in this work, and with what values, is described in Sec. 3.1.10.

The second is the energy term. It penalises the distance between the controls and a reference $ctrl_{ref}$, weighted by r_e. When no reference is available the reference is zero, and the term then selects, among the combinations that produce the required torques, the one of least overall activity. This is a choice, not a measurement: it resolves the redundancy of Sec. 3.1.8 by assuming that the participant activates the muscles as little as the movement allows, which is a common assumption and is systematically wrong whenever the participant co-contracts. The whole of Sec. 3.3 exists because the reference need not be zero: if an sEMG measurement is available it can be placed there, and the estimation is then no longer free to assume minimal activity.

The third is the smoothness term, weighted by r_s, which penalises the change of each control with respect to the previous simulation step and suppresses oscillations

in the solution.

Numerical formulation

The minimisation is solved at every simulation step, so it has to be fast. It is cast as a quadratic program and solved with an operator-splitting solver [49]. The activation dynamics of Eq. 3.4 are nonlinear and would prevent that formulation, so within the minimisation, and only there, the author of the method replaces them by a linear approximation [48]; the physics engine continues to integrate the exact dynamics when the estimated controls are applied. The estimated activations reported in this work are those produced by the model in response to the estimated controls.

The controls are confined to the interval the actuators admit, and, at each step, to a neighbourhood of the current activation, so that no control can jump discontinuously from one step to the next. This second bound acts on the solution instead of acting on the cost, and it is what keeps the estimate from tracking the required torque through changes faster than a muscle can produce.

The missing term of comparison

MjID has been demonstrated on MyoHand on a synthetic target trajectory, where the kinematic replication is accurate, but the quality of the estimated *activations* was left unassessed for that model, for a stated reason: no selective measurement of the activity of the individual forearm muscles was available. Surface recordings suffer from crosstalk, and intramuscular recordings are too invasive to be used during free dynamic movements. The present work addresses exactly that gap. It does not propose a new estimation method; it supplies the missing term of comparison, on measured kinematics, accepting the limitations of surface recording and declaring them, so that the activations MjID produces on a hand model can be confronted with a measurement rather than only with the movement they reproduce.

3.1.10 Subject-specific parameter optimisation

The method of Sec. 3.1.9 has free parameters: the two controller gains k_p and k_v of Eq. 3.6, and the three regularisation weights r_r, r_e and r_s of Eq. 3.8. They are not physiological quantities and cannot be measured; they govern how closely the model is asked to follow the target trajectory and how strongly the redundancy is resolved towards low and smooth activity. Values for this model are reported by the author of the method [48], but they were obtained on a synthetic trajectory, whereas the trajectories used here are measured, noisier and different from one participant to another. The parameters were therefore fitted.

Six values are fitted: k_p, k_v, r_e, r_s, and r_r split into two values, one for the wrist DOFs and one for the finger flexion joints. The split reflects the fact that the two groups do not have to be reproduced with the same accuracy: in a wrist task the model is driven by a single wrist angle and the fingers only have to stay where they are put, whereas in a task that moves the digits the joints of the moving digits carry

the whole movement. Fitting one weight per group instead of one per joint keeps the number of free values small enough to be identified from the available data.

The criterion is kinematic. For a candidate set of parameters the whole procedure is run on the trajectory, the resulting controls are applied to the model, and the configuration the model actually reaches is compared with the target; the cost is the root-mean-square error (RMSE) between the two, accumulated over the trajectory. The cost is computed over all DOFs of the model and not only over those the task drives, so that a parameter set which follows the intended movement while making some other joint tremble is penalised for it. And it contains no term involving the sEMG: the parameters are chosen without ever looking at the measured muscle activity, which is what allows the comparison of Sec. 3.2 to be a comparison and not a fit.

Two mild penalties discourage the energy and smoothness weights from collapsing towards zero. Without them the search is free to remove both regularisations, which improves the kinematic cost (the controls become entirely free to chase the required torques) while producing solutions in which the estimated activation of a muscle jumps up and down between consecutive samples without the movement requiring anything of the sort. Those solutions reproduce the trajectory and mean nothing physiologically, so the region is penalised.

The search is a differential evolution over bounded ranges for the six values [50]. A global method is used because the cost cannot be differentiated: evaluating it once requires running a full simulation of the trajectory, so no analytical gradient exists, and approximating one numerically would cost one additional simulation per parameter at every step of the search. Differential evolution needs neither a gradient nor an initial guess. The bounds and the settings of the search are reported in Table A.5 and Table A.6, and the values the search converged to in Table A.7. Those values are given for reproducibility and not for interpretation: several parameters settle at the edge of their admitted range rather than inside it, most often the proportional gain at its lower bound. Where that happens the value reflects the range that was allowed as much as the data, so the parameters should not be read as estimates of anything.

The fit is performed independently for every participant and for every task, and yields one parameter set per participant and task. It is the most permissive choice of this stage. The alternative, one set per participant, would ask a single tuning to serve movements as different as a wrist deviation and a finger flexion, which involve different joints, different inertias and different required accuracies. Fitting per task avoids that compromise, at the cost of allowing the procedure more freedom, and of preventing any statement about the parameters themselves being compared across tasks. Since the parameters are not physiological, no such statement is attempted.

The target used for the fit is not a single repetition but the average of all the repetitions of that trial that were accepted in Sec. 3.1.7, resampled to a common length. Averaging first makes the fit insensitive to the peculiarities of any one repetition and gives it the movement the participant performed typically in that trial.

It also means that the trajectory used to choose the parameters and the trajectories on which the model is subsequently run are not independent: the same repetitions contribute to both. The parameters are tuned on the material they are then applied to, so the kinematic accuracy reported for a trial is an optimistic estimate of what the same parameters would achieve on a trial they had not seen. It does not affect the comparison with the sEMG, which is the object of Sec. 3.2, because no sEMG enters the fit at any point.

3.2 Validation against sEMG

The procedure of Sec. 3.1 produces, for every repetition of every trial, an estimate of the activation of each of the thirty-nine muscles of the model, obtained from movement alone. This section describes how those estimates are placed beside the muscle activity that was actually recorded, and what is measured in the comparison.

No recorded muscle activity entered the procedure at any point: neither the model inputs of Sec. 3.1.7 nor the parameter fit of Sec. 3.1.10 use it, so the envelopes serve here as an independent term of comparison. Six muscles were recorded, however, and thirty-nine are estimated, so the comparison covers only a fraction of what the model produces; what the model does on the remaining thirty-three is the subject of Sec. 3.4. Fist closure appears neither in this section nor in the two that follow. It is driven through the model like the other tasks; whether it is carried further is decided on the outcome of that step, which is reported in Sec. 4.2.1.

3.2.1 Reproduction of the imposed trajectories

The activations examined in this section are those the model produces while reproducing a measured trajectory. If it fails to reproduce that trajectory, its activations describe a different movement from the one the participant performed, and the movement of interest is lost at the outset. How closely the model follows the trajectory it is given is therefore established first.

This check is not a validation. A validation confronts an estimate with an independent measurement of the same quantity, which is what the comparison with the envelopes does. Here there is no independent measurement: the trajectory the model is asked to follow is the same one that was used to build its input, so the agreement between the two expresses how well the model is able to reproduce a given movement from the inputs it receives.

For each repetition the configuration reached by the model is compared with the target configuration, joint by joint, over the same normalised cycle described in Sec. 3.2.3. The metrics are defined in Sec. 3.2.4 together with those of the muscle comparison.

Driven and held degrees of freedom

The model has twenty-three DOFs. Seventeen of them have a measured counterpart through the correspondence of Sec. 3.1.7; the remaining six (forearm pronation–supination and the five abduction DOFs of the thumb and fingers) are not recovered by the markerless tracking and are therefore not measured at all.

The target is not of the same nature for all of them, and the distinction governs which metric is meaningful:

- **Driven** DOFs are those the task requires the participant to move, and their target is the measured trajectory itself. There is one per wrist task, flexion–extension or radial–ulnar deviation according to the task; there are three in the two middle-finger tasks, namely the MCP, PIP and DIP joints of that finger; and there are fifteen in fist closure, the three joints of each of the four fingers together with the three of the thumb.
- **Held** DOFs are all the others, and their target is a constant. For the measured ones the constant is that participant's own resting angle, estimated as described in Sec. 3.1.7; for the six unmeasured ones it is the neutral value of the model. In both cases the model is asked to keep the joint where it was placed instead of moving it.

Held DOFs are reported in three groups: the wrist DOF that the task does not drive, the finger and thumb joints, and the six unmeasured ones. The three are not equivalent. The first two are held at a posture that was observed on the participant, whereas the last are held at a posture that was assumed, so a departure from the target means something different in the two cases.

Reporting by group

The error is reported separately for each of these groups and is never averaged across them. On a driven joint the error measures how well the model follows a movement; on a held joint it measures how well it refrains from moving. The two answer different questions and a single value combining them would describe neither.

The restriction concerns the aggregation across joints only. Repetitions are still averaged within a participant and participants are still aggregated with one another.

The grouping also determines which metric can be used. The RMSE is reported for every group: it is expressed in degrees, so it is interpretable on a joint that moves and on one that should not, and it is comparable across joints. The coefficient of determination is reported for the driven DOFs only: since it is one minus the ratio between the residual sum of squares and the total sum of squares of the target about its own mean, on a held joint its denominator is identically zero (the target is a constant, so it does not deviate from its own mean at any sample). The coefficient is therefore not defined there, which is a different situation from a poor value.

3.2.2 Modelled activations and measured envelopes

The comparison places the activation the model assigns to a muscle beside the normalised envelope recorded from that muscle, for the six channels that were instrumented, each matched with the actuator of the same name in the model.

The two are nominally the same quantity, a dimensionless fraction of the maximal activation of that muscle: in one case produced by a Hill-type actuator, in the other estimated from the recordings. On the modelled side, thirty-nine actuators act on twenty-three DOFs, so many patterns of activation produce the same movement, and the criterion of Sec. 3.1.9 selects among them by penalising muscle activity and its variation; methods of this family are documented to underestimate the activity of muscles that do not contribute directly to the required moment, since activating them adds to the cost without reducing the tracking error [51]. On the measured side the scale rests on the normalisation reference, which is itself an estimate: for most channels it is the largest amplitude reached during the session instead of a maximum voluntary contraction.

How the difference in scale between the two is handled in the comparison is described in Sec. 3.2.4.

3.2.3 Time normalisation and alignment

Repetitions differ in duration and the movement does not fall at the same point of every repetition, so the signals to be compared are not directly superimposable. They are aligned with one another by the same two operations already described in Sec. 3.1.6, that is, resampling onto a common normalised axis and alignment of the repetitions on a designated anchor channel, so only what differs at this stage is detailed here.

In particular, two things differ. The normalised axis spans the movement alone, from onset to offset, sampled at one hundred points: the half-second margins that Sec. 3.1.6 retains on either side, so that the resting level remains visible in the group averages, are not needed here, because the comparison concerns the movement. The interval and the grid are the same for every signal, but each is resampled onto it from its own sampling rate, which is 2 kHz for the envelopes and about 30 Hz for the kinematics.

The second difference is that four sets of curves are involved: the recorded envelopes and the joint angles given to the model, which are measured, and the estimated activations and the joint angles the model actually reaches, which are produced by it. Each of the four is aligned independently, and two anchors are used: a muscle channel for the two sets of muscle signals and a joint channel for the two sets of angles. Within each set the shift is estimated on the anchor and applied to all its channels, so that the relative timing of the muscles, and of the joints, is left as it was.

Aligning each set to its own average removes the jitter and leaves the lag. The jitter is the repetition-to-repetition variability in when the movement falls within the

cycle; it has no systematic sign, and removing it is the purpose of the operation. The lag is a systematic difference in timing between the two modalities, such as the delay between electrical activity and the movement it produces; it survives the alignment, because each set keeps the mean phase of its own repetitions and only the dispersion around that mean is reduced.

On the kinematic side the anchor is the same at both stages, and is in every case the DOF the task drives: the wrist angle in the plane of the movement for the four wrist tasks, and the MCP angle of the middle finger for the finger tasks. On the muscle side one choice does change. The channel designated as the anchor is not the same at the two stages for two of the six tasks: in wrist flexion and wrist extension the first stage uses the radial flexor and the radial extensor, the second the ulnar flexor and the finger extensor. Recomputing the comparison of Sec. 3.2 for those two tasks with the anchors of the first stage changes the reported correlations by at most 0.03, and by 0.01 or less on the muscles that carry them; no value changes sign and no ordering between muscles is altered.

3.2.4 Metrics

This section defines the metrics used in the two comparisons of Sec. 3.2. The two answer different questions, as set out above, and do not take the same metrics throughout, although some are shared; which applies where is stated below.

Reporting a measure of agreement in shape together with a measure of the magnitude of the error is the usual practice when the output of a musculoskeletal model is confronted with an independently obtained quantity: the work that established the EMG-driven approach validates its predicted joint moments against moments obtained from inverse dynamics, and reports a coefficient of determination, a slope and an intercept, and a mean residual error [37]. The structure adopted here is the same, with one difference. Where the coefficient is accompanied by a slope, a systematic difference in scale between the two quantities appears in the slope and leaves the coefficient largely unaffected; in that work the mean slope is 0.89 while the mean coefficient is 0.91. No slope is reported here: the coefficient of Eq. 3.10 is computed about the measurement itself, so a systematic difference in scale is charged to it and it is accompanied instead by the gain of Eq. 3.12, which measures that difference directly.

In what follows a repetition-averaged curve is sampled at $N = 100$ points over the normalised cycle; m_i denotes the measured quantity and e_i the estimated one, both referring to the i-th sample of the cycle, and $\bar{m}$ and $\bar{e}$ their means over the cycle.

Reproduction of the trajectories

Agreement between the configuration reached by the model and the target configuration is measured by the RMSE, computed for every joint and every repetition and expressed in degrees,

$$\mathrm{RMSE} = \sqrt{\frac{1}{N}\sum_{i=1}^{N}(m_i - e_i)^2}, \tag{3.9}$$

and reported by group of DOFs as set out in Sec. 3.2.1.

For the driven DOFs the coefficient of determination is also reported,

$$R^2 = 1 - \frac{\sum_{i=1}^{N}(m_i - e_i)^2}{\sum_{i=1}^{N}(m_i - \bar{m})^2}, \tag{3.10}$$

which expresses the residual as a fraction of the variation the target itself displays. Its denominator is what restricts it to the driven DOFs: on a held joint the target is constant, every term of that sum is zero, and the quantity is undefined.

Comparison with the envelopes

Three quantities are computed for each participant, task and monitored muscle, on the average of that participant's aligned repetitions.

The Pearson correlation

$$r = \frac{\sum_{i=1}^{N}(m_i - \bar{m})(e_i - \bar{e})}{\sqrt{\sum_{i=1}^{N}(m_i - \bar{m})^2 \sum_{i=1}^{N}(e_i - \bar{e})^2}} \tag{3.11}$$

summarises the agreement in shape and timing. It is invariant to a multiplicative constant, and therefore unaffected by any difference in scale. It is the only one of the three that can be applied to every channel regardless of how strongly that muscle participates in the task, and it is accordingly the primary metric of this comparison.

The coefficient of determination of Eq. 3.10, computed with m_i the measured envelope and e_i the estimated muscle activation, summarises the agreement in shape and amplitude together. The amplitude is part of what a validation should examine, so the quantity is reported, with two cautions: it is unbounded below, so a single poor case can dominate an arithmetic mean and it conflates two distinct failures: a curve of the right shape on the wrong scale and a curve of the wrong shape are both penalised, and the coefficient alone does not distinguish them.

The gain

$$g = \frac{\max_i m_i}{\max_i e_i} \tag{3.12}$$

gives an indication of the difference in scale between the two signals. Reporting it alongside the coefficient of determination allows a negative coefficient to be attributed to its cause: a gain far from unity indicates a disagreement in amplitude, whereas a gain close to unity accompanied by a negative coefficient indicates a disagreement in shape. Like the coefficient of determination, it is unbounded above.

Muscle involvement

The metrics above say how well a muscle is reproduced but not whether that muscle does anything in the task, and the second is needed to read the first. For each participant and task, the involvement of the monitored muscle j is the peak of its average normalised envelope as a share of the same quantity summed over the six monitored channels,

$$\iota_j = \frac{\max_i m_i^{(j)}}{\sum_{k=1}^{6} \max_i m_i^{(k)}} , \tag{3.13}$$

so that the six values sum to unity and express which muscles carry the recorded activity of that task. It is computed on the measured envelopes alone and is therefore independent of the model; the same definition is used in Sec. 3.3.6.

3.2.5 Aggregation and statistics

Every metric is computed on the individual repetition. The values of a participant are then reduced to a single central value, and those are aggregated across participants. The alternative, averaging the repetitions of a participant first and computing the metric once on the resulting mean curve, was not adopted.

A correlation computed on averaged observations depends on how many observations were averaged, growing with their number according to the Spearman–Brown relation, whereas the average of the individual correlations does not; the latter is therefore the quantity that represents the experimental unit actually used, and it has the smaller sampling variability of the two [52]. The distinction matters here because the number of repetitions that reaches the model is not the same for every participant and task, ranging from seven to twenty-one (Appendix B, Table B.1), so a quantity that depends on that number would not be comparable across them.

Therefore the metric is computed on each repetition, the repetitions of a participant are summarised by a central value, and the participants are aggregated. Where a distribution is displayed the unit is the participant: repetitions are never pooled across participants, since repetitions of the same participant are not independent observations and pooling them would let a participant who contributed more of them weigh more heavily.

Relation to the figures

Alongside the choice just described on how the values are aggregated, the curves shown in this work are averages: the repetitions of a participant are averaged, and the participants are averaged with one another. The metrics are not computed on those curves. In most cases the two nearly coincide, but they need not, and a figure showing two averaged curves in poor agreement is compatible with individual repetitions that agree well, or the reverse; Fig. B.2 shows both situations on real data. Therefore it is not necessarily true that the reported values can be read off the figures.

Conventions adopted throughout this work

The conventions below apply to every quantity reported from this point onwards, including those of Sec. 3.3 and Sec. 3.4.

Quantities that are bounded and approximately symmetric are summarised as mean ± standard deviation; quantities that are unbounded or markedly asymmetric, such as the coefficient of determination and the gain, are summarised by median and interquartile range (IQR), defined from the first to the third quartile. This follows the usual recommendation for reporting biomedical data [53].

Which of the two applies was not decided by testing each set of values for normality. With seven participants such a test has almost no power: it would fail to reject in nearly every case, including those where the distribution is plainly not normal. The choice rests instead on the properties of the quantities themselves, which are known in advance and do not depend on the sample, and on the distributions actually obtained here, many of which are markedly asymmetric. Median and IQR are accordingly the summary used for most of the values reported in this work.

Number of participants

Seven participants contribute to each task, with two exceptions that follow from the exclusions of Sec. 3.1.5: ulnar deviation, for which four participants remain, and middle-finger extension, for which six do. The counts of repetitions behind each of these averages are reported in Table B.1.

3.3 EMG-informed MjID: a feasibility test

Sec. 3.2 compares an estimate (the activations from the model) with a recording (the sEMG envelopes) that did not produce it. Here the recorded envelopes are given to the estimation procedure instead, so the independence between what enters the estimate and what it is compared with is lost, and none of the comparisons reported in this section is a validation. What the section examines is how closely the estimation can follow the two signals it is given at once, the imposed kinematics and the recorded envelopes.

The subsections that follow set out how the recorded envelopes are turned into a reference for the solver, how that reference enters the estimation problem, which quantities are fitted again once it is present, and which quantities are compared. The last one describes a control experiment in which one recorded muscle at a time is withheld from the reference, which tests whether the information supplied on the remaining channels reaches the muscle that was left out.

3.3.1 Measured envelopes as an input

In Sec. 3.2 the recorded envelopes are the term against which the estimate is judged, and nothing of them enters the estimate itself. This section reverses that relation: the same envelopes are supplied to the estimation as an input, and what is examined

is whether the movement and the measured muscle activity can be accounted for at the same time. The envelope is used here in the form it was given in Sec. 3.1.4, since the linear envelope is the conventional representation of recorded activity when that activity is used as an excitation term [37]; no other form of the signal was considered.

The question is therefore one of compatibility rather than of accuracy. The procedure of Sec. 3.1.9 chooses, among the many activation patterns that reproduce a given movement, the one that a cost function prefers. The recorded activity is added to that cost as a reference, and not as a direct input in the way the kinematics is: the solver is asked to stay close to it, not required to match it. Once the parameters are fitted again in the presence of that term, as described in Sec. 3.3.4, the reference may be followed without the trajectory being reproduced any less accurately, or it may not be followed as expected and the reproduction of the trajectory may deteriorate.

The purpose of this section is not to evaluate the activations the model produces. The recorded signal is now both an input to the estimation and the term it would be compared with, so agreement with it would say only that the solver followed the reference it was given. What is quantified instead is how closely the controls follow the reference they are given, and how much accuracy in reproducing the movement is given up in order to follow it. The activations themselves are not examined here: they are obtained from the controls through the deterministic first-order dynamics of Eq. 3.4, so comparing them would repeat the comparison already made on the controls.

3.3.2 Construction of the control reference

The reference given to the solver is a reference for the muscle *controls*, not for the activations, as anticipated in Sec. 3.1.4. Specifically, a reference is built for all six recorded muscles and for every repetition, from the detrended and normalised envelopes of Sec. 3.1.4 resampled onto the control time base. A reference is also built for the repetition on which the parameters are fitted, by averaging the repetitions of that trial in the same way as the kinematic target of Sec. 3.1.10.

Nothing else is altered. The joint angles supplied to the model, both the per-repetition trajectories and the averaged one used by the fit, are the same as those of Sec. 3.1.7, without the fist task as anticipated in Sec. 3.2. Any difference between the results of this section and those of Sec. 3.2 is therefore attributable to the reference alone.

The reference has still to be given an amplitude. As anticipated in Sec. 3.2.2, the amplitude of a normalised envelope and that of a modelled control may not be on the same scale, so two choices are compared here.

- **Uniform anchoring.** A single factor is applied to all six channels, the median of the six ratios between the amplitude of the control that the model produced for a muscle in the unconstrained solution of Sec. 3.1 and the amplitude of the envelope recorded from it. The absolute scale is thereby anchored to the typical behaviour of the unconstrained solution, while the relative amplitudes between

muscles are preserved. An earlier version applied each of those six ratios to its own channel, which anchored every muscle to its own previous amplitude and in doing so removed the relative amplitudes between them, that is, precisely the information that the involvement of Sec. 3.2.4 expresses; it was replaced by the single factor for that reason and is not among the configurations compared below.

- **Direct use of the normalised envelope.** The envelope is used as it is, being already expressed as a fraction of a reference contraction (Sec. 3.1.4), so no factor is applied. Unlike the uniform anchoring, this reference does not depend on anything the model has produced in the previous simulations.

3.3.3 Supplying the reference to the solver

The reference is not imposed as a hard constraint, and it requires no change to the formulation of the MjID algorithm. The minimisation of Eq. 3.8 already contains the term $\|ctrl - ctrl_{ref}\|^2_{r_e}$, whose reference is zero in Sec. 3.1 and which there selects the solution of least overall activity. Here $ctrl_{ref}$ is the recorded envelope of Sec. 3.3.2, and the same term therefore expresses a preference for controls close to the measured activity instead of zero. Its weight r_e, which Sec. 3.1.10 treats as one of the fitted parameters, becomes the quantity that decides how strongly the recording is allowed to shape the solution.

This has a consequence for the fit. The criterion by which the parameters are chosen is kinematic, as in Sec. 3.1.10, and the recorded activity does not appear in it. An optimisation free to choose r_e would therefore be free to set it near zero, since following the envelope can make the trajectory harder to reproduce; at that value the estimation would be the unconstrained one of Sec. 3.2 again, and the section would no longer be testing anything. A lower bound on r_e prevents it.

A second decision concerns the muscles that were not recorded. The weight acts on all thirty-nine actuators, so raising it to make the recorded muscles follow their reference also penalises the activity of the thirty-three that have no reference and are simply drawn towards zero. Two treatments were implemented: a single weight for all actuators, and a weight split into two values, one for the recorded muscles and one for the others, the second either fixed at a low value or left to the fit.

3.3.4 Re-optimisation of the parameters

The parameters are fitted again with the reference present, by the procedure of Sec. 3.1.10 and with the same algorithm and criterion. Six values are fitted where the weight is single and seven where it is split; the weight r_e is searched on a logarithmic scale, since it spans several orders of magnitude.

One quantity is deliberately excluded from the fit: the scale factor that gives the reference its amplitude. In the anchored modality of Sec. 3.3.2 that factor is derived from the controls of the unconstrained solution; it is computed once, before the fit begins, and held fixed throughout. Recomputing it while the parameters are being

adjusted would make the reference depend on the controls that are being adjusted to follow it, so that the target would move at every iteration.

3.3.5 Comparisons and metrics

Several estimations with a reference were run, differing in how the reference is built and in how much weight it is given. This subsection states which ones, what they are compared with, and by which quantities.

Configurations evaluated

Every configuration is compared with the unconstrained solution of Sec. 3.1, which is the same estimate examined in Sec. 3.2 and serves here as the baseline. Five configurations were run in addition to it, so that six sets of simulations are compared in all. Each of the five has its own parameters: the fit of Sec. 3.3.4 is repeated for every configuration and for every participant, and the simulations of that configuration are then run with the parameters it produced.

The five configurations differ from one another in two "axes", and only one axis at a time. The first axis is the amplitude given to the reference, treated in Sec. 3.3.2; the second is the way the weight r_e is distributed between the recorded muscles and the thirty-three that have no reference, treated in Sec. 3.3.3. One of them combines uniform anchoring with a single weight for all actuators. Three differ from it only along the second axis, the weight of the unrecorded muscles being fixed at either of two low values or left to the fit. The fifth differs from it only along the first, the reference being the normalised envelope used directly.

No configuration combines the direct use of the envelope with a split weight, so the comparison speaks to each axis separately and not to their interaction. And since the reference depends only on the first axis, the five configurations share two references between them: uniform anchoring for four of them and the normalised envelope for the fifth.

The comparison has a purpose beyond describing the five configurations. One of them is carried forward to the control experiment of Sec. 3.3.6 and to Sec. 3.4, and this comparison is what selects it. The selection is argued on two aspects: how each reference is constructed, which Sec. 3.3.2 has already established, and how the configurations behave, which is what follows from the quantities defined below.

Metrics

Four quantities are computed, two on each side of the trade-off, and all of them use the definitions and the conventions of Sec. 3.2.4 and Sec. 3.2.5 without modification.

Two measure how closely the solution follows the reference it was given: the correlation r and the coefficient of determination R^2 between the controls the solver produced and $ctrl_{ref}$, on the six recorded muscles. This is not a validation, for the reason given in Sec. 3.3.1: it measures whether the solver did what it was asked. It has no counterpart in Sec. 3.2, where controls are never compared with a reference.

Both are computed for the baseline as well, whose controls were produced with no reference at all: the values they take there are the agreement that arises without any reference being supplied, and the five configurations have to be read against them.

The other two measure what the agreement costs, on the comparison between the configuration reached by the model and the target trajectory that Sec. 3.2.1 applies to the unconstrained solution. The first is the coefficient of determination on the driven DOFs, defined as in Sec. 3.2.4. The second is the tracking error over all the DOFs, reported in normalised form as the normalised root-mean-square error (nRMSE),

$$nRMSE = \frac{RMSE}{ROM}, \tag{3.14}$$

where ROM is the anatomical ROM of that DOF in the model, taken from Table A.4. Dividing by the range makes the error comparable between joints whose excursions differ by an order of magnitude, so that a single value can summarise all the DOFs of a repetition. Here the difference from the baseline is again the relevant quantity, and the target is common to all six sets of simulations.

Read together the four describe a trade-off, and the purpose of the comparison is to establish how much agreement with the control reference can be obtained, and how much accuracy in reproducing the movement is given up to obtain it.

Statistical comparison

Two comparisons are made, and both work on one value per participant, obtained by the aggregation of Sec. 3.2.5.

The first compares the six sets of simulations with one another, on each of the four quantities. The six are first compared together with the Friedman test, which is the non-parametric test for repeated measurements on the same participants and does not assume normality, and the comparison then proceeds on a set of pairs fixed in advance with the Wilcoxon signed-rank test. The pairs are based on the design: each configuration against the baseline that shares its reference, and each configuration against the one with uniform anchoring and a single weight, which is the configuration the other four differ from along one axis only. Pairs that differ along both axes are not tested, since the result could not be attributed to either.

The second compares the selected configuration with the baseline alone, and is resolved by task rather than pooled over tasks: muscle by muscle on the agreement between controls and reference, then by group of DOFs, using the four groups of Sec. 3.2.1, and finally on the three driven joints of the two finger tasks taken one at a time, which is where the cost of the constraint concentrates. Each comparison is a Wilcoxon signed-rank test on the participants present in both conditions.

In both comparisons the p values are adjusted for multiplicity by the Benjamini–Hochberg correction, and the threshold is 0.05 on the adjusted value. In the second the adjustment is applied within each task, since the question asked there is a question about one task at a time.

3.3.6 Robustness to a missing constraint: the LOMO test

A further question is asked of the configuration that the comparison above selects and focuses about how much of what the model produces on a recorded muscle comes from the reference given to that muscle, and how much from the references given to the other five. A control experiment is therefore run, in which one recorded muscle at a time is withheld from the reference and the model is asked to account for it without having been shown it. The procedure is referred to in what follows as the leave-one-muscle-out (LOMO) test.

The experiment uses the single configuration selected in Sec. 3.3.5, and its design does not depend on which one that is. For every movement, each of the six recorded muscles is removed in turn from the reference, the remaining five are supplied as before with the parameters already fitted for that configuration, and no quantity is refitted. The activity the solver assigns to the withheld muscle is then compared with the envelope recorded from that muscle, which took no part in producing it.

The four quantities of Sec. 3.3.5 are used again without change: r and R^2 between the controls and the reference on the muscle side, the nRMSE of Eq. 3.14 and the coefficient of determination on the driven DOFs on the kinematic side. On the muscle side a LOMO run contains two kinds of muscle that have to be kept apart, the one that was withheld and the five that remained constrained, so four conditions are compared: the baseline, the selected configuration, and the two halves of the LOMO run. On the kinematic side no such distinction exists, because a simulation reproduces one trajectory however its reference was composed, and there the LOMO run is therefore a single condition beside the other two.

That comparison is read between two levels that are known in advance. The lower one is the baseline of Sec. 3.3.5, which is what the model produces for that muscle when no reference is supplied at all. The upper one is the selected configuration, in which the same muscle is constrained by its own reference. A withheld muscle recovered at the lower level would mean that the constraint on the other five carries no information about it; one recovered at the upper level would mean that its own reference was never needed. Where the result falls between the two is the quantity of interest, and the five muscles that remained constrained establish that removing one did not disturb the others.

These conditions are compared with the same tests as in Sec. 3.3.5, on one value per participant: the Friedman test across the conditions of each panel, then the Wilcoxon signed-rank test on the pairs of interest, which are the withheld muscle against the baseline and against the selected configuration, with the p values adjusted by the Benjamini–Hochberg procedure.

The result is further read against the involvement of the withheld muscle, defined in Sec. 3.2.4. This is examined muscle by muscle, with one value per participant and per task. A straight line fitted to the task medians is drawn in each panel as a visual summary of the direction the points take, and the relation itself is quantified by the Spearman rank correlation between r and the involvement, computed on all

the per-participant points and reported with its p value. The rank correlation is used because each muscle contributes few points and the involvement is bounded between zero and one with a skewed distribution, so that a single extreme value would weigh heavily on a correlation computed on the values themselves.

3.4 Mechanical characterisation of the model

Sec. 3.3 supplies the model with six recorded channels and asks whether it can follow them while still reproducing the movement. It says nothing about the thirty-three muscles that were not recorded, which the model also drives and about which no measurement exists.

For those muscles no ground truth exists, so no measure of agreement can be computed. What remains possible is to describe what the model makes them do, and to base that description on a quantity that is comparable between the recorded muscles and the unrecorded ones: the joint torque. This section is therefore a description of the mechanics of the model instead of a further test of its accuracy, and the question it builds towards is how much of those mechanics the six recorded channels account for.

The question is the coverage problem raised in Sec. 2.5 and left open there. It is not a question about the accuracy of the estimate but about its reach: even an estimate that agreed perfectly with all six recordings would leave unconstrained whatever the other thirty-three muscles do.

3.4.1 Joint torque as the measure of what a muscle does

The analysis of this section is carried out on joint torques instead of activations, and the reason is that activations do not answer the question.

An activation is a fraction of a muscle's own maximum. Two muscles at the same activation are not doing the same thing, because what a muscle does at a joint is its force multiplied by its moment arm at that joint, and both differ by an order of magnitude across the thirty-nine actuators of the model. Above all, activations cannot be added: the sum of the activations of a set of muscles is not a quantity with a mechanical meaning, so it cannot express what fraction of anything that set covers. Torques can be added, and their sum over the muscles crossing a joint is the torque at that joint, which is what moves it. The torque is also the only quantity on which a recorded muscle and an unrecorded one can be placed side by side, since it expresses the action of both in the same units at the same DOF.

Joint moments are commonly the quantity against which an MSK model is validated, by comparing the moments it predicts with those obtained from inverse dynamics, as established by reference works in the field [37, 12], but that is not their role here, since no measured moment is available for these recordings. The torques of this section are internal quantities of the model, read from the forces its muscles produced and from its geometry, and they serve to decompose what the model does

among its muscles and its joints.

3.4.2 Joint torque read from the simulation

The torque that muscle j exerts about DOF k is

$$\tau_{jk} = a_{jk} F_j \,, \tag{3.15}$$

with a_{jk} the moment arm of that muscle about that DOF in the posture the model holds at that instant and F_j the force the muscle develops there; the force already includes the dependence of the Hill model on activation, on fibre length and on the rate of shortening.

Both factors come from the simulations of Sec. 3.1 and Sec. 3.3. The configuration of the model and the force of every actuator are available at every integration sub-step, that is every two milliseconds and sixteen times per video frame, and the moment arm follows from the configuration alone, so Eq. 3.15 is evaluated exactly at every sub-step. Sixteen sub-steps of two milliseconds cover thirty-two milliseconds against a video frame lasting between 33.3 and 33.6 ms, so each trajectory is traversed in about four per cent less simulated time than it was recorded in, by the same factor for every muscle and every DOF.

The torque time courses are then resampled to a common grid of one hundred and one points over the movement cycle, as in Sec. 3.2.3; quantities are aggregated as set out in Sec. 3.2.5, so over the samples of a repetition, then over the repetitions of a participant, then across participants. Unlike most of the values reported in this work, the quantities of this section are shares bounded between zero and one hundred, so the convention of Sec. 3.2.5 summarises them by the mean across participants and its standard deviation rather than by median and IQR. The number of participants is seven except where Sec. 3.1.5 excludes a combination, which leaves four for ulnar deviation and six for middle-finger extension.

3.4.3 Anatomical capacity of the model

Before the torques produced during the recorded movements are examined, the anatomical capacity of the model is established. For each muscle and each DOF the capacity is Eq. 3.15 evaluated with the activation set to one, that is the magnitude of the moment arm multiplied by the force the muscle develops at full activation in that posture, averaged over twenty thousand postures drawn at random from the joint ranges of the model. Because the moment arm varies with the posture, that average depends on which postures are sampled. The postures actually assumed along the recorded trajectories were therefore compared with the random sample, to establish whether the capacity computed over the range of the model is representative of the configurations in which the model is operated in the rest of this work.

The DOFs are then grouped into regions, namely wrist flexion and extension, radial and ulnar deviation, forearm rotation, each of the four fingers and the thumb, and the capacity along every muscle is normalised to its own total, so that the result

reads as the distribution of a muscle's mechanical action across the hand. Only the mean over the sampled postures is reported, without its dispersion.

This quantity expresses how much each muscle weighs on each group of DOFs and depends only on the geometry and the strength parameters of the model: no participant and no recording enter it, and the recorded trajectories are used only to check that the sampled postures are representative. It is reported as a property of the model rather than as a result about the hand.

3.4.4 Partition between movement and stabilisation

Not all of the torque a hand produces serves to move it. In a wrist task the fingers are held still, and holding them still requires torque; the same is true of the wrist during a finger task. The second quantity of this section separates the two.

Here the net torque at each joint is used, instead of the contribution of individual muscles, and the twenty-three DOFs are assigned to three groups for each task: those the task moves, the remaining wrist DOFs, and the remaining finger and thumb DOFs. The magnitude of the net torque is averaged over the cycle and summed within each group, and each group is expressed as a fraction of the total, so that the three fractions sum to unity for every participant.

The partition is computed both for the unconstrained solution of Sec. 3.1, which serves here as the baseline as it does in Sec. 3.3.5, and for the selected EMG-informed configuration of Sec. 3.3. Comparing the two shows whether supplying the recording changes how the effort is distributed between moving and holding.

3.4.5 Share of the torque carried by the recorded muscles

This is the quantity that answers the question of the section. For a given task, the torque produced on a DOF is decomposed among the thirty-nine muscles, and the share carried by the six that were recorded is what the six electrodes account for.

It is computed separately on the DOFs the task drives and on those it holds still. The six recorded muscles are for the most part wrist muscles, so they may account for a large part of the torque that moves the wrist and for a small part of the torque that holds the fingers still; a single value averaged over both would hide that difference.

On the driven side the decomposition is carried out on the primary driven DOF, which is the wrist DOF of the task or, in the finger tasks, the MCP joint of the middle finger. Each muscle is expressed as a percentage of the summed magnitude of the torque of all thirty-nine, after averaging its torque along the repetition. On the held DOFs the same is done after summing, for each muscle, the magnitudes of its torques over the DOFs of the region being stabilised. Summing magnitudes instead of signed values is deliberate: two muscles that oppose one another both spend effort, and a signed sum would report their joint action as nothing.

The same decomposition is also reported resolved over the movement cycle and over each driven DOF separately, which shows when in the movement the unrecorded muscles contribute rather than only how much.

Both decompositions group the muscles by role as well as by whether they were recorded, which gives four categories. A muscle is called agonist on a driven DOF when its action there has the same sign as the excursion of that DOF in that task, and antagonist otherwise. Muscles whose torque on that DOF falls below a small threshold are excluded from the assignment, since a muscle with almost no moment arm there has no role to play. The assignment is made once, on the EMG-informed solution, and applied unchanged to the baseline, so that a difference between the two is a difference in what the muscles did and not in how they were labelled.

The aggregate decomposition reports magnitudes, so its values carry no sign. The decomposition resolved over the cycle reports the torques themselves, and there the sign is physical and is not adjusted for presentation: in a task whose excursion is negative, such as extension or ulnar deviation, the agonists appear with negative torque.

3.4.6 Co-contraction

The last quantity summarises how much of the activity of the recorded muscles is spent by agonists and antagonists acting against each other, and unlike the quantities above it is computed on the activations.

The activation is the level at which the mechanics described in this section is generated, since the torques of Eq. 3.15 are produced by the forces the activations develop, and it is also the level at which the model was set beside the recording in Sec. 3.2. Sec. 3.3.1 declines to examine the activations there, because they are a deterministic filter of the controls and comparing them would repeat a comparison already made on the controls. That objection does not apply here: the index below is a ratio between two groups of the same signal rather than a comparison with the reference, and the comparison it serves is with the unconstrained solution, which has no reference.

The index used is the co-contraction index of Falconer and Winter [54]. Writing A_i and B_i for the summed activation of the agonists and of the antagonists among the six recorded muscles at the i-th of the N points of the normalised cycle, it is

$$\mathrm{CCI} = \frac{\sum_{i=1}^{N} 2\min\left(A_i, B_i\right)}{\sum_{i=1}^{N}\left(A_i + B_i\right)}\,, \tag{3.16}$$

which is zero when only one of the two groups is ever active and one when the two are equal throughout. The two groups are formed by the rule of Sec. 3.4.5, applied to the primary DOF: a recorded muscle counts as an agonist when its torque there has the same sign as the excursion of that DOF in the task, as an antagonist when it has the opposite sign, and it is left out of both groups when its torque there falls below the threshold. Being a ratio, the index describes how the activity is divided between the two groups and not how much activity there is: two muscles balanced at a low level and two balanced at a high level receive the same value.

It is computed on three signals: the recorded envelopes, the unconstrained solution and the EMG-informed solution. The recorded envelopes are a property

of the participant and serve as the term of comparison. Two comparisons carry information: the recorded envelopes against the unconstrained solution, which asks whether the model left to itself produces co-contraction of the kind the participants produced, and the unconstrained solution against the EMG-informed one, which asks how far supplying the recording moves it. The third comparison, the EMG-informed solution against the recording, is not independent. The envelopes of those same six muscles were supplied to that solution as the reference for its controls, of which the activations are a filtered image, so agreement is in part built in, for the reason Sec. 3.3.1 sets out. The reference is the recorded envelope, but following it exactly is not required of the solver, which is only asked to stay close to the reference, so a solution that follows it is expected to replicate that value, but this is not trivial.

Chapter 4

Results

This chapter reports the results in the order in which Chapter 3 presents the methods that produce them, one section against one. Sec. 4.1 corresponds to Sec. 3.1 and reports the quality of the signals through the SNR (Sec. 3.1.5) and what the muscle activity looks like on the sEMG envelopes (Sec. 3.1.6). Sec. 4.2 corresponds to Sec. 3.2: it first establishes how closely the model reproduces the imposed movement of each task, and then, on that basis, compares the activations it estimates with the measured envelopes. Sec. 4.3 corresponds to Sec. 3.3 and assesses whether the EMG-informed approach is feasible and, through the LOMO test, how robust it is. Sec. 4.4 corresponds to Sec. 3.4 and describes the mechanics of the model, including the thirty-three muscles no electrode observes. The conventions for summarising and comparing quantities are the ones stated in Sec. 3.2.5, except where a section states otherwise.

4.1 The recorded signals

This section reports what the recordings look like once the processing of Sec. 3.1 has been applied to them, before any of it reaches the model: how much signal the six sEMG channels and the kinematic channels carry in each task, what the detrending removed, and what the measured muscle activity looks like across the seven participants.

Throughout the section every quantity is summarised as the median across participants of the per-participant median, with the interquartile range across participants alongside, following the hierarchy of Sec. 3.2.5. The number of participants is seven, except for ulnar deviation, where it is four, and finger extension, where it is six, following the exclusions of Sec. 3.1.5.

4.1.1 Signal quality

Table 4.1 reports the SNR of the sEMG envelopes for every task and channel, and Table 4.2 the same quantity for the kinematic channels that each task guides. Both follow the definition of Sec. 3.1.5.

Table 4.1: Signal-to-noise ratio of the sEMG envelopes, in decibels. Median across the seven participants, interquartile range below. Defined in Sec. 3.1.5.

Task	FDS	FCR	FCU	EDC	ECRL	ECU
Wrist flexion	15.9	15.7	17.3	1.2	0.2	4.6
	(14.3, 16.9)	(12.3, 21.6)	(13.9, 21.7)	(−2.2, 3.7)	(−0.3, 2.9)	(3.8, 9.2)
Wrist extension	9.3	9.2	8.5	12.3	10.0	13.5
	(5.6, 10.6)	(5.9, 10.6)	(7.3, 11.4)	(10.9, 13.3)	(9.2, 12.6)	(11.0, 14.2)
Radial deviation	10.2	11.8	4.8	4.6	9.3	5.0
	(9.7, 14.2)	(8.8, 12.6)	(1.6, 6.9)	(2.2, 6.0)	(6.2, 12.3)	(2.8, 5.8)
Ulnar deviation	7.0	8.0	14.2	8.7	3.6	13.2
	(5.8, 8.4)	(5.6, 8.8)	(10.0, 16.7)	(3.4, 11.2)	(2.5, 4.4)	(12.3, 15.0)
Finger flexion	10.6	10.1	5.0	3.5	2.9	1.3
	(6.0, 13.6)	(5.1, 10.7)	(0.4, 8.1)	(2.3, 4.6)	(1.3, 4.5)	(0.5, 4.6)
Finger extension	7.3	4.1	7.1	10.7	5.8	3.6
	(2.0, 13.4)	(2.6, 7.3)	(4.9, 8.9)	(8.4, 12.5)	(3.9, 8.7)	(−0.7, 4.2)
Fist closure	15.1	9.1	14.6	5.8	9.0	9.3
	(9.8, 16.9)	(6.5, 10.4)	(10.0, 15.0)	(4.9, 6.3)	(7.4, 11.3)	(8.0, 11.7)

The sEMG table uses all seven participants in every task. The four participant–task combinations excluded elsewhere are excluded because their kinematics could not be segmented, which leaves their sEMG unaffected; they already contribute to the group averages for the same reason (Sec. 3.1.6). The kinematic table, in contrast, follows the exclusions set in Sec. 3.1.7.

The sEMG ratio is computed before the detrending since after the clipping the RMS of the resting phase falls by two orders of magnitude on the channels a task does not recruit, which would make the table describe the recordings in a misleading way.

Read across a row, Table 4.1 recovers the expected agonists of each task. The three flexor channels reach 15.7 dB to 17.3 dB in wrist flexion while the three extensor channels remain between 0.2 dB and 4.6 dB, and the ordering reverses in wrist extension. The clearest case is ulnar deviation, where the two channels that stand out are FCU at 14.2 dB and ECU at 13.2 dB, a flexor and an extensor that deviate the wrist together while cancelling each other in flexion and extension. The two finger tasks are the weakest of the set: their strongest channel reaches 10.6 dB and 10.7 dB, against 11.8 dB in radial deviation and more than 13 dB in each of the other three wrist tasks.

Fist closure stands apart from the other six. It is the only task in which every channel carries signal, from 5.8 dB on EDC to 15.1 dB on FDS, with no channel left at its resting level, which is consistent with a movement that engages the whole hand instead of one digit or one wrist axis.

A value near zero does not mean that the channel was noisy: it means that in that task the muscle stayed at its own resting level, and the reading proposed here is that it was not recruited. ECRL in wrist flexion, at 0.2 dB, is such a case. The

ratio therefore conflates recruitment with recording quality, and the table is best read as a map of which muscles each task involves. As such it anticipates Sec. 4.2: where a channel carries little signal, any disagreement between the model and the measurement concerns a quantity that was barely measured in the first place.

About kinematics, only the DOF that each task moves is reported: the remaining wrist DOF is held still by instruction. The ratios of Table 4.2 are far higher than those of the sEMG, between about 20 dB and 38 dB, as expected of a direct measurement of the quantity of interest rather than of an indirect correlate of it. Within the two finger tasks the DIP is the weakest of the three joints, at 19.9 dB and 22.0 dB against 24.3 dB to 38.5 dB for the MCP and the PIP: it is the smallest segment of the hand and the most often occluded. Among the wrist channels the lowest values are those of ulnar deviation.

Table 4.2: Signal-to-noise ratio of the kinematic channels, in decibels, for the DOF each task moves. Median across participants, interquartile range below.

Task	Wrist F/E	Wrist R/U	MCP3	PIP3	DIP3
Wrist flexion	27.9 (26.4, 34.9)	–	–	–	–
Wrist extension	30.1 (27.5, 31.8)	–	–	–	–
Radial deviation	–	25.5 (23.3, 29.4)	–	–	–
Ulnar deviation	–	23.7 (22.4, 24.0)	–	–	–
Finger flexion	–	–	35.7 (29.0, 39.2)	38.5 (31.8, 39.7)	19.9 (19.4, 29.8)
Finger extension	–	–	30.6 (28.6, 31.0)	24.3 (19.5, 27.2)	22.0 (19.0, 25.5)

Fist closure guides fifteen joints so its kinematic channels are reported separately in Table 4.3. The ratios are of the same order as those of the other tasks, between 21.2 dB and 44.5 dB, and they are ordered along each digit: the MCP joints of the four fingers are the strongest channels of the whole set, between 39.7 dB and 44.5 dB, while their interphalangeal joints fall to between 24.4 dB and 30.9 dB, and the thumb is the weakest digit throughout. The pattern follows how far each joint is from the back of the hand, which is what the two views see best.

Wrist extension is the only condition in which the tracking loses frames systematically: five of the seven participants have gaps there, up to 17.9 % of the samples of a trial, whereas every other task stays below one per cent. The interpolation of Sec. 3.1.3 fills them, but the input of wrist extension rests on a reconstruction of a larger part of the recording than the others do.

Table 4.3: Signal-to-noise ratio of the kinematic channels in fist closure, in decibels, for the fifteen DOFs the task guides. Median across participants, interquartile range below.

Digit	Proximal	Middle	Distal
Thumb	21.2	27.6	23.7
	(20.0, 25.5)	(24.9, 30.1)	(21.5, 26.5)
Index	39.7	28.2	26.8
	(29.3, 41.7)	(25.5, 36.8)	(25.5, 29.2)
Middle	42.2	26.6	30.9
	(32.7, 47.3)	(25.2, 30.3)	(24.1, 32.7)
Ring	44.5	24.8	26.2
	(33.4, 48.1)	(24.4, 29.2)	(22.7, 27.8)
Little	43.6	24.4	24.9
	(36.0, 45.7)	(24.1, 29.5)	(24.2, 26.9)

4.1.2 Effect of the detrending

Fig. 4.1 shows what the detrending of Sec. 3.1.5 does, on one kinematic channel (wrist flexion–extension) of one trial (wrist extension). The resting posture drifts by more than ten degrees over the four minutes of the recording, so the same movement starts and ends at a different angle depending on when in the trial it was performed. The estimated trend follows that drift without following the repetitions, which are interpolated across, and subtracting it returns the trial to a common resting level while leaving the excursion of each repetition where it was.

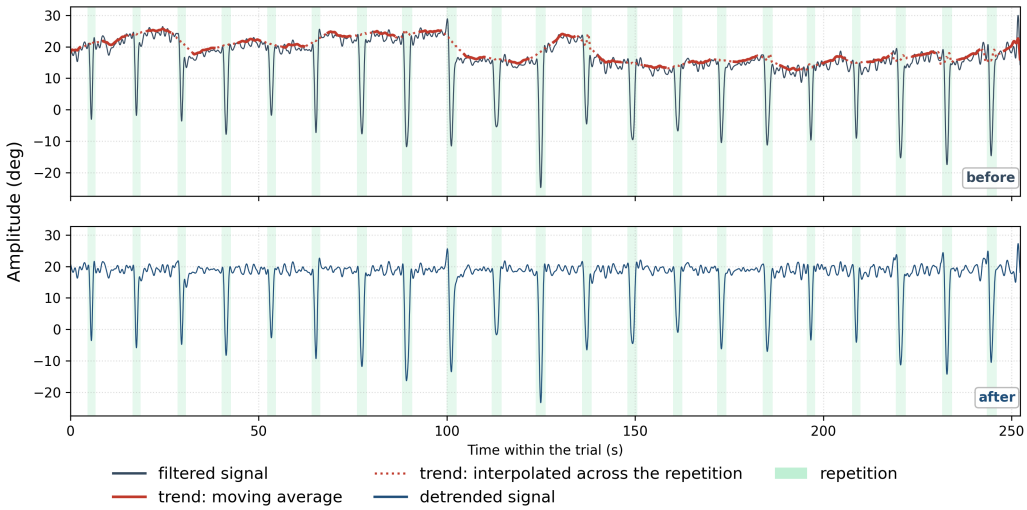


Figure 4.1: Effect of the detrending on one kinematic channel (wrist flexion–extension) of one trial (wrist extension) on a random subject. The trend is drawn superimposed on the signal it is subtracted from.

The size of the correction differs across the hand. The two wrist DOFs drift by 4.4° to 11.0° over a trial, whichever task is being performed. The finger joints drift by more: from 3.5° at the MCP up to 24.6° at the DIP in finger extension, and,

across the fifteen joints that fist closure guides, from 3.7° at the thumb to 22.0° at the PIP joint of the little finger, with a median of 11.6°. The interphalangeal angles are those whose resting value the tracking estimates least stably over four minutes, and this is what the correction removes. These are the measured angles: the ones that later drive the model at the two interphalangeal joints are reconstructed from the MCP (Sec. 3.1.7), and the drift reported here on the measured interphalangeal angles is one of the reasons that reconstruction exists.

On the sEMG the same operation is applied for the same reason, with a statistic suited to a non-negative signal. The amount removed is small in absolute terms, a median of 0.89 % of the MVC, and it follows the channel rather than the task: FCR sits between 0.51 and 0.60 % in every one of the seven tasks and ECU between 1.83 and 3.30 %, whatever that task asks of either muscle, the six channels together spanning that same range from 0.51 to 3.30 %. What is removed is therefore an offset that belongs to the electrode and its site, not a component of the contraction.

4.1.3 Group averages of the envelopes

Fig. 4.2 shows the group averages. Every task produces a single burst that rises after the onset, peaks near the position of the kinematic peak and returns to the resting level before the offset, so the activity is synchronised with the movement instead of being spread over the cycle.

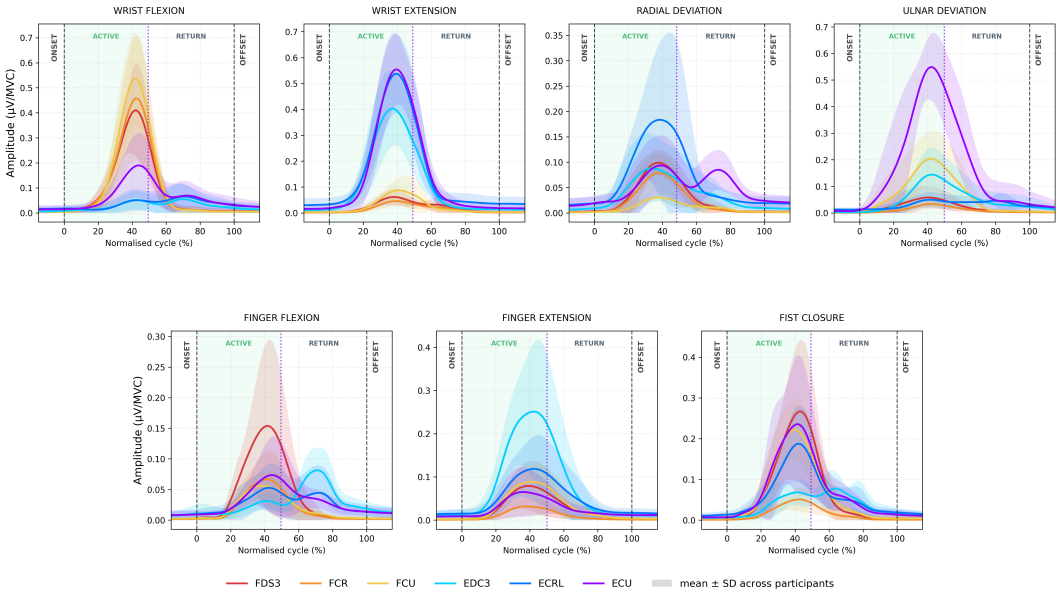


Figure 4.2: Group averages of the sEMG envelopes, one panel per task, seven participants in each. Dashed lines mark the onset and the offset of the repetition, the dotted line the mean position of the kinematic peak.

The muscle that leads each panel is the one anatomy predicts. Wrist flexion is led by the three flexors, FCU highest at 0.54 of the MVC; wrist extension by ECU and ECRL, which overlap almost exactly; ulnar deviation by ECU, at 0.55 against 0.20 for the next channel, which makes it the panel where a single muscle separates

most clearly from the others. Radial deviation behaves differently: ECRL leads, at 0.18, but the remaining five channels lie between 0.08 and 0.10, so the leader stands out by a factor of two while the rest of the hand stays at a comparable level. The two finger tasks are led by FDS and EDC respectively.

The tasks divide into two groups by amplitude rather than by joint. Wrist flexion, wrist extension and ulnar deviation all peak at about half of the MVC; radial deviation and the two finger tasks stay between 0.15 and 0.25, with radial deviation the lowest of the seven. The dispersion across participants follows the same divide: relative to the peak it is smallest in the three high-amplitude tasks and largest where the amplitude is low, so the tasks that demand little are also those in which the participants differ most in how strongly they drive the leading muscle.

In some tasks a second and smaller peak appears during the return phase, carried by the muscle that opposes the movement. It is modest, well below the peak of the agonist, but it is present in wrist flexion, radial deviation and finger flexion.

In fist closure the four wrist muscles reach comparable amplitudes, FDS at 0.27, ECU at 0.24, FCU at 0.22 and ECRL at 0.19, while EDC and FCR stay below 0.08. Flexors and extensors of the wrist are thus active together while the finger extensor is not, which is what closing the whole hand asks of the forearm.

Differences between muscles

Fig. 4.3 summarises the same recordings as amplitudes in the active phase, and tests whether the six muscles differ within each task. The Friedman test is significant in all seven tasks, from $p < 0.0001$ in wrist flexion, wrist extension and ulnar deviation to $p = 0.044$ in finger flexion. The post-hoc comparisons are more selective.

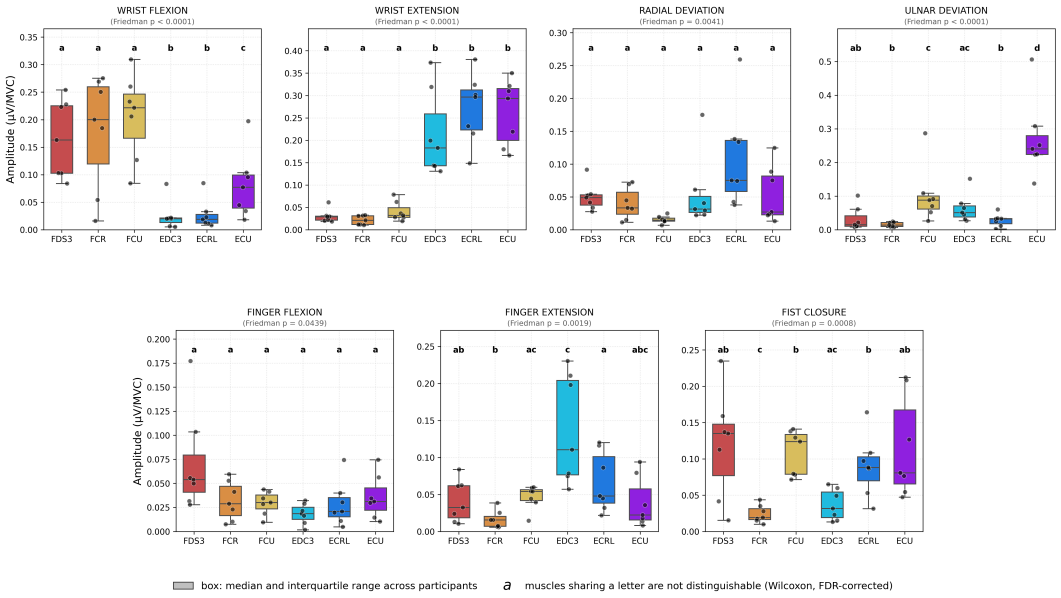


Figure 4.3: Mean amplitude in the active phase, one point per participant, with the Friedman test for each task and the compact letter display of the post-hoc comparisons. Seven participants in each panel.

In wrist flexion, wrist extension and ulnar deviation the muscles separate into distinguishable groups, and those groups correspond to the agonists and antagonists of each task. In radial deviation and in finger flexion no pair survives the correction for multiple comparisons, so the muscles are not distinguishable from one another even though the global test rejects equality. These are the same two tasks that carry the smallest amplitudes and the widest dispersion, and the three observations are consistent with one another: where the demand is low, the participants differ enough that the ordering between muscles cannot be established on seven of them.

The two remaining tasks separate their muscles in different ways. In finger extension EDC leads with a median amplitude of 0.11 of the MVC, and is distinguishable from FDS, FCR and ECRL, though not from the two remaining wrist muscles. Fist closure has no single leader, as FDS and FCU reach 0.13 and 0.12 with ECRL and ECU following at 0.09 and 0.08: closing the fist thus recruits the flexors of the fingers and the wrist muscles of both directions together.

Read as a whole, the group averages behave as the literature on forearm muscle function would lead one to expect, which is the check Sec. 3.1.6 sets out as one of their purposes: each task recruits the muscles that cross the joint it moves, in the direction it moves it, and the two ulnar deviators act together. This is the reference against which the estimates of the model are judged in Sec. 4.2.

4.2 Validation against sEMG

This section reports the comparison set out in Sec. 3.2. It first establishes that the model reproduces the trajectories it was asked to follow, since an error there would propagate into everything that follows, and then confronts the activations it estimates with the envelopes measured on the six monitored muscles. Fist closure is carried only as far as the first of those two steps: the reproduction of its trajectories, reported in Sec. 4.2.1, is what withdraws it from every following analysis.

Metrics are computed on the individual repetition, summarised within each participant and then aggregated across participants, as explained in Sec. 3.2.5. They are all summarised by their median and IQR across participants.

4.2.1 Reproduction of the imposed trajectories

Agreement between the configuration reached by the model and the target is summarised by Table 4.4, and for every task it is reported by four group of DOFs. On the driven DOF of the four wrist tasks it stays between 0.43° and 0.69° and the coefficient of determination is between 0.986 and 0.997. The imposed wrist trajectory is reproduced almost exactly.

On the held DOFs the question is the opposite one: whether the model refrains from moving a joint that should not move. Across all tasks, the wrist DOF that the task does not drive is held to within 0.09° to 0.47°, below the error on the driven DOF of the same task in every case; the finger and thumb joints depart by 1.01° to

Table 4.4: Tracking error on each task by group of DOFs: RMSE in degrees, and coefficient of determination, defined on the driven DOFs only. The thumb is counted among the fingers. Median across participants, IQR below.

Task	Driven	Held, wrist	Held, fingers	Held, unmeasured	R^2 (driven)
Wrist flexion	0.59	0.38	2.08	1.28	0.995
	(0.40, 1.98)	(0.32, 0.39)	(1.47, 2.94)	(1.19, 1.64)	(0.992, 0.998)
Wrist extension	0.50	0.47	2.27	1.44	0.997
	(0.38, 0.70)	(0.40, 0.64)	(1.87, 2.61)	(1.40, 1.71)	(0.997, 1.000)
Radial deviation	0.69	0.12	1.01	1.31	0.995
	(0.59, 0.84)	(0.10, 0.17)	(0.91, 1.62)	(1.15, 1.40)	(0.987, 0.997)
Ulnar deviation	0.43	0.09	1.03	0.89	0.986
	(0.42, 0.97)	(0.08, 0.17)	(0.92, 1.19)	(0.81, 1.00)	(0.948, 0.989)
Finger flexion	7.58	0.28	1.12	1.16	0.812
	(5.88, 9.42)	(0.23, 0.29)	(1.06, 1.25)	(0.95, 1.26)	(−0.246, 0.883)
Finger extension	1.94	0.22	1.05	0.99	0.895
	(1.67, 2.15)	(0.19, 0.27)	(1.00, 1.08)	(0.88, 1.12)	(0.253, 0.924)
Fist closure	7.16	0.60	—	4.13	0.756
	(6.80, 12.91)	(0.58, 3.02)		(2.78, 5.77)	(−0.225, 0.887)

2.27°, and the six unmeasured DOFs by 0.89° to 1.44°, even reaching 4.13° in the fist task.

The held finger joints depart most in the two tasks that flex and extend the wrist, 2.08° and 2.27°, and least in the two deviations, 1.01° and 1.03°, even though none of the four involves moving the fingers. Holding the fingers still therefore costs more when the wrist rotates about one axis than about the other.

The tasks that move the fingers are less accurate, and the error is not spread evenly over the joints they drive. On the driven DOFs taken together the coefficient of determination falls from the 0.986 to 0.997 of the wrist tasks to 0.812 in finger flexion and 0.895 in finger extension, and to 0.756 in fist closure, the lowest of the seven. The coefficient separates the joints more sharply than the RMSE does, and it is the quantity reported joint by joint in Table 4.5. There, a negative value means that the simulated angle departs from the target by more than the target departs from its own mean, so that the model does not follow the imposed trajectory in any useful sense.

Read down the table, the two middle-finger tasks fail on one joint out of three, the DIP, at −0.432 and −0.023, and are followed closely on the two joints above it. Fist closure fails clearly on six of the fifteen joints it drives. Three of those six are the DIP joints of the middle, the ring and the little finger, at −0.295, −0.033 and 0.042, which is the same joint that fails in the two retained tasks; the fingers themselves are otherwise followed as closely as in those tasks, and the index is followed more closely still, at 0.645 on its distal joint.

The thumb is what sets fist closure apart. Its three joints reach −1.838, −14.395 and −38.015, values of a different order from anything else in the table. The thumb

Table 4.5: Coefficient of determination between the imposed and the reproduced trajectory, for every driven DOFs of the three tasks that move the fingers; median across participants, IQR below.

	Proximal	Middle	Distal
Finger flexion	0.894	0.812	−0.432
	(0.888, 0.902)	(0.784, 0.849)	(−0.461, −0.323)
Finger extension	0.895	0.940	−0.023
	(0.854, 0.907)	(0.922, 0.961)	(−0.160, 0.178)
Fist closure			
Thumb	−1.838	−14.395	−38.015
	(−3.363, 0.094)	(−28.728, −7.881)	(−85.431, −13.720)
Index	0.867	0.941	0.645
	(0.851, 0.924)	(0.921, 0.959)	(0.548, 0.683)
Middle	0.885	0.756	−0.295
	(0.813, 0.912)	(0.659, 0.783)	(−0.732, −0.027)
Ring	0.909	0.815	−0.033
	(0.908, 0.931)	(0.784, 0.862)	(−0.576, 0.466)
Little	0.857	0.906	0.042
	(0.833, 0.899)	(0.842, 0.918)	(−0.209, 0.140)

is asked to stay almost still while the fingers close, with an imposed excursion of four to eight degrees, and the model moves it by between one and a half and three and a half times that amount, depending on the joint. These three numbers need to be read with caution, because the coefficient of determination divides by the variance of the target: where the target barely moves, as here, the ratio grows without the absolute error having to be large, so the three values establish that the thumb is not held where it was asked to be.

Watching the simulated trajectories of fist closure, the fingers do not close together: they lag one another visibly, and the motion carries sudden jerks, most markedly during the return to the resting posture. The metrics used throughout this section compare each joint with its own target, one at a time, and are therefore blind to a disagreement that lies in the relation between joints; this is reported here as an observation, and the mechanism that would account for it is taken up in Chapter 5.

Taken together, these results are the reason fist closure is not carried onto on further analyses. The task is reproduced fairly well on the fingers and not on the thumb, and there is a lack of coordination between the fingers during the movement. Muscle activity estimated from a movement the model did not reproduce would describe that movement rather than the one that was recorded, and then a disagreement with the envelopes could not be told apart from it. Fist closure therefore appears in this subsection and in no other: everything that follows, in this section and in the two after it, concerns the six tasks that remain.

Overall, the tracking error is small on the wrist, larger on the fingers and

concentrated on their most distal joint, and in fist closure it extends to the thumb and to the coordination between the digits. Except on the DIP of the two finger tasks and on the whole fist task, the model reaches the configuration it was asked to reach, so what follows can be read as a statement about the activations on the retained movements.

4.2.2 Agreement in shape

Table 4.6 reports the Pearson correlation between the estimated activations and the measured envelopes, for every task and monitored muscle. The involvement of Eq. 3.13 is given beside it, as a measure of how much of the recorded activity of the task that muscle carries.

Table 4.6: Pearson correlation and muscle involvement (second line of each task) between the estimated activations and the measured envelopes; median across participants, IQR below. Negative correlations are shaded blue and positive ones red.

Task	FDS	FCR	FCU	EDC	ECRL	ECU
Wrist flexion	0.76 (0.71, 0.77)	−0.17 (−0.20, −0.10)	0.86 (0.84, 0.89)	0.59 (0.25, 0.66)	0.45 (0.36, 0.54)	0.33 (0.21, 0.41)
involvement	0.23 (0.19, 0.30)	0.25 (0.22, 0.29)	0.32 (0.26, 0.35)	0.03 (0.02, 0.05)	0.04 (0.03, 0.07)	0.08 (0.07, 0.11)
Wrist extension	−0.01 (−0.19, 0.24)	0.46 (−0.07, 0.71)	−0.47 (−0.63, −0.28)	0.36 (0.21, 0.50)	−0.43 (−0.52, −0.04)	0.12 (−0.04, 0.37)
involvement	0.03 (0.03, 0.05)	0.03 (0.02, 0.03)	0.04 (0.04, 0.06)	0.25 (0.20, 0.27)	0.31 (0.29, 0.33)	0.33 (0.31, 0.34)
Radial deviation	0.55 (0.35, 0.60)	0.16 (0.01, 0.22)	0.04 (−0.15, 0.07)	0.56 (0.44, 0.64)	0.49 (0.36, 0.54)	0.24 (0.12, 0.57)
involvement	0.17 (0.14, 0.19)	0.12 (0.09, 0.20)	0.05 (0.05, 0.07)	0.13 (0.10, 0.17)	0.23 (0.20, 0.38)	0.20 (0.13, 0.24)
Ulnar deviation	0.26 (0.02, 0.36)	0.14 (0.05, 0.23)	0.73 (0.63, 0.81)	0.62 (0.39, 0.62)	0.13 (0.00, 0.25)	0.80 (0.76, 0.83)
involvement	0.07 (0.04, 0.11)	0.03 (0.02, 0.03)	0.24 (0.17, 0.31)	0.10 (0.07, 0.12)	0.06 (0.05, 0.07)	0.51 (0.47, 0.55)
Finger flexion	0.61 (0.49, 0.72)	−0.21 (−0.28, −0.14)	−0.23 (−0.29, −0.20)	0.77 (0.73, 0.77)	0.23 (−0.01, 0.32)	0.34 (0.13, 0.37)
involvement	0.32 (0.23, 0.38)	0.14 (0.10, 0.16)	0.10 (0.09, 0.19)	0.16 (0.11, 0.21)	0.10 (0.09, 0.14)	0.15 (0.10, 0.19)
Finger extension	0.31 (0.26, 0.52)	0.48 (0.41, 0.60)	0.48 (0.46, 0.52)	0.85 (0.78, 0.89)	−0.04 (−0.07, 0.01)	−0.42 (−0.58, −0.39)
involvement	0.15 (0.07, 0.16)	0.05 (0.03, 0.06)	0.13 (0.10, 0.19)	0.35 (0.28, 0.43)	0.18 (0.12, 0.21)	0.11 (0.04, 0.22)

The muscle that carries most of the recorded activity is reproduced well in four tasks out of six. In wrist flexion FCU, with an involvement of 0.32, reaches a correlation of 0.86; in ulnar deviation ECU, which carries the largest involvement recorded in any task (0.51) reaches 0.80; in middle-finger extension EDC reaches 0.85; in middle-finger flexion FDS reaches 0.61. In all four the IQR stays clear of zero, so no participant reverses the sign, and the correlation is higher than that of any other muscle of the same task except in finger flexion, discussed below.

Radial deviation is weaker, as ECRL reaches 0.49, which is consistent with what Sec. 4.1.3 reports for that task: the leading muscle stands out from the others by a factor of two only, and the recorded activity is spread across the hand.

Wrist extension is the exception, and it is not a weak task in terms of peak amplitude as a fraction of the MVC. Its envelopes are among the largest recorded, with ECU and ECRL at involvements of 0.33 and 0.31, yet ECU correlates at 0.12

and ECRL at −0.43, with the sign reversed in five participants out of seven. The model does not reproduce the timing of the extensors in the one task that is named after them.

Negative correlations occur on channels that carry little of the activity of their task: FCU in wrist extension, at an involvement of 0.04, correlates at −0.47, and FCR and FCU in middle-finger flexion, at 0.14 and 0.10, correlate at −0.21 and −0.23. Interestingly, in wrist flexion FCR carries an involvement of 0.25, second only to FCU in the same task, and still correlates at −0.17. There, the envelopes show that the muscle carries a large part of the activity of the task, while the model does not follow it.

Additionally, in middle-finger flexion the extensor EDC correlates at 0.77, higher than the flexor FDS that leads the task at 0.61, and it does so with an involvement of 0.16 against 0.32. On this task the model reproduces the timing of the antagonist better than that of the agonist.

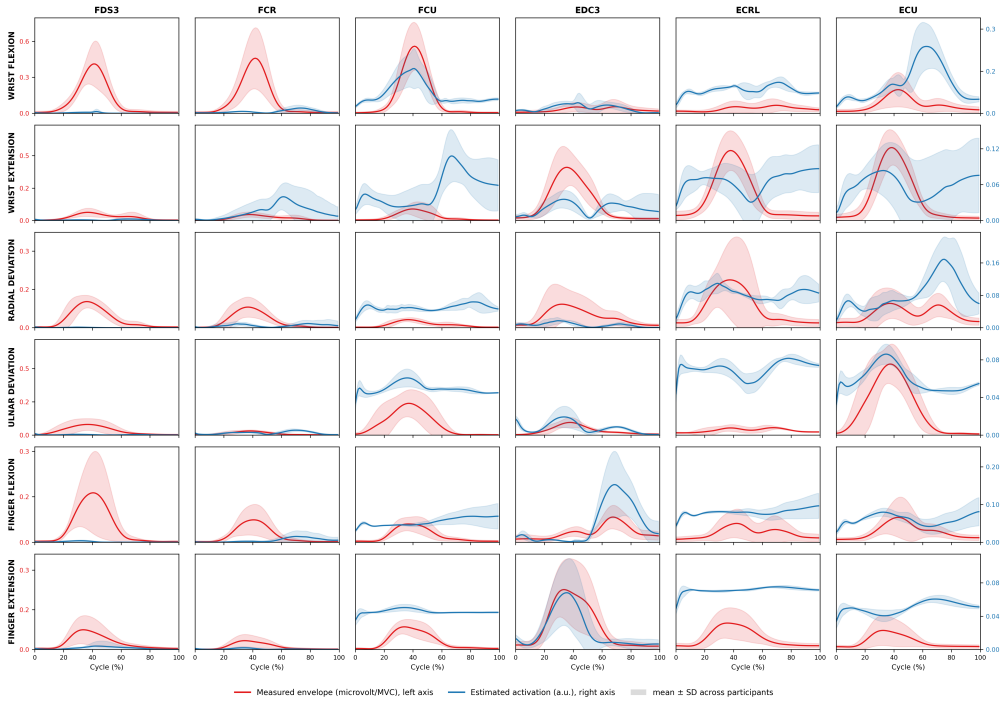


Figure 4.4: Estimated activations and measured sEMG envelopes, one row per task and one column per monitored muscle, mean and standard deviation across participants. The two scales are independent and common to the row: the envelope is read on the left axis and the activation on the right axis.

Fig. 4.4 shows the time course of estimated activations against envelopes, task by task and muscle by muscle. On several channels the estimated activation stays close to zero over the whole cycle, while the measured envelope carries a burst. The correlation summarises agreement in shape alone, since it is invariant to a multiplicative constant, and the figure already shows why: the two curves of a cell are drawn on scales that had to be set independently of one another. Whether they also agree in amplitude is the subject of Sec. 4.2.3.

4.2.3 Agreement in amplitude

Table 4.7 reports the coefficient of determination, which measures agreement in shape and amplitude together, and the gain, which measures the difference in amplitude alone.

Table 4.7: Coefficient of determination and gain (second line of each task); median across participants, IQR below. Cells are shaded by R^2, red where it is positive and blue where negative, saturating below -1.

Task	FDS	FCR	FCU	EDC	ECRL	ECU
Wrist flexion	−0.45	−0.47	0.26	−0.68	−11.27	−2.12
	(−0.60, −0.40)	(−0.63, −0.39)	(0.10, 0.33)	(−0.92, −0.41)	(−33.34, −1.75)	(−9.28, −1.07)
gain	101.91	19.39	3.18	1.28	0.64	0.72
	(96.80, 123.62)	(9.62, 42.49)	(2.79, 4.72)	(0.91, 2.65)	(0.40, 0.97)	(0.41, 0.83)
Wrist extension	−0.94	−0.40	−2.30	−0.55	−0.47	−0.36
	(−1.16, −0.71)	(−0.88, −0.06)	(−6.38, −1.09)	(−0.60, −0.43)	(−0.52, −0.38)	(−0.43, −0.28)
gain	31.47	1.43	0.99	10.81	6.30	7.78
	(17.43, 33.04)	(0.99, 2.11)	(0.56, 1.39)	(8.71, 12.18)	(4.70, 6.91)	(6.81, 8.29)
Radial deviation	−0.63	−0.41	−12.34	−0.91	−0.33	−1.53
	(−0.76, −0.58)	(−0.58, −0.39)	(−13.22, −11.01)	(−1.23, −0.41)	(−1.54, −0.25)	(−8.42, −1.04)
gain	46.78	4.62	0.53	3.54	1.25	0.59
	(29.45, 58.98)	(3.19, 6.56)	(0.46, 0.58)	(3.31, 4.66)	(1.05, 2.00)	(0.42, 0.94)
Ulnar deviation	−0.77	−0.46	−0.08	−0.46	−5.33	−0.18
	(−0.97, −0.71)	(−0.67, −0.31)	(−0.36, 0.11)	(−0.98, −0.18)	(−10.50, −0.95)	(−0.42, −0.04)
gain	47.97	5.55	3.51	3.86	0.91	7.58
	(37.57, 65.31)	(4.52, 6.45)	(2.82, 5.05)	(2.58, 5.38)	(0.58, 1.25)	(5.76, 8.98)
Finger flexion	−0.61	−0.57	−1.77	−3.64	−16.52	−0.68
	(−0.71, −0.53)	(−0.74, −0.53)	(−5.35, −1.07)	(−5.56, −0.86)	(−49.49, −3.30)	(−15.84, −0.22)
gain	42.00	5.46	1.17	0.49	0.63	0.98
	(22.59, 69.78)	(4.12, 13.22)	(0.88, 1.43)	(0.41, 0.89)	(0.37, 1.24)	(0.60, 1.33)
Finger extension	−0.85	−0.60	−0.15	−0.12	−2.68	−6.75
	(−1.15, −0.64)	(−0.67, −0.55)	(−0.36, −0.04)	(−0.47, 0.09)	(−3.93, −0.67)	(−37.60, −0.68)
gain	20.17	14.77	1.85	3.53	1.05	0.84
	(11.20, 31.63)	(9.87, 25.13)	(1.61, 1.97)	(3.10, 6.58)	(0.92, 2.14)	(0.40, 2.00)

The coefficient tells a different story from the correlation. It is negative in thirty-five of the thirty-six task–muscle pairs, the single exception being FCU in wrist flexion at 0.26; even the pairs that correlate above 0.8 have coefficients of −0.18 for ECU in ulnar deviation and −0.12 for EDC in middle-finger extension. Agreement in shape, where it exists, is not accompanied by agreement in amplitude.

The gain exceeds unity in most of the pairs, and by a wide margin on the muscles that lead their task: 3.18 for FCU in wrist flexion, 7.58 for ECU in ulnar deviation, 3.53 for EDC in finger extension, 10.81 for EDC in wrist extension. On those muscles the peak of the estimated activation is therefore three to ten times lower than the peak of the measured envelope. FDS is the extreme case: its gain is 101.91 in wrist flexion and 42.00 in the middle-finger flexion it leads.

On some of the muscles the task does not recruit the ratio goes the other way. ECRL and ECU in wrist flexion have gains of 0.64 and 0.72, and EDC in finger flexion of 0.49: there the model assigns more activity, relative to the peak, than the

envelopes show.

A gain close to unity does not by itself indicate agreement. FCU in wrist extension has a gain of 0.99 and a coefficient of −2.30, ECRL in ulnar deviation 0.91 and −5.33, ECRL in finger extension 1.05 and −2.68: the two signals reach comparable peaks and still disagree over the cycle.

Read by muscle rather than by task, the gain separates the channels. FDS is above 20 in every one of the six tasks, from 20.17 to 101.91, whereas ECRL stays close to unity in three of them, at 0.91, 1.05 and 1.25.

This section answers the first two research questions, within the limits set by the reproduction of the imposed movement. Activations of the forearm and hand muscles can be estimated from markerless kinematics alone, through the inverse dynamics of the model and without recourse to any measurement of muscle activity, to the extent that the model reaches the configuration it is asked to reach: this holds at the wrist, holds on the fingers except at their distal joint, and does not hold in fist closure, which is withdrawn from the comparison (RQ1). Those estimates agree with the measured envelopes in shape, for the muscle that leads the task, in four tasks out of six, and do not agree in amplitude in any of them. The agreement is therefore partial and depends on the task and on the muscle: an estimated activation may agree with the measured envelope in timing, but not in amplitude (RQ2).

4.3 EMG-informed MjID: a feasibility test

Sec. 4.2 established that the estimated activations follow the measured envelopes in shape on most of the muscles that lead a task, and not in amplitude. This section asks whether two requirements can be satisfied at once: whether the solver can be drawn towards the recorded activity while still reproducing the movement, and what the first costs in terms of the second.

All quantities use the definitions and the aggregation of Sec. 3.2.4 and Sec. 3.2.5, computed on the individual repetition and summarised by median and IQR across participants.

4.3.1 Comparison of the configurations

Fig. 4.5 compares the five configurations with the unconstrained solution of Sec. 4.2, which serves here as the baseline. The Friedman test rejects equality on all four quantities, from $p < 10^{-5}$ on the correlation with the reference to $p = 5 \cdot 10^{-4}$ on the coefficient of determination of the driven DOFs.

The five are named after the two axes of Sec. 3.3.5. The prefix gives the amplitude of the reference: *avg* for uniform anchoring, *MVC* for the normalised envelope used directly. The suffix gives the weight of the thirty-three unrecorded muscles: *floor100* for a single weight shared with the recorded ones, *low5* and *low30* for a weight fixed at those two values, *free* for a weight left to the fit. The weight of the recorded muscles is fitted above the same lower bound in all five.

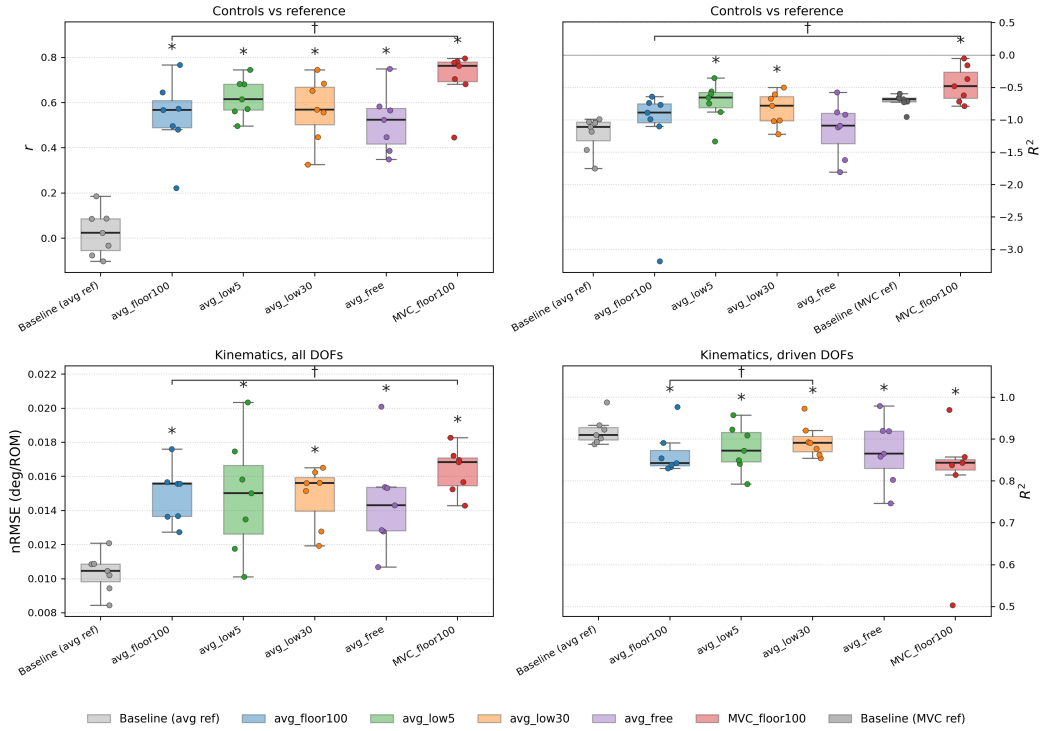


Figure 4.5: Agreement with the reference and cost in tracking accuracy, one point per participant, pooled over muscles, joints and tasks. Asterisk: differs from the corresponding baseline. Dagger: differs from *avg_floor100*; both $p < 0.05$.

The first result concerns the baseline itself. Its controls were produced without any reference, and their correlation with the reference is 0.02: the unconstrained solution does not resemble the recorded activity to any useful degree. Whatever agreement the other five configurations show is therefore produced by the term that was added, and not inherited from the estimation that preceded it. The correlations of this section are not those of Sec. 4.2 aggregated. There the activations the model produces were placed beside the measured envelopes; here the controls the solver produced are placed beside the reference built from those envelopes, and the two pairs of quantities are separated by the activation dynamics of Eq. 3.4. The two sets of values are not comparable and are not meant to be.

All five configurations depart from the baseline on the correlation, on the normalised tracking error and on the coefficient of determination of the driven DOFs. The coefficient of determination between the controls and the reference behaves otherwise: only three configurations depart from the baseline (*avg_low5*, *avg_low30* and *MVC_floor100*) while *avg_floor100* and *avg_free* do not. Agreement in amplitude is thus harder to obtain than agreement in shape even when the reference is supplied explicitly, which extends to this section the dissociation reported in Sec. 4.2.3.

Both quantities are computed against the reference of the configuration being examined, as Sec. 3.3.5 prescribes, and the baseline has none of its own: it is therefore evaluated twice, once against each of the two references in use, and each configuration is read against the value that corresponds to it. On the correlation the two coincide,

since r is invariant to the scale factor that distinguishes the two references, and only one is drawn. On the coefficient of determination they do not, the baseline reaching -1.11 against the anchored reference and -0.68 against the normalised envelope: the anchoring factor is itself estimated from the model, participant by participant, and carries its own variability into the metric.

The two axes along which Sec. 3.3.5 varies the five configurations behave differently. Along the weight given to the unrecorded muscles, the three configurations that vary it are all but indistinguishable from *avg_floor100*: of the twelve comparisons only one separates, *avg_low30* reaching a higher coefficient of determination on the driven DOFs, 0.891 against 0.843. Along the amplitude of the reference, *MVC_floor100* differs from *avg_floor100* on the correlation, on the coefficient of determination with the reference and on the tracking error, though not on the coefficient of determination of the driven DOFs. Its correlation with the reference reaches 0.76 against 0.57, and its normalised tracking error 1.68% of the anatomical range against 1.56%. The axis that the design varies by a single scale factor is therefore the one that matters.

The configuration carried forward is *MVC_floor100*. Uniform anchoring derives the amplitude of the reference from the controls the model itself produced at the previous stage, so the signal towards which the model is subsequently drawn is built in part from the model; the normalised envelope does not, its amplitude coming from the reference amplitude of Sec. 3.1.4. That the same configuration also binds most closely is reported here, but it is not what decides the choice, nor is it the best on every count, since it is also the configuration with the highest normalised tracking error of the five, 1.68% against 1.43–1.56%, and one of the lowest coefficients of determination on the driven DOFs, 0.844 against 0.843–0.891. It is chosen because its reference does not depend on the model, and it is carried forward with that cost attached.

4.3.2 Agreement with the control reference

Fig. 4.6 reports the agreement, in terms of Pearson correlation, between the controls the solver produced for the adopted *MVC_floor100* configuration and the reference it was given, resolved by muscle and by task. Both conditions are evaluated against the same reference, so the quantity compared is the same on the two sides of each pair; what separates them is the formulation of Sec. 3.3.3 taken as a whole.

The soft constraint moves the correlation upward in thirty-five of the thirty-six task–muscle pairs. The solver does what it was asked, and does it on most of the muscles of most of the tasks.

Where the shift is large it is also systematic. Wrist extension is the clearest case: all six channels improve in every participant, ECRL from -0.65 to 0.95, ECU from -0.47 to 0.98, FCU from -0.61 to 0.85, and all six survive the correction. Wrist flexion contributes two further significant pairs, FCR rising from -0.23 to 0.84 and ECU from 0.40 to 0.85. Ulnar deviation shows shifts of comparable size, FCR from -0.10 to 0.86 and ECU from 0.59 to 0.99, and carries no mark at all: three of the

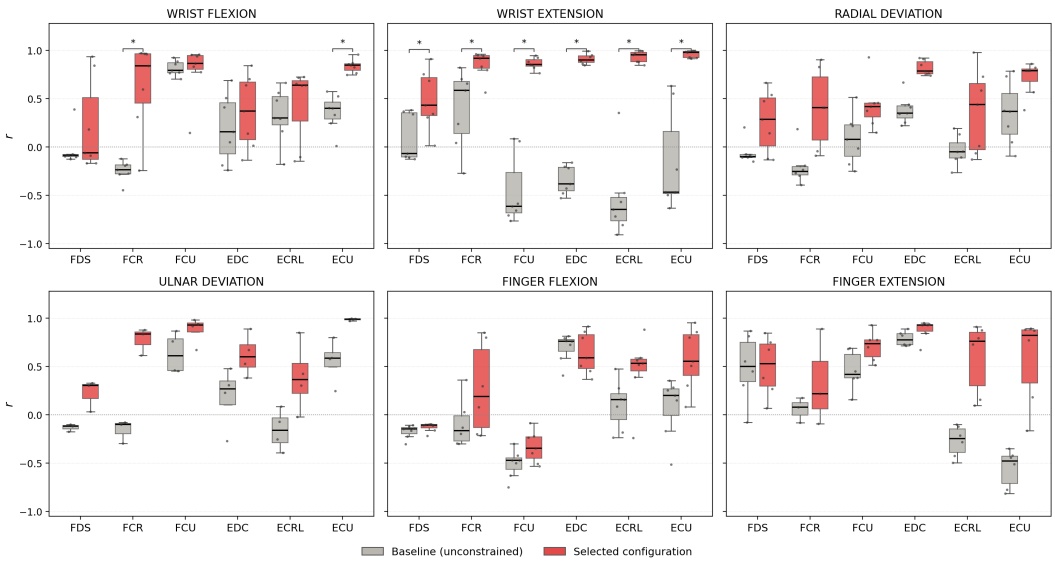


Figure 4.6: Correlation between the controls produced by the solver and the reference it was given, one panel per task and one pair of boxes per muscle: unconstrained baseline on the left, selected configuration on the right; one point per participant. Asterisk: the two conditions differ, with $p < 0.05$.

seven participant–task combinations of that task were excluded before the model (Sec. 3.1.5), leaving four, and with four paired observations the smallest value the two-sided exact test can return is 0.125.

The constraint does not take hold everywhere. FDS in middle-finger flexion moves from −0.14 to −0.11 and FCU from −0.47 to −0.34: on that task the constraint takes hold on the extensors and not on the flexor that leads it, and EDC is the single pair of the whole figure that moves downward, from 0.76 to 0.59; also FDS in wrist flexion moves by 0.03. FDS is in fact the least well bound channel overall, staying below 0.55 in every task. The solver can be brought to follow a recorded signal on the wrist musculature and, on the finger flexor, largely cannot.

The coefficient of determination behaves otherwise, and Fig. 4.5 reports it pooled rather than by muscle because the effect is not resolved at that level. Its improvement is small compared with that of the correlation, and the reading of Sec. 4.2.3 carries over: the solver reproduces the shape of the signal it is given and not its scale, even when that signal is placed in the objective.

A pattern that will matter in Sec. 4.3.4 is already visible: in wrist flexion the constraint raises FCR to 0.84 and ECU to 0.85, values comparable with FCU at 0.86, on muscles whose contribution to that movement is secondary. The reference is followed on the channels that are given it, whether or not the movement requires them.

The figure covers the six recorded channels, the only ones a reference exists for. What the solver does with the remaining thirty-three, which receive none, is shown task by task in Appendix B, Fig. B.3 onwards, where the control of every actuator of the model is drawn in both conditions.

4.3.3 Cost of the agreement

Fig. 4.7 reports the RMSE between the configuration the model reaches and the trajectory it was asked to follow, for the selected configuration and for the unconstrained solution, with the difference between the two boxes of each pair being what the agreement of Sec. 4.3.2 costs.

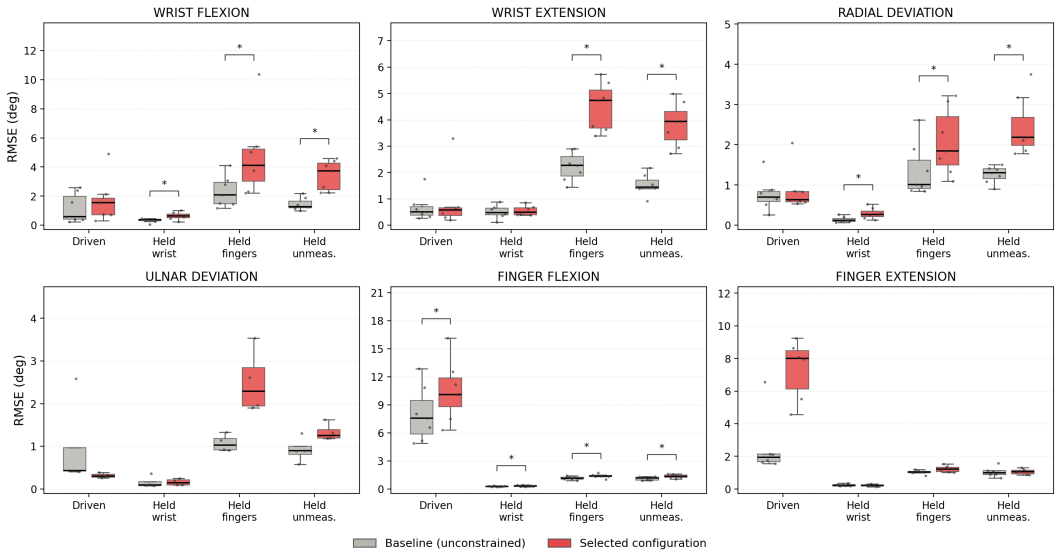


Figure 4.7: RMSE between the configuration reached by the model and the imposed target, by group of DOFs and by task, before and after the constraint; one point per participant, each panel has its own vertical axis. Asterisk: the two conditions differ, with $p < 0.05$.

On the driven DOFs of the four wrist tasks nothing happens. The error stays below 1.6° and in two of the four tasks the EMG-informed solution tracks better than the baseline, and in none of the four does the comparison approach significance: the movement the model was asked to reproduce is reproduced as before. What degrades in those tasks are the joints that were merely to be held: the finger and thumb joints and the six unmeasured DOFs roughly double in wrist flexion, wrist extension and radial deviation, in every participant and with the difference surviving the correction, while the held wrist DOF grows by a quarter of a degree at most. The model gains its agreement with the recording out of the posture, not out of the movement.

The finger tasks behave otherwise, and there the movement itself is spent. The driven error rises from 7.58° to 10.08° in flexion, in all seven participants, and from 1.94° to 8.01° in extension, in all six; the second shift is the larger of the two in relative terms, but with six paired participants and the correction applied over the four groups of the task it does not clear the threshold, so it should be read from its size and its consistency instead of being read only from the significance mark.

Fig. 4.8 resolves those two figures over the three driven DOFs, and the cost turns out to be ordered exactly as the reconstruction is. In middle-finger flexion the MCP angle, the one that is measured, degrades from 4.03° to 5.33° with the coefficient of determination falling from 0.894 to 0.817; the PIP from 6.59° to 8.77° and from 0.812

to 0.548; the DIP from 11.32° to 17.60° and from −0.432 to −1.183. Middle-finger extension shows the same ordering more sharply: the MCP loses little, from 1.18° to 1.90° and from 0.895 to 0.851, whereas the PIP goes from 0.84° to 6.49° and from 0.940 to −0.169, and the DIP from 3.41° to 14.99° and from −0.023 to −11.360. Every one of the twelve comparisons is significant, both on the error and on the coefficient.

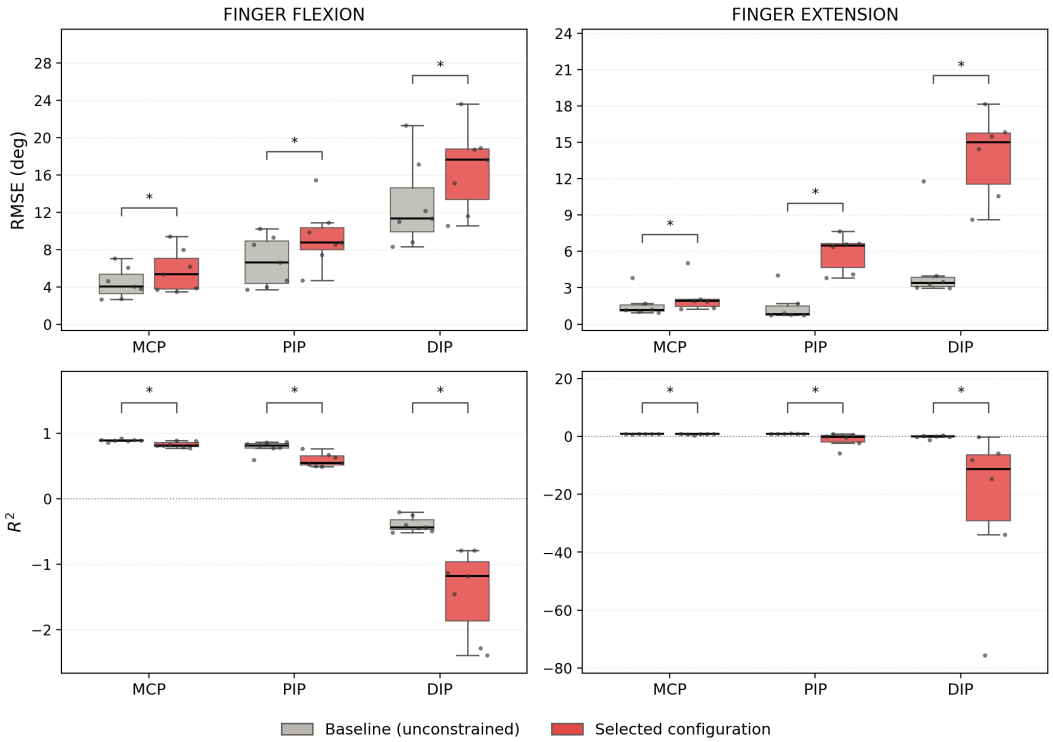


Figure 4.8: The three driven DOFs of the finger tasks taken separately: RMSE above, coefficient of determination below; each panel has its own vertical axis. Asterisk: the two conditions differ, with $p < 0.05$.

On the DIP the error is no longer small with respect to the movement: the imposed excursion of that joint is 27.1° in flexion and 9.98° in extension, against errors of 17.6° and 15.0°, so in extension the residual exceeds the whole displacement the joint was asked to make. The coefficients below −1 are not an artefact of a small denominator alone. The dispersion adds something the medians do not say. On the MCP the constraint leaves the participants where they were, tightly grouped above 0.75 in both conditions; on the PIP and the DIP of middle-finger extension it does not only lower the coefficient but also scatters it with two participants falling below −30 on the DIP and one below −75. The constrained solution on those joints is not uniformly worse, but rather, it is unstable from one participant to the next.

Sec. 4.2.1 had already found the DIP to be the weakest link of the unconstrained model, with a coefficient of −0.432 in flexion before any constraint was applied. The constraint does not create that weakness; it concentrates its own cost on the joint that had the least margin to absorb it.

Read together, Sec. 4.3.2 and this subsection outline the shape of the trade-off

that Sec. 3.3.5 set out to establish, and answer to what cost to the tracking of the measured kinematics the recorded activity can be introduced (RQ3). Agreement with the recording can be obtained, and at the wrist it is obtained almost for free, the price being paid by joints that were only being held still; on the finger it is obtained at the price of the movement itself, and precisely on the reconstructed DOFs driving the movement.

4.3.4 Robustness check: the LOMO test

The experiment of Sec. 3.3.6 runs the selected configuration six times per movement, withholding one recorded muscle at a time and keeping the parameters fitted for that configuration, without any re-optimisation. Over the thirty-eight participant–task combinations that enter the feasibility experiment this amounts to almost four thousand simulated repetitions, evaluated with the same quantities and the same aggregation used for the feasibility comparison in Sec. 3.3.5.

The weight of Sec. 3.3.3 acts on every actuator, so a muscle removed from the reference is not left free, but drawn towards zero as the muscles with no reference are. The unconstrained baseline remains the right floor to compare against, since it is what the model produces for that muscle with no recording at all, but the comparison is conservative.

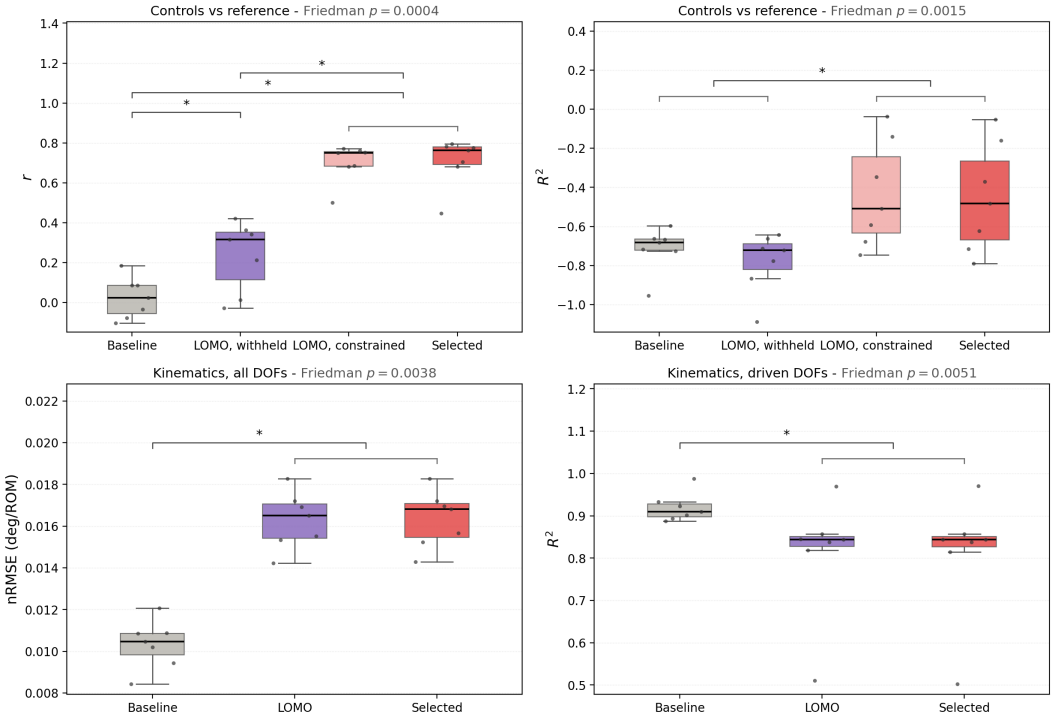


Figure 4.9: The LOMO test. Agreement with the reference and cost in tracking accuracy, one point per participant, pooled over muscles, joints and tasks. Asterisk: the two conditions differ, with $p < 0.05$.

Fig. 4.9 reduces every participant to a single value per condition, pooling the six withheld muscles and the six tasks together. The Friedman test rejects equality

on all four quantities, from $p = 4 \cdot 10^{-4}$ on the correlation with the reference to $p = 4 \cdot 10^{-3}$ and $p = 5 \cdot 10^{-3}$ on the two kinematic ones.

The first result is that the experiment costs nothing kinematically. The normalised tracking error of the LOMO runs is 1.65% of the anatomical range against 1.68% for the selected configuration, and the coefficient of determination on the driven DOFs is 0.844 in both. The two conditions are indistinguishable from one another and both differ from the unconstrained baseline, by the same amount and in the same direction that Sec. 4.3.3 reported for the selected configuration.

The second result is that removing one constraint does not disturb the other five. Their correlation with the reference is 0.75, against 0.76 when all six are constrained, and their coefficient of determination -0.51 against -0.48; neither comparison approaches significance. The five muscles that kept their reference behave as if the sixth had never been removed.

The third result is the one the experiment was built for, and it falls between the two levels. The withheld muscle reaches a correlation of 0.32 with the recording it was never shown, against 0.02 for the unconstrained baseline and 0.76 when it is constrained by its own reference; it differs from both, and after correction both differences survive. The constraint on the other five therefore carries some information about the sixth: the solution is not the one the model would have produced without any recording. But the recovery is much closer to the floor than to the ceiling, so most of what a muscle's own constraint contributes cannot be supplied by the other five.

On amplitude there is no recovery at all. The coefficient of determination of the withheld muscle is -0.72, indistinguishable from the -0.68 of the unconstrained baseline and clearly below the -0.48 of the selected configuration. The dissociation of Sec. 4.2.3 thus holds in this experiment as well, and in its sharpest form: whatever is transferred from one muscle to another concerns the timing and nothing of the scale.

Fig. 4.10 then separates the six muscles and reports the withheld one alone; the task is shown as the colour of the marker, so that its contribution remains visible, but no comparison is made between tasks.

The figure shows that the median of 0.32 hides two opposite behaviours. The wrist extensors recover well (ECU at 0.43, ECRL at 0.41, EDC at 0.33) while the flexors do not: FCR at -0.09, FDS at -0.06, FCU at 0.08. The same split appears when the recovery is read against how much of the recorded activity the muscle carries. For FDS, FCR and FCU the two move in opposite directions, with $\rho = -0.47$, -0.63 and -0.54: in the movements where those muscles are marginal the model reproduces them at 0.75 to 0.90, and in the movements they lead it fails, dropping below zero. For ECRL and ECU there is no such dependence, $\rho = -0.03$ and -0.08, and the recovery stays between 0.3 and 0.7 whatever the involvement, while EDC sits between the two groups, $\rho = -0.30$.

These two levels together address the question about how much of the information carried by one recorded channel is redundant with the other five (RQ3). Pooled over

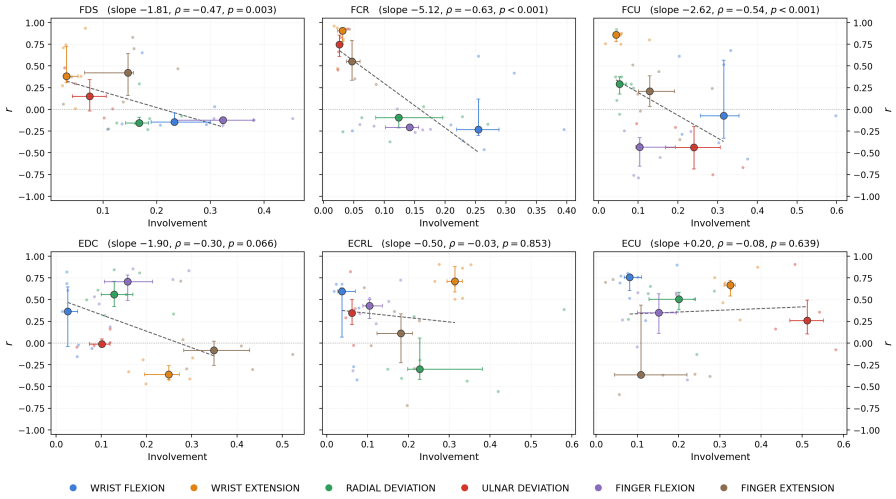


Figure 4.10: Correlation between the activity assigned to the withheld muscle and the recording it was never shown, against the involvement of that muscle in the movement, one panel per muscle. Small markers are single participants in a single task, the large marker their median and the bars its IQR, on both axes.

the six, some of it is and most of it is not. Resolved by muscle, the redundancy is not distributed evenly: on the wrist extensors it is substantial and it does not depend on the role the muscle plays, while on the flexors it is confined to the movements in which the muscle contributes little, and disappears exactly where the muscle matters, which is the condition under which the recovery would have been useful. EDC lies between the two, recovering at the level of the extensors, but with the dependence on involvement of the flexors. What one channel can be reconstructed from the others is therefore a property of the muscle and of its part in the movement, and not of the method.

4.4 Mechanical characterisation of the model

The two previous sections judged the model against a recording, and could do so only on the six muscles that were recorded. They say nothing about the thirty-three that were not, which the model drives all the same and for which no measurement exists. This section describes what the model makes all thirty-nine muscles do, and builds towards the question Sec. 3.4 sets out: how much of that mechanics the six electrodes account for.

The quantity examined throughout this section is the joint torque of Eq. 3.15, read from the simulation at every integration sub-step rather than reconstructed afterwards from the frames, so the muscle force is the one the model actually produced. Two conditions are compared everywhere, as in Sec. 4.3: the unconstrained solution, which is the estimate validated in Sec. 4.2, and the selected EMG-informed configuration.

Values are means across participants with their standard deviation. The convention differs from that of the two previous sections, and the reason is the one stated in Sec. 3.2.5: the quantities reported there were correlations, coefficients of determi-

nation and gains, which are unbounded or asymmetric by construction, whereas the quantities reported here are shares bounded between zero and one hundred whose values lie far from both bounds.

4.4.1 Anatomical capacity

Fig. 4.11 reports what the muscles of the model are able to do, before asking how the model employed them in reproducing the movement of each task, in terms of torques generated by the muscles. Each row is normalised to the total capacity of that muscle, as Sec. 3.4.3 prescribes, so the figure reads as the distribution of a muscle's mechanical action across the hand. The columns are the DOFs of the model, grouped where a group is more informative than its parts. The two DOFs of the wrist are kept apart, since a muscle can act on flexion and extension without acting on deviation, and forearm rotation is a third DOF of its own; each of the four fingers and the thumb, on the other hand, carries four DOFs. The eight columns therefore account for the twenty-three DOFs of the model.

The capacity is an average of the moment arms over sampled postures, so in principle it could depend on which postures are sampled. Empirically, it does not: the map computed over postures drawn at random from the joint ranges and the map weighted by the postures actually visited in the recorded trajectories correlate at $r = 0.99$. The distribution reported here is therefore a property of the model itself, and not of the particular configurations in which it was operated.

Reading the rows, the six recorded muscles do not all behave in the same way. The four wrist muscles are almost purely wrist muscles: ECU places 69% of its capacity on deviation and FCR 66% on flexion and extension, and neither reaches any finger. The two finger channels are not the mirror image of that. FDS places 47% of its capacity on the middle finger but 34% on wrist flexion and extension, and EDC 44% on the middle finger against 42% at the wrist: the two channels used to observe the finger act on the wrist about as much as they act on the finger. The nine intrinsic muscles of the fingers and the four of the thumb, by contrast, are entirely confined to their own digit.

The muscles that act on the fingers and on the thumb are not confined to them: each compartment of FDS, flexor digitorum profundus (FDP) and EDC places about half of its capacity on the wrist, and between a quarter and two fifths of it on flexion and extension alone. Conversely, several of the rows that place all of their capacity on a single digit belong to muscles that no electrode observes.

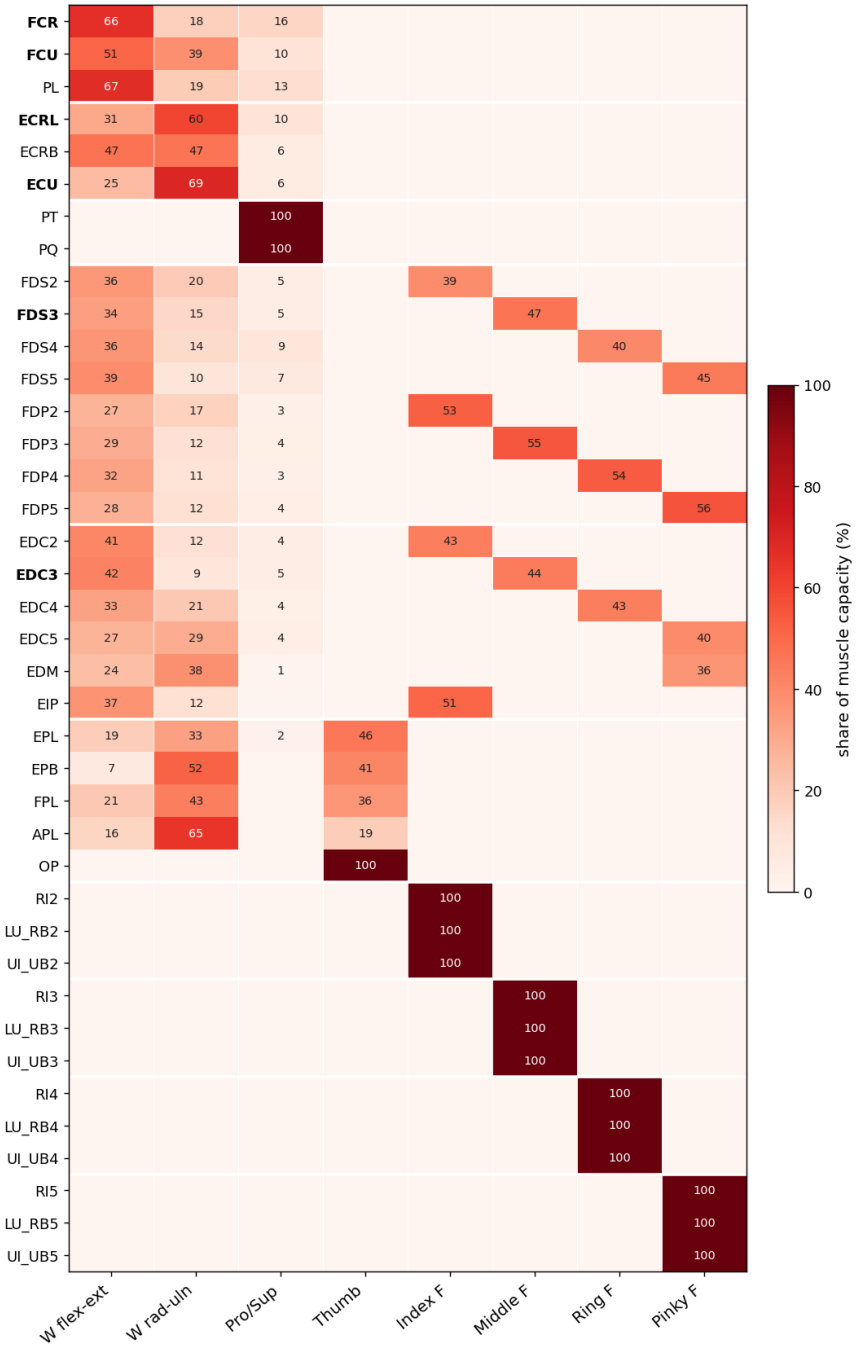


Figure 4.11: Capacity of each muscle on each region of the hand, as a percentage of that muscle's own total torque capacity, so that every row sums to one hundred. Bold: the six recorded muscles.

4.4.2 Movement and stabilisation

Fig. 4.12 separates the torque that moves the hand as the movement requires from the torque that the movement asks to hold the fingers and the wrist still. The first result is that the partition is essentially a property of the task and not of the solution: the two conditions give the same picture in all six tasks, and the only systematic difference between them is a shift from wrist to finger stabilisation in the wrist tasks. Every figure below is therefore given as a pair, the unconstrained solution first and

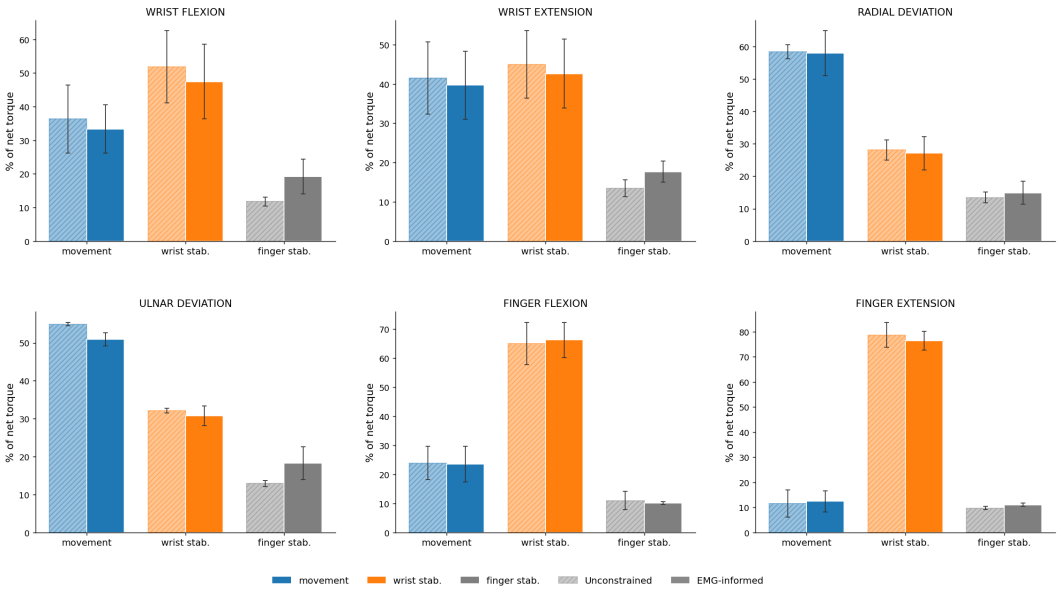


Figure 4.12: Net joint torque split into the DOFs the task moves, the remaining wrist DOFs and the remaining finger and thumb DOFs, as a percentage of the total. Bar: mean across participants; whisker: standard deviation.

the EMG-informed one second.

In the wrist tasks the movement takes between a third and three fifths of the total, 36.3%–33.4% in flexion and 58.4%–58.0% in radial deviation, and the wrist DOF that is not driven takes most of the remainder. In the finger tasks the proportion collapses: moving the middle finger takes 23.9%–23.6% of the torque in flexion and 11.5%–12.5% in extension, while holding the wrist takes 65.0%–66.3% and 78.7%–76.5%. Bending one finger costs the model far more in holding the wrist than in effectively bending the finger, which is a consequence of the geometry reported in Sec. 4.4.1: the muscles that flex the finger cross the wrist, so their action there has to be cancelled.

Supplying the recording changes this partition little, and in one direction. The share spent on the movement falls in five tasks out of six, and the share spent holding the fingers rises in all four wrist tasks, from 11.8% to 19.2% in flexion, from 13.5% to 17.7% in extension, from 13.4% to 14.9% in radial deviation and from 12.9% to 18.3% in ulnar deviation. No test is applied to these differences, here or anywhere else in this section, so what follows is an observation and not a result: the direction of the shift is the same one Sec. 4.3.3 measured on the tracking error, seen now on the torque that produces it, with the constrained solution spending more effort on the hand joints it was supposed to be keeping still.

The dispersion is large, and the shares are not as well separated as their means suggest. In wrist extension the participants range, in the EMG-informed configuration, from 28% to 52% on the movement and from 33% to 57% on the wrist stabilisation, so the two regions overlap almost completely; in radial deviation and in the two finger tasks they do not overlap at all. The partition is therefore a property of the task more than a stable property of the model.

4.4.3 Agonists, antagonists and the recorded muscles

The previous subsection asked where the torque goes, in terms of DOFs related to the movement itself and to the stabilisation of the hand. This one explicitly asks who produces it: how it divides between the muscles that drive the movement and those that oppose it, and how much of each side the six electrodes observe.

Fig. 4.13 shows the torque each group produces on the driven DOFs over the normalised cycle, in the two conditions side by side. At the wrist the picture is the same in all four tasks. The recorded agonists and the recorded antagonists are almost mirror images of each other and peak at the same instant, around 40% of the cycle, so the net torque is an order of magnitude smaller than either: in wrist extension the two reach $+2.37$ and -2.17 N·m while the net peaks at 0.38 N·m. Comparing the two columns, the amplitude is not the only thing that changes. In the unconstrained solution the curves are comparatively flat, and flattest in the two deviation tasks, while under the EMG constraint the same curves acquire a marked peak. The recording therefore does two things at once: it concentrates the torque in time, and it scales it up, the peak of the recorded agonists rising by a factor between one and a half and five depending on the task, with the recorded antagonists rising alongside them. What does not change is the net, which moves by less than 0.05 N·m in all four tasks, and the unrecorded contributions, which stay small and flat. This is the mechanical statement of the trade-off: the kinematics is prescribed, so the net torque is fixed by it, and everything the recording adds is added in equal and opposite amounts on the two sides of the joint.

In wrist flexion and in radial deviation the EMG-informed curves do not return to rest monotonically: after the main peak the torque rises again at about 75% of the cycle, on the agonist and on the antagonist together. In flexion it is a distinct second peak, 0.38 N·m on the agonist against a main peak of 1.62 and 0.56 against 1.01 on the antagonist; in radial deviation it is a peak on the antagonist (0.64 against 0.93) while on the agonist it appears as a plateau near 0.54 rather than as a peak. Its position corresponds to the return phase of the movement, when the joint is brought back to the resting configuration.

The finger rows read differently, and in three respects. The first is scale: the three finger DOFs carry between 0.03 and 0.12 N·m against the 1 to 2.4 N·m of the wrist, a factor of twenty to fifty, which is the same asymmetry Sec. 4.4.2 reported as a share. The second is that the unrecorded muscles are not spectators here. In middle-finger flexion, at the MCP and the PIP joints the unrecorded agonists are as large as the recorded ones, and in the unconstrained solution larger (0.061 against 0.015 N·m at the MCP): a substantial part of the torque that bends the finger is produced by muscles no electrode observes. The third concerns the DIP joint specifically: there the unrecorded agonists are identically zero and no recorded antagonist appears, so the joint is left with a single non-negligible curve, the recorded agonist, opposed only by a small unrecorded contribution below 0.02 N·m. The DIP is therefore the one joint whose torque the constraint fixes almost completely, and in middle-finger

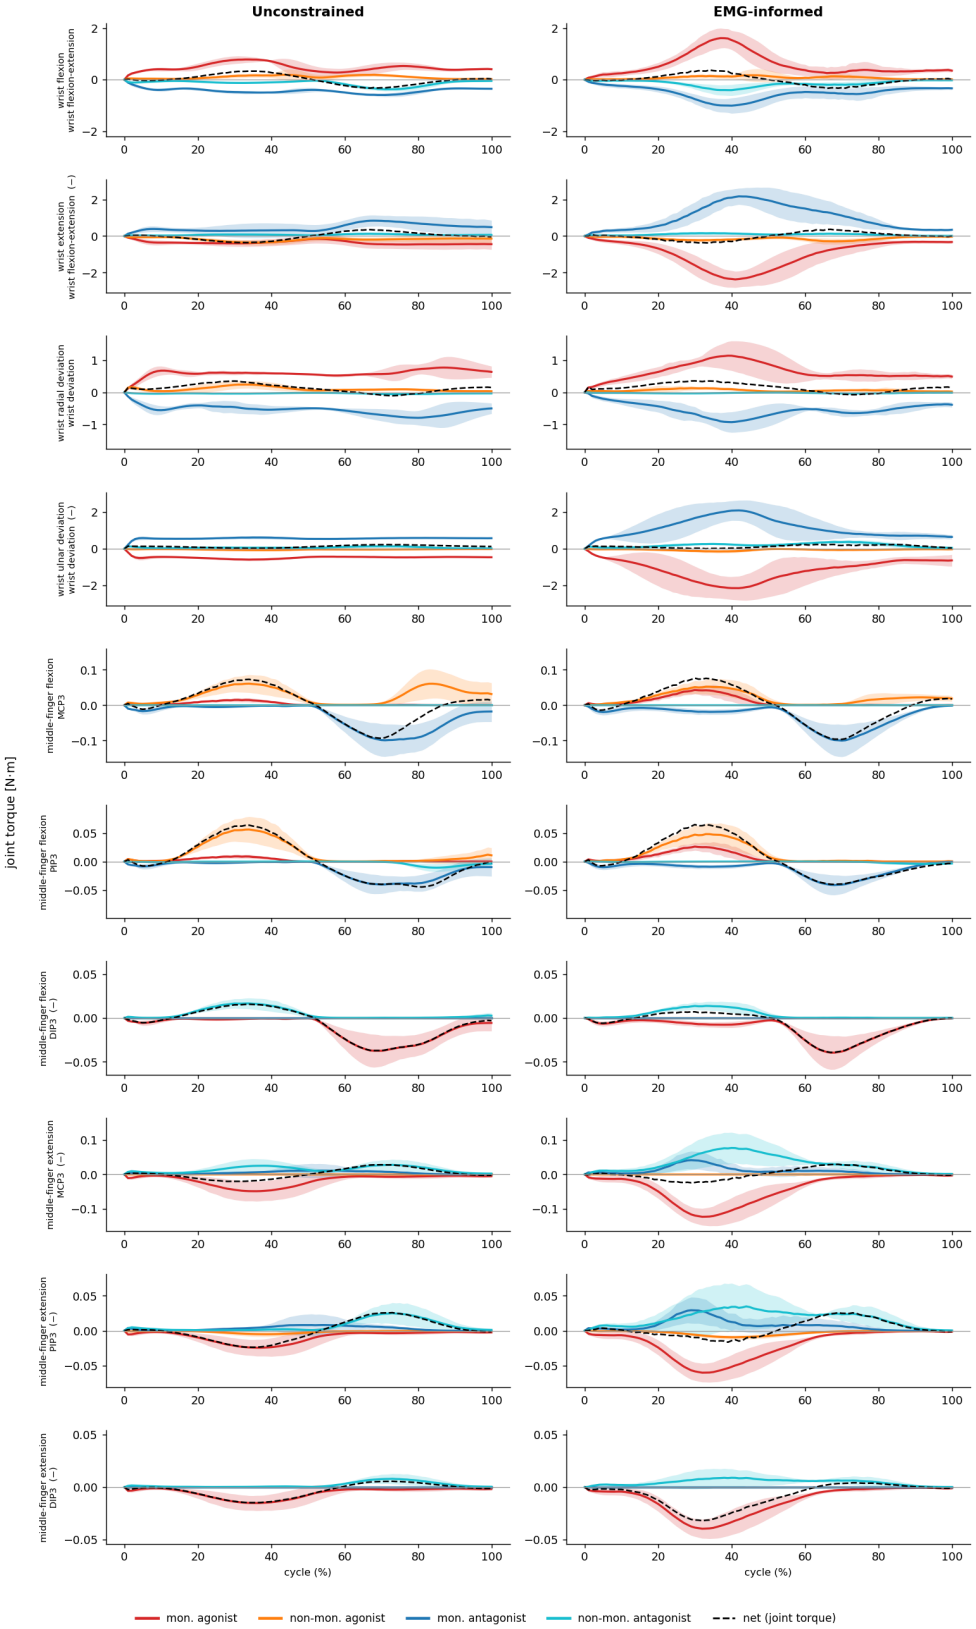


Figure 4.13: Torque on the driven DOFs over the normalised cycle, grouped into monitored and non monitored muscles and into agonists and antagonists. Mean and standard deviation across participants. Rows marked (−) are tasks whose movement direction is negative in the joint frame, so their agonist runs below zero.

extension it is also the only place where the net torque itself differs between the two conditions, doubling from 0.015 to 0.032 N·m. A net torque that changes means the two solutions are not producing the same movement, which is exactly the tracking failure Sec. 4.3.3 measured at that joint in that task, seen here on the quantity that causes it.

Three barplots take the three regions of Fig. 4.12 in turn and break each of them down muscle by muscle. Fig. 4.14 corresponds to its blue bars, the DOFs the task moves, with the muscles sorted inside the four categories. In the four wrist tasks the six recorded muscles carry between 77.6% and 90.6% of the torque without the recording and between 82.9% and 93.7% with it; in the two finger tasks they carry 48.1% and 49.3% without it, and 58.2% and 50.6% with it. Half of the torque that moves the middle finger is produced by muscles no electrode observes, and the figure names the most important of them: FDP3, which alone carries about a fifth, then extensor pollicis longus (EPL) at 11% to 14% and the ulnar interosseous of the third finger at 5% to 10%.

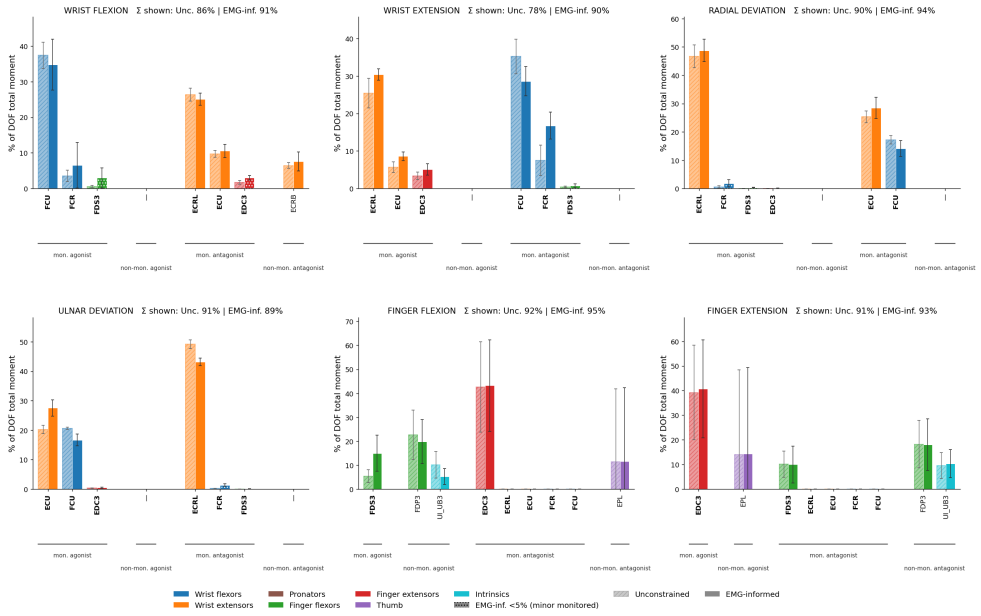


Figure 4.14: Share of the torque on the driven DOF carried by each muscle, as a percentage of the summed magnitude over the thirty-nine. Mean and standard deviation across participants. Bold: recorded muscles.

Fig. 4.15 does the same for the orange bars, the wrist DOFs the task holds still, and stabilisation divides more sharply than the movement did. The six recorded muscles carry between 82.3% and 92.6% of that torque without the recording and between 83.9% and 92.8% with it, in every one of the six tasks: what holds the wrist is the wrist musculature, and the electrode placement sees almost all of it. Three or four bars are enough to reach the 90% threshold, and they are the recorded wrist flexors and extensors.

Fig. 4.16 completes the set with the grey bars, the finger and thumb DOFs held still, and the picture is the opposite of the wrist one in two respects. The first is

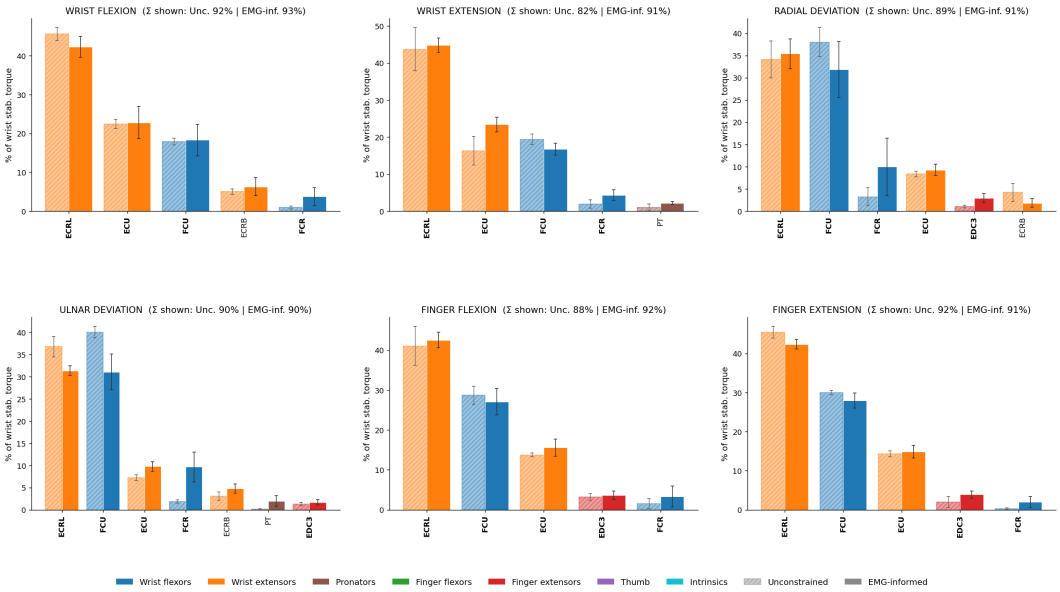


Figure 4.15: The same decomposition on the wrist DOFs the task holds still, ordered by the EMG-informed configuration and truncated where the cumulative share reaches 90%. Mean and standard deviation across participants.

that no muscle dominates: the largest single contributor reaches 10% to 18% of the region's torque without the recording and 11% to 22% with it, so the 90% threshold needs a dozen bars instead of three or four. These joints are held by many small contributions spread over the extensor compartment, the deep flexors and the thumb muscles, none of which is recorded. The second is that the coverage depends strongly on the task, and it is intrinsically limited here: of the six recorded muscles only EDC3 and FDS3 have any moment arm on the finger joints, the other four being wrist muscles. Without the recording the two carry 11.2% and 12.0% of this torque in wrist flexion and extension but only 2.3% and 2.2% in the two finger tasks, where the fingers are the joints of interest and the model holds them almost entirely with muscles the electrodes never see. Supplying the recording raises those shares to 30.0%–30.6% and 4.3%–4.1%, and does so by promoting EDC3, which without the recording holds between 6% and 9% of this torque in the four wrist tasks and with it between 11% and 22%, enough to make it the largest single contributor in all four: this is the same shift toward finger stabilisation that Sec. 4.4.2 measured on the partition, now identified mostly in a single actuator.

The direction of the comparison between the two conditions is consistent across the three regions: supplying the recording raises the share attributed to the recorded muscles in five of the six tasks on the driven DOFs, in four of six on the wrist stabilisation and in all six on the finger stabilisation. The model constrained by six channels attributes more of the mechanics to those same six channels, which is expected and is not evidence about the hand. Taken together, the three regions answer the last research question. The six recorded channels account for most of the mechanics of the wrist, whether the wrist is moving or being held, for about half of the torque that moves the middle finger, and for a fraction of the torque that holds

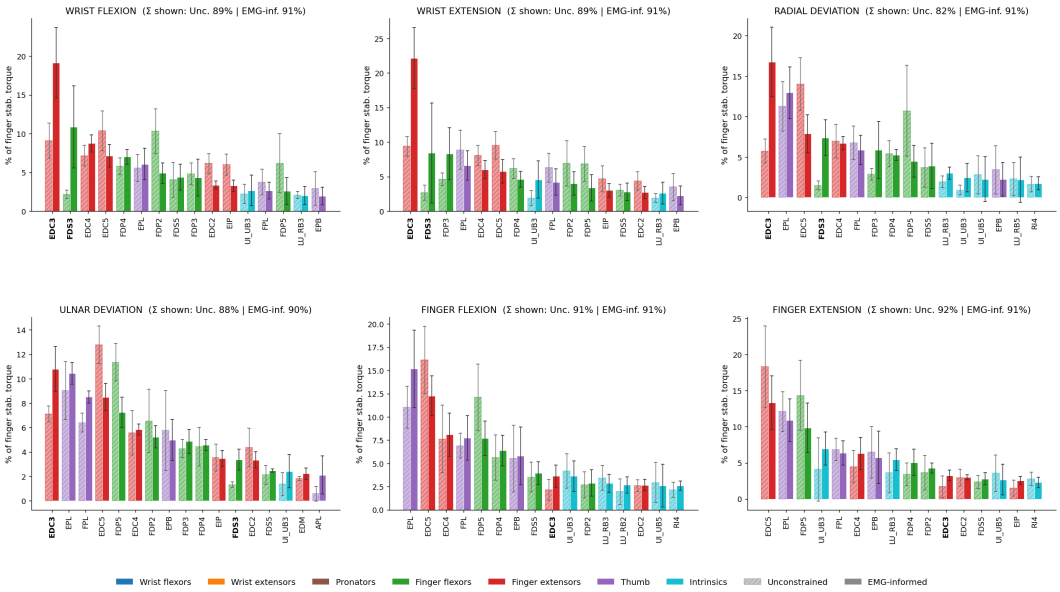


Figure 4.16: The same decomposition on the finger and thumb DOFs the task holds still, with the same ordering and the same 90% threshold. Mean and standard deviation across participants.

the fingers and the thumb still that ranges from a few per cent in the finger tasks to about a third in the wrist tasks: what the six electrodes leave unobserved is therefore small at the wrist and large at the fingers (RQ4).

4.4.4 Co-contraction

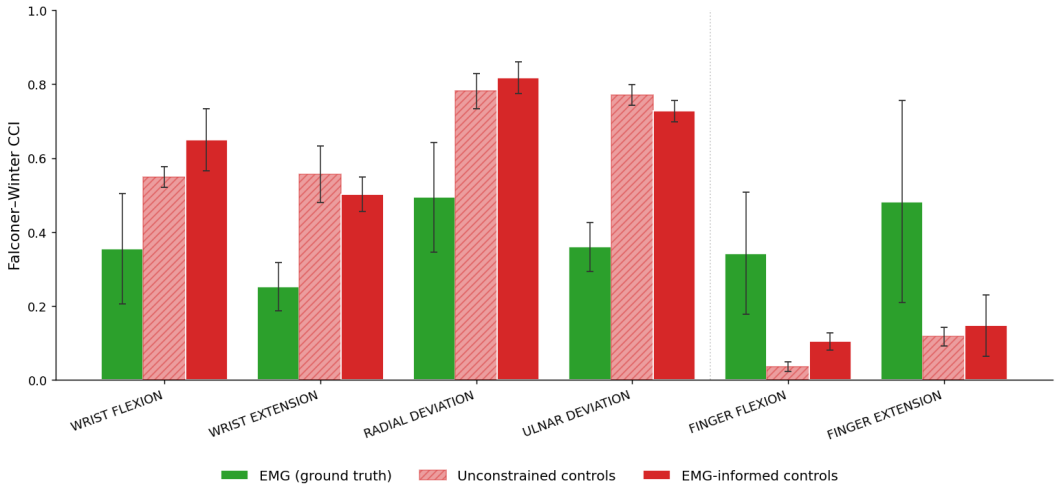


Figure 4.17: Falconer–Winter index on the six recorded muscles, for the measured envelopes, the unconstrained solution and the EMG-informed configuration. Bar: mean across participants; whisker: standard deviation.

Fig. 4.17 reports how much of the activity of the six recorded muscles is spent by agonists and antagonists acting against each other, measured by the co-contraction index of Eq. 3.16: zero when only one of the two groups is ever active, one when the two are equal throughout.

Table 4.8: Assignment of the six recorded muscles to the agonist and antagonist groups of Eq. 3.16. The assignment is made once, on the EMG-informed solution, and is the same for every participant.

Task	Agonists	Antagonists
Wrist flexion	FCR, FCU, FDS3	ECRL, ECU, EDC3
Wrist extension	ECRL, ECU, EDC3	FCR, FCU, FDS3
Radial deviation	FCR, FDS3, ECRL, EDC3	FCU, ECU
Ulnar deviation	FCU, ECU, EDC3	FCR, FDS3, ECRL
Middle-finger flexion	FDS3	EDC3
Middle-finger extension	EDC3	FDS3

The two groups are those of Table 4.8, assigned by the rule of Sec. 3.4.5 from the torque each muscle exerts on the primary DOF of the task. In the two finger tasks that DOF is the MCP joint of the middle finger, on which only FDS3 and EDC3 have a moment arm: the four wrist channels are left out of both groups, and the index there measures the balance of a single pair.

The assignment is made on the model and not on the recordings, but the two agree in the wrist tasks. In wrist flexion the agonists are the three flexors, which are also the three channels that lead the group averages of Sec. 4.1.3, and in wrist extension the same correspondence holds with the roles reversed; in ulnar deviation the agonists include ECU and FCU, the two channels that stand out there. Radial deviation is the exception, and in the expected direction: the model places four of the six channels on the agonist side, and the recording is the least selective of the set, its leading muscle standing out from the others by a factor of two only.

The index is given for the recording and for the two model solutions. The comparison that carries most information is the first pair, the recording against the unconstrained solution, and it fails in two opposite ways depending on the task.

In the four wrist tasks the model co-contracts about twice as much as the participants: the recorded envelopes give an index between 0.25 and 0.49, the unconstrained solution between 0.55 and 0.78. In the two finger tasks the model co-contracts far less: the recording gives 0.34 and 0.48 and the model 0.04 and 0.12. The discrepancy is therefore not a bias in one direction but a failure to reproduce the balance between agonists and antagonists at all, in opposite senses at the wrist and at the finger.

The dispersion adds a second failure to the first. On the recorded envelopes the participants differ widely, covering 0.09 to 0.83 in middle-finger extension and 0.09 to 0.53 in middle-finger flexion; the unconstrained solution, in the same tasks, stays within 0.08 to 0.15 and 0.02 to 0.06. It does not reproduce the level of co-contraction, and it does not reproduce the variability between people either: it returns nearly the same value for every participant where the measurement returns very different ones.

Supplying the recording moves the index very little, from 0.55 to 0.65 in wrist flexion, from 0.56 to 0.50 in extension, from 0.04 to 0.10 in finger flexion. It does not close the gap in either family. The third comparison, the EMG-informed solution against the recording, is not an independent one, for the reason given in Sec. 3.4.6:

those same envelopes were supplied to that solution as the reference for its controls, so any agreement would be in part built in, but it is surprisingly informative, since the index still remains far from the recorded value.

Chapter 5

Discussion

This chapter reads the results of Chapter 4 against the four research questions and against one another, building on three main notions. The first is that muscle activity can be estimated from markerless kinematics alone, and that the estimate follows the recording in time on the muscles that lead a task while never following it in amplitude. The second is that the recording can be introduced into the estimation as a soft constraint, at a kinematic cost that is negligible at the wrist and structural at the finger, and that the six channels carry information that cannot be recovered from one another where each of them matters most. The third is that both outcomes follow from what the six electrodes cover of the mechanics of the hand: nearly all of it at the wrist, about half of the movement at the finger, and almost none of what holds the fingers still. The chapter closes with the choices and limitations that bound these claims and with consequent implications.

5.1 Validation against sEMG: timing without amplitude

The signal quality assessment of Sec. 4.1 is designed to describe the recordings after preprocessing, and does not anticipate how the model will perform: what enters the model is not those traces but the segmented repetitions and, in the finger tasks, the reconstructed distal joints. The group averages have a more direct use, as they identify the muscle that leads each task while also giving a visual counterpart to the involvement of Eq. 3.13, which is the quantity against which the agreement of Sec. 4.2 has to be read. The correspondence is qualitative only, since the involvement is computed for each participant and pooled afterwards.

Activations of the forearm and hand muscles can be estimated from markerless kinematics alone, through the inverse dynamics of the model and without any measurement of muscle activity (RQ1). The method is not new, having been introduced and demonstrated on this model by Prestia [48]; what this work establishes is that it runs on the kinematics of real participants extracted from video. The answer holds only where the imposed trajectories are closely reproduced, which is the case at the wrist and on the fingers except at the distal joint, and is not the case in fist closure, withdrawn from the comparison for that reason.

How accurate those activations are is a separate question, and the answer is partial (RQ2). Agreement exists in terms of temporal correlation and on a subset of the channels: the muscle that carry most of the recorded activity, as indicated by the muscle involvement, are reproduced in four tasks out of six, as FCU and FDS in wrist flexion, ECU and FCU in ulnar deviation, EDC in finger extension and FDS in finger flexion. In amplitude there is no agreement in any muscle of any task: the coefficient of determination is negative in thirty-five of the thirty-six pairs, and the gain on the leading muscles goes from about three to about ten, reaching two orders of magnitude on FDS.

The two tasks where the agreement in time fails are not alike. Radial deviation is the weaker rather than the failing one and Sec. 4.1.3 shows why: there the leading muscle stands out from the others by a factor of two, the activity is spread across the hand, and it is one of the two tasks in which no pair of muscles survives the post-hoc comparison. Where the measurement itself does not separate the muscles, there is little reason for the model to. Wrist extension is a different case and remains open. Its envelopes are among the largest recorded, ECU and ECRL are among the muscles the task involves most, and the trajectory is reproduced as accurately as in the other wrist tasks, yet the two correlate weakly or negatively with the activations estimated for them, so low amplitude, marginal involvement and poor tracking are therefore all excluded as explanations.

The disagreement in amplitude may have two explanations, the first one being the scale. The envelope is normalised to the maximal voluntary contraction of that participant and that muscle, whereas the activation is a fraction of the maximal isometric force of the corresponding actuator, which comes from the parameters of the model and is not made participant specific, since the optimisation of Sec. 3.1.10 fits gains and weights and not muscle strength. The two signals are therefore bounded above by different quantities, the maximal isometric force of the actuator on one side and the maximal voluntary contraction of the participant on the other, and they do not share a range.

Scale alone may not account for what is observed, and a second explanation concerns the inverse dynamics algorithm itself. The imposed kinematics fixes the net torque at each DOF and not the individual muscle forces that produce it, and the split among those muscles is settled by a cost that asks for low and smooth activity, so what sets the level of each muscle is a criterion, resting on a minimisation assumption. Consistently with this, it is possible to observe that the model distributes the activity more evenly than the participants do, under-activating the muscles that drive the movement and over-activating the muscles the movement does not directly require. The gain is well above unity on the muscles that lead their task, as for FCU in wrist flexion and EDC in wrist extension, and it falls below unity on the muscles the task does not recruit, as for ECRL and ECU in wrist flexion and EDC in finger flexion.

Taken together, the two comparisons answer RQ2 in a single sentence: the shape and timing of the recorded activity is reproduced, but only on some muscles and in some tasks, and its amplitude is not reproduced anywhere.

5.2 EMG-informed MjID: agreement, cost and redundancy

In the estimation discussed above the model never saw the recorded muscle activity. The experiment of Sec. 4.3 introduces that activity as a soft constraint and evaluates the trade-off between a closer agreement with the supplied envelopes and a loss in the reproduction of the imposed trajectories (RQ3), while also asking how much of what one channel carries is already contained in the other five.

The recorded activity can be introduced, and the model follows it. The soft constraint raises the correlation with the reference in thirty-five of the thirty-six task–muscle pairs, and although most of those shifts do not reach the significance threshold, the result rests on their consistency. What matters is that the envelopes enter as a reference at all, instead of how the term is configured: the five configurations of Sec. 4.3.1 differ from the unconstrained solution far more than they differ from one another. The one adopted here, which uses the normalised envelope directly, separates significantly from the baseline, and it does so in both directions of the trade-off, binding more closely to the reference and tracking the imposed trajectory slightly worse. Amplitude is not part of that trade-off, since it behaves as it did without the reference in this configuration as in the other four.

The kinematic price depends on the task. At the wrist the movement is reproduced as it was before the constraint, and what degrades are the joints that were only to be held still, whose error roughly doubles in three of the four tasks. In the finger tasks the movement itself is spent, the error on the driven DOFs rising considerably and most of all in extension, and the degradation is ordered along the finger as the reconstruction of Sec. 3.1.7 is, small at the measured MCP and large at the PIP and at the DIP. The cost is therefore paid out of posture at the wrist and out of the movement at the finger.

Wrist extension is where the constraint changes the most: all six channels rise significantly in every participant, ECRL, ECU and FCU among the most improved, so the task on which the unconstrained model failed in timing is the one the recording repairs most completely. The constraint also fails to take hold on FDS, one of the two recorded channels that act on the fingers, and this matter is taken up in Sec. 5.3.

The second question is whether the information carried by the six constrained channels is redundant across them. Withholding one of them at a time, the withheld channel improves significantly on the case in which no recording is supplied at all, while remaining well below the case in which its own reference is present, and the other five behave as if nothing had been removed. Moreover, the reproduction of the imposed trajectories remains indistinguishable from the baseline. Thus, part of what a channel carries is contained in the others, and most of it is not.

Resolved by muscle, the redundancy across the six channels is not distributed evenly. The three flexors show a significant trend of losing their recoverability (intended as correlation) as their involvement in the movement grows: where they contribute little the model reproduces them well, and where they lead the movement

the correlation falls below zero. The extensors show no comparable dependence, their recovery holding at an intermediate level whatever their involvement. On the flexors the six channels are therefore not redundant in the one condition in which the redundancy would be useful.

The two experiments answer RQ3: the constraint changes the solution at a kinematic cost that is negligible on the movement at the wrist and not at the finger, and the six channels cannot stand in for one another where each of them leads.

5.3 Mechanical characterisation: the coverage of the six electrodes

The two sections above compare the model with a measurement. This one looks inside the model instead, and asks how the model distributes mechanical effort across the muscles for which no measurement exists, and how much of the joint torque the tasks require the six recorded muscles account for (RQ4). Nothing in Sec. 4.4 is tested against a recording, so what follows describes the model and, in doing so, explains part of what the two previous sections observed.

The starting point is how each muscle distributes its own mechanical action over the hand: the four wrist channels are wrist muscles and reach no finger, whereas the two channels used to observe the finger place about as much of their capacity on the wrist as on the digit they are named after, and the muscles that are confined to a single finger or to the thumb are almost all muscles that no electrode observes. This distribution belongs to the model itself since the capacity computed over postures drawn at random and the capacity weighted by the postures actually imposed agree almost exactly. The results that follow are read against it.

The torque the model spends divides into moving the joints the task drives and holding the others still, and only in the two deviations the movement itself takes the larger part. Everywhere else holding the wrist dominates, and it does so most sharply in the finger tasks, where bending one finger costs the model more in wrist stabilisation than in the movement itself. This follows from the same anatomy: the muscles that flex and extend the finger cross the wrist, so the torque they produce there has to be cancelled by other muscles, and the cancellation is not free.

How much of that torque the six recorded channels carry divides the hand the same way. What moves the wrist and what holds it still are almost entirely carried by them, so at the wrist the electrodes see nearly all of the mechanics. On the finger they do not: about half of the torque that moves the middle finger is produced by muscles no electrode observes, the deep flexor of that digit first among them, and the torque that holds the fingers and the thumb still is spread over many small contributions from the extensor compartment, the deep flexors and the thumb muscles, none of which is recorded. This is the answer to RQ4, and it is asymmetric: what the six electrodes leave unobserved is small at the wrist and large at the finger.

That asymmetry is also where the two previous sections met their limits. The

joints on which the constraint costs most, the PIP and DIP joints, are the joints whose torque the recorded channels cover least, and at the distal joint the recorded agonist is left almost alone, so the constraint fixes that joint's torque nearly by itself and the net torque of the two conditions is no longer the same in finger extension. A net torque that differs means the two solutions are not producing the same movement, which is the tracking failure of Sec. 4.3.3 seen on the quantity that causes it.

The same distribution accounts for the behaviour of FDS, the channel the constraint takes hold on least. It carries only part of the torque that moves the finger, it acts on the wrist about as much as on the finger, and the muscles that could produce the same finger torque in its place are not observed. Asking the solver to follow it is therefore asking for a signal whose mechanical role the model can fill in other ways, which is consistent with FDS being the channel that binds least, the one the other five recover worst, and the one whose disagreement in amplitude is the largest.

Co-contraction is where the model departs from what the literature leads one to expect. Solutions obtained by minimising muscle activity are known to under-represent antagonist activity, and this is the standard argument for supplying recorded activity to the solver [51]. Here that expectation holds only on the fingers, where the model co-contracts far less than the participants; at the wrist it is reversed, and the model co-contracts considerably more than they do. This may be accounted for by what the criterion is asked to resolve: in the work taken as an example a single joint rotates and the rest of the limb is free, whereas here all twenty-three DOFs of the hand are held to a prescribed configuration, some of them following the imposed trajectories and some of them only told to stay still, and, as the partition above shows, most of the torque goes into holding joints still rather than into moving them, which is work that opposing muscles have to do together. A further difference between the wrist and the finger tasks is that the index is computed on six channels only, so at the wrist it is read on three muscles against three, all of them recorded and all of them carrying most of the mechanics, whereas at the finger it rests on a single pair which is not even the pair that carries most of the recorded activity there. Where the model is asked to hold instead of moving, and where the electrodes see the muscles that do it, the minimisation produces too much co-contraction rather than too little. Supplying the recording, finally, moves the index very little, so that the unconstrained and the EMG-informed solutions differ hardly at all on it, which confirms that the constraint acts on the shape of each activation and not on the balance between the two groups.

It is also possible to make a last hypothesis concerning the wrist extension task. Overall in the wrist tasks the net torque on the driven DOF is almost the same with and without the reference, which is the mechanical counterpart of the kinematic reproduction being preserved. Additionally, the LOMO test adds that the extensors are not independent of one another, since they are the channels the remaining five recover best, so more than one distribution among them produces that same net torque. Together these two observations suggest why wrist extension is the task

the recording improves most. Without a reference nothing in the problem selects one distribution over another, and the solution settles where the cost is lowest; the envelopes supply the direction that was missing, and the agreement rises on all six channels at once. Testing this would mean constraining different subsets of the extensors and observing whether the solution moves (which was not done here).

Therefore, all these observations answer RQ4. The six electrodes account for most of the mechanics of the wrist and for about half of what moves the middle finger, they see almost nothing of what holds the fingers and the thumb still, and this distribution, fixed by the anatomy of the model, is what decides where the estimation and the constraint succeed.

5.4 Methodological choices and limitations

The kinematics that drives the model is extracted from video, and with this setup the extraction is reliable at the wrist and not at the fingers. The distal joint angles it returns are unusable and had to be replaced by fixed proportions of the MCP angle, and part of the repetitions even had to be discarded before the model, based on the failure of the MCP angle itself. Two of the three angles that drive a finger task are therefore not measured, but entirely reconstructed from the third and identical for every participant, so also the variability of finger coordination between people is absent from the input by construction. The same assumption is stretched furthest in fist closure, where eight interphalangeal angles are derived at once from a single pair of proportions applied to four digits, whose coupling is not necessarily the same. That task was withdrawn because the model did not reproduce the trajectory it was given, and since that trajectory is itself in large part assumed, a failure of the model and an inconsistent target cannot be told apart. This is why the exclusion is motivated as a failure to reproduce the imposed movement, but not as a statement about the kinematics of the participant's hand.

A second limitation comes from where the cameras are placed. Radial and ulnar deviation develop out of the image plane of the view from which the wrist angles are obtained, so in some participants the movement is not captured well enough for the repetitions to be located on the kinematics: three of the four participant–task combinations discarded before the model are ulnar deviations, and the fourth is a finger extension. A second camera does view the hand from the medial side, but the wrist angles cannot be recovered from it, since they are defined between the forearm and the palm and therefore require the body landmarks, which that view does not frame well enough for the pose model to estimate them reliably. Therefore, the deviation is tracked better from above than from the side, even though it develops out of the plane of that view, and a camera arrangement that observes the plane of deviation while keeping the body in frame would remove the restriction.

A second group of choices concerns what was deliberately not modelled. In every task the finger joints the task does not drive are held at the resting posture, whereas in the participant they move a little: the tenodesis effect displaces the

fingers when the wrist rotates, and the extrinsic muscles are multi-tendoned, so an isolated finger movement carries the adjacent digits with it. That residual movement falls on the channels where the kinematics are recovered worst, which is why it was discarded. The movement given to the model is therefore not exactly the movement the participants performed: the hand is asked to hold perfectly still what in reality moved a little.

Moreover, the envelope plays two roles in this work, as the term of comparison for the activations and as the reference for the controls, which sit on either side of the activation dynamics; and being low-pass filtered at 2 Hz it carries no fast variation, so it is easier to follow than a rougher signal would be, which may have favoured the agreement reported after the feasibility test.

Finally, two things are not quantified. Crosstalk between channels was not measured: an independent component analysis was attempted and abandoned, because the protocol contains too few and too similar movements for it to be informative. The simulated time covers slightly less than the video frame it represents, as previously declared, and no correction is applied for it.

5.5 Implications

What this work suggests follows from the asymmetry between the wrist and the finger. A six-channel electrode setup of forearm muscles is enough to describe the mechanics of the wrist, and it is not enough for the finger, where about half of the torque is produced by muscles no surface electrode can reach, and a study concerned with the fingers would need those muscles. The EMG-informed approach can be recommended at the wrist, where it makes the estimated activity resemble the sEMG activity at a negligible kinematic cost, and cannot be recommended on the fingers as long as the distal trajectories are reconstructed instead of being measured. In either case the estimated activations should be valued for their shape and their timing and not for their level, unless the strength of the model is scaled to the participant. The number of channels cannot be reduced by relying on what the others carry, at least not on the flexors, where a channel is least recoverable in the movements it leads. Co-contraction should be measured rather than inferred from the model: the disagreement with the recording does not have a fixed direction, being an overestimation at the wrist and an underestimation at the finger, so a solution of this kind cannot be relied on for the balance between agonists and antagonists in a task that has not been recorded.

Chapter 6

Conclusions and future work

6.1 Conclusions

This work brought together markerless kinematics extracted from video and six-channel sEMG recorded on the forearm of seven participants, and used them to validate and to drive an inverse-dynamics estimation of muscle activity on a MSK model of the hand and wrist. Three experiments were built on that basis and each rests on the one before it: the estimated activations were compared with the measured envelopes, the envelopes were then supplied to the solver as a soft constraint and the cost of doing so was measured, and the mechanics of the model itself was described in order to ask how much of it the six electrodes account for.

Muscle activations can be estimated from markerless kinematics alone, wherever the model reproduces the movement it is given, which is the case at the wrist and on the fingers except at the most distal joint (RQ1). Their agreement with the recording is partial (RQ2): the timing of the activity is reproduced on the muscles that lead a task, in some of the tasks, while the amplitude is not reproduced on any muscle.

Introduced into the estimation as a soft constraint, the recorded activity changes the solution and the model follows it (RQ3). The kinematic cost of doing so depends on the task: at the wrist the movement is reproduced as it was before, and what degrades are joints that were only being held still, whereas in the finger tasks the movement itself is spent, most of all on the joints whose trajectory is reconstructed. Withholding one channel at a time shows that part of what a channel carries is contained in the other five and most of it is not, and that on the flexors it is precisely in the movements they lead that a channel cannot be recovered from the others.

The mechanical description of the model answers the last question and explains the other three (RQ4). The six recorded channels account for nearly all of the torque that moves the wrist and holds it still, for about half of the torque that moves the middle finger, and for very little of what keeps the fingers in place, since the muscles that do that work are deep or intrinsic and no surface electrode reaches them.

Three further points emerge from this work. The first is that the estimation method now has a term of comparison it did not have, that is an independent measurement of the quantity it produces, on real participants rather than in simulation.

The second is that the dissociation between timing and amplitude survives placing the envelopes in the inverse-dynamics algorithm. The third is that the coverage map is knowable in advance: it depends on the anatomy of the model and not on the data, so it can be consulted before an experiment to decide where a surface electrode placement is sufficient and where it is not. The limitations that bound these conclusions are set out in Sec. 5.4, and the reconstruction of the distal finger joints is the assumption that bounds them most.

6.2 Future work

Three developments follow directly from the discussion. The first is to scale the strength of the model to the participant, which would be the safest way to overcome the disagreement in amplitude: as long as the maximal isometric force of every actuator is the same for everyone, the estimated activations and the normalised envelopes are bounded above by different quantities. The second is to evaluate the fitted parameters out of sample, since they are at present fitted and used on the same movement, so nothing establishes that they transfer to a task the optimisation has not seen. The third is to test the reading proposed for the wrist extensors, by constraining different subsets of them and observing whether the solution moves: if it does not, the choice among those muscles is indeed weakly determined without a sEMG recording.

Three developments concern the acquisition setup. A camera arrangement that observes the plane of radial and ulnar deviation while keeping the body in frame would remove the restriction that cost this work three participant–task combinations in that task alone. Measuring the finger joints instead of reconstructing them is the condition for everything that concerns the fingers: it is what would allow the constraint to be applied there without spending the movement, and it is also what would allow fist closure to be reconsidered. A protocol containing more varied movements would make an independent component analysis informative and would allow the crosstalk between channels to be quantified. One question still remains open: how the deep flexors and the intrinsic muscles could be observed at all, given that surface electrodes do not reach them and that intramuscular recording is not a practical option. Until it has an answer, the half of the finger mechanics they carry stays unobserved.

A longer-term direction concerns what happens beyond the hand. The same framework, that is markerless kinematics driving a MSK model whose activations are read against sEMG recordings, is not specific to the hand and could be transferred to other districts, the lower limb first among them, where personalised neuromusculoskeletal pipelines are established [55]. Those pipelines have been built within OpenSim, so whether an equivalent can be assembled around a MuJoCo model remains open. The hand represents a place to verify such an approach first, since it offers a case in which the same muscles can be both recorded and modelled, so that the estimate can be checked before it is used where no check is available.

Appendix A

Reproducibility material

This appendix collects the settings needed to reproduce the processing pipeline described in Chapter 3. Each table is referenced from the section that introduces the corresponding step, and reports the values actually used, including the choices that were made by inspection rather than by an automatic rule.

Software and hardware

The whole pipeline is written in Python. The recorded signals and the simulation outputs are stored in HDF5 files and handled with `numpy`, `scipy`, `pandas` and `h5py`; the figures are produced with `matplotlib` and the video files are read with `opencv`.

The kinematics of Sec. 3.1.2 is extracted with MediaPipe, using two of its landmark models, *pose_landmarker_heavy* for the body and *hand_landmarker* for the hand. The two model files are obtained once from the public MediaPipe endpoints and reused from then on. The address they are obtained from always serves the current release, so a pipeline rebuilt from scratch at a later date may receive a different version of them.

The MSK model is MyoHand as distributed with MyoSuite [29], and the copy actually used is the one included with the code of this work. The simulations run in MuJoCo [31], and the quadratic program of Sec. 3.1.9 is solved with OSQP. The subject-specific parameters of Sec. 3.1.10 are fitted with the differential evolution implementation of `scipy.optimize`, which evaluates the population in parallel over the available cores.

The parameter optimisation was run on a laboratory server, under Linux and with thirty-two cores available, since it repeats a full simulation of the movement at every evaluation of the cost. Every other step, from the extraction of the kinematics to the analyses and the figures, was run on a laptop with an eight-core processor and 16 GB of memory, under Windows 11.

Normalisation reference

Table A.1 lists, for every participant and channel, the two candidate normalisation references and the one that was adopted. As described in Sec. 3.1.4, the reference is the larger of the amplitude reached during the dedicated calibration contraction and the peak reached during the recorded tasks, so for a substantial fraction of the channels the denominator of the MVC scale is a task peak rather than a calibration value.

The two candidates are not the same statistic: the calibration value is a mean over a 0.5 s window centred on the peak, whereas the task value is an instantaneous maximum. The last column reports which candidate would have been selected had the comparison been made between homogeneous statistics, on either side. The reference changes in a small minority of channels, which indicates that the asymmetry of the statistics is not what drives the selection: the tasks genuinely exceed the calibration.

Table A.1: Normalisation reference adopted for each participant and channel. The last column reports which candidate would have been selected had the comparison been made between homogeneous statistics, either as two instantaneous peaks or as two means over 0.5 s.

Subj.	Muscle	Calib. (μV)	Task (μV)	Reference	If uniform (peak/mean)
S01	FDS	244.2	104.2	calib	calib/calib
S01	FCR	338.1	124.2	calib	calib/calib
S01	FCU	94.4	115.4	task	task/calib
S01	EDC	251.3	153.8	calib	calib/calib
S01	ECRL	246.7	209.2	calib	calib/calib
S01	ECU	137.3	108.2	calib	calib/calib
S02	FDS	138.7	91.0	calib	calib/calib
S02	FCR	319.2	36.4	calib	calib/calib
S02	FCU	117.5	131.7	task	task/calib
S02	EDC	583.2	376.5	calib	calib/calib
S02	ECRL	140.4	251.8	task	task/task
S02	ECU	152.5	303.8	task	task/task
S03	FDS	427.9	314.0	calib	calib/calib
S03	FCR	288.2	321.4	task	task/calib
S03	FCU	128.5	82.4	calib	calib/calib
S03	EDC	477.7	401.5	calib	calib/calib
S03	ECRL	259.0	371.2	task	task/task
S03	ECU	232.5	281.1	task	task/task
S04	FDS	93.6	262.9	task	task/task
S04	FCR	163.8	181.8	task	task/calib
S04	FCU	105.7	171.3	task	task/task
S04	EDC	364.5	537.0	task	task/task
S04	ECRL	116.9	143.9	task	task/task
S04	ECU	193.0	185.7	calib	calib/calib
S05	FDS	107.5	132.9	task	task/calib

Continued on the next page

Continued from the previous page

Subj.	Muscle	Calib. (μV)	Task (μV)	Reference	If uniform (peak/mean)
S05	FCR	141.7	286.0	task	task/task
S05	FCU	126.3	272.4	task	task/task
S05	EDC	468.6	279.4	calib	calib/calib
S05	ECRL	127.1	362.0	task	task/task
S05	ECU	297.2	251.0	calib	calib/calib
S06	FDS	231.6	318.5	task	task/task
S06	FCR	256.3	445.2	task	task/task
S06	FCU	155.9	325.1	task	task/task
S06	EDC	724.0	690.3	calib	calib/calib
S06	ECRL	189.4	248.7	task	task/task
S06	ECU	253.6	379.6	task	task/task
S07	FDS	131.9	122.2	calib	calib/calib
S07	FCR	95.2	187.9	task	task/task
S07	FCU	166.6	179.1	task	calib/calib
S07	EDC	133.2	198.4	task	task/task
S07	ECRL	120.4	155.4	task	task/task
S07	ECU	88.6	78.2	calib	calib/calib

Segmentation settings

Table A.2 reports the settings used to segment the repetitions within each trial, described in Sec. 3.1.5. Thresholds were tuned per participant and task, because the amplitude and the regularity of the movement vary considerably across participants; where the reference kinematic channel did not allow the repetitions to be located, the segmentation was performed on the sEMG envelope of the expected agonist instead.

The threshold is applied to the detrended reference channel and is given in one of two forms: adaptive, as k standard deviations of that channel away from its initial level, or fixed at an absolute value. Fixed thresholds are in degrees where the reference channel is kinematic and in normalised envelope units where it is an sEMG envelope, the latter marked *env*, while k is dimensionless in both cases. One channel is marked *inv*: there the movement of interest appears as a decrease, so the sign is reversed before thresholding.

The two ends of the trial are delimited differently. *Pre-roll* is the amount of signal kept before the first repetition annotated during the session, whereas *end trim* is the amount discarded from the end, with *abs.* marking the trials whose end was given as an absolute timestamp of the session. Both have a default that applies wherever the value was not set explicitly, so no trial is left untrimmed, and for the first participant both ends are fixed directly in samples.

The last column lists the repetitions that were discarded by visual inspection, which is the only part of the segmentation that cannot be reproduced from the settings alone.

Table A.2: Settings used to locate the repetitions within each trial, for every participant and task.

Subj.	Task	Threshold	Margin (s)	Pre-roll (s)	End trim (s)	Excl. reps
S01	Wrist flex.	$k = 1.5$	0.800	fixed	fixed	–
S01	Wrist ext.	$k = 1.875$	0.700	fixed	fixed	–
S01	Radial dev.	$30°$	0.700	fixed	fixed	–
S01	Ulnar dev.	$k = 1.5$	0.700	fixed	fixed	–
S01	Finger flex.	$k = 1.875$	0.700	fixed	fixed	–
S01	Finger ext.	$k = 1.65$	0.650	fixed	fixed	–
S01	Fist	$k = 0.375$	0.650	fixed	fixed	–
S02	Wrist flex.	$k = 1.2$	0.800	5	3	–
S02	Wrist ext.	$k = 1.5$	0.700	5	3	–
S02	Radial dev.	$k = 1.125$	0.700	5	3	–
S02	Ulnar dev.	$-2.1°$	0.700	5	5	–
S02	Finger flex.	$k = 1.8$	0.700	5	3	–
S02	Finger ext.	$k = 0.75$	0.650	4	25	–
S02	Fist	$k = 0.375$	0.650	5	3	–
S03	Wrist flex.	$k = 1.5$	0.800	5	3	–
S03	Wrist ext.	$k = 1.425$	1.300	5	3	–
S03	Radial dev.	$k = 1.125$	0.700	5	3	–
S03	Ulnar dev.	$k = 1.5$ (env)	0.700	5	3	–
S03	Finger flex.	$k = 1.35$	0.700	5	3	–
S03	Finger ext.	$k = 1.17$ (inv)	0.975	5	3	1
S03	Fist	$k = 0.375$	0.650	5	3	–
S04	Wrist flex.	$49°$	0.800	5	abs.	–
S04	Wrist ext.	$k = 1.875$	0.700	5	3	–
S04	Radial dev.	$k = 1.125$	0.700	5	abs.	–
S04	Ulnar dev.	$k = 1.5$ (env)	0.700	5	3	–
S04	Finger flex.	$k = 1.05$	0.700	5	abs.	–
S04	Finger ext.	$k = 0.75$	0.650	5	3	–
S04	Fist	$k = 0.375$	0.650	5	3	1
S05	Wrist flex.	$k = 1.5$	0.800	5	3	–
S05	Wrist ext.	$k = 1.875$	0.700	5	3	–
S05	Radial dev.	$k = 1.125$	0.700	5	3	–
S05	Ulnar dev.	$k = 1.5$	0.700	5	abs.	–
S05	Finger flex.	$k = 1.875$	0.700	5	3	–
S05	Finger ext.	$k = 0.75$	0.650	5	3	–
S05	Fist	$k = 0.375$	0.650	5	3	–
S06	Wrist flex.	$k = 1.5$	0.800	5	3	–
S06	Wrist ext.	$k = 1.575$	0.700	5	3	–
S06	Radial dev.	$k = 2.19$	0.700	5	3	–
S06	Ulnar dev.	$-1°$	0.700	5	abs.	–
S06	Finger flex.	$k = 1.65$	0.700	5	3	–
S06	Finger ext.	0.2 (env)	0.650	5	3	–
S06	Fist	$k = 0.375$	0.650	5	3	–
S07	Wrist flex.	$40°$	0.800	5	3	1, 2, 6, 18
S07	Wrist ext.	$-5°$	0.700	3.5	3	1, 2, 3, 4, 7

Continued on the next page

Continued from the previous page

Subj.	Task	Threshold	Margin (s)	Pre-roll (s)	End trim (s)	Excl. reps
S07	Radial dev.	35°	0.700	5	3	1
S07	Ulnar dev.	0.3 (env)	0.700	1	3	–
S07	Finger flex.	$k = 0.6$	0.700	2	6	–
S07	Finger ext.	162.5°	0.650	5	3	1
S07	Fist	$k = 0.375$	0.650	5	3	–

Corrections to the resting posture

Table A.3 lists the joints whose resting angle was modified before being used as the reference configuration of the model, as described in Sec. 3.1.7. Two corrections are applied. DIP angles estimated above 60° were capped at that value, because at rest they reflect the occlusion of the distal phalanx rather than an actual flexion; the remaining corrections clip the estimated angle to the range admitted by the MSK model.

Table A.3: Joints whose resting-posture angle was modified before being passed to the model. *Estimated* is the mean angle measured over the quiet baseline interval; *adopted* is the value actually imposed.

Subj.	Joint	Estimated (°)	Adopted (°)	Reason
S01	Thumb CMC	56.9	51.1	Model limit
S01	Index DIP	63.2	60.0	DIP cap (60°)
S01	Middle DIP	77.1	60.0	DIP cap (60°)
S01	Ring DIP	90.6	60.0	DIP cap (60°)
S01	Pinky DIP	94.8	60.0	DIP cap (60°)
S03	Middle DIP	83.1	60.0	DIP cap (60°)
S03	Ring DIP	97.2	60.0	DIP cap (60°)
S03	Pinky DIP	101.2	60.0	DIP cap (60°)
S04	Pinky PIP	109.3	105.0	Model limit
S06	Thumb CMC	52.3	51.1	Model limit
S06	Ring DIP	83.5	60.0	DIP cap (60°)
S06	Pinky DIP	87.4	60.0	DIP cap (60°)

SUBJ_05 is an exception and is not covered by the table, since its baseline recording captured the hand in a partially clenched position, which would have placed the reference configuration far from a neutral posture; the resting angles of the whole hand were therefore replaced with those estimated for *SUBJ_02*.

Joint ranges of the model

The model of Sec. 3.1.8 is used as distributed in its geometry and its inertial properties, but not in its joint limits. Table A.4 gives, for every DOF, the range distributed with MyoHand beside the range adopted here.

Seventeen of the twenty-three DOFs were widened for a practical reason: a target configuration that falls outside the admitted range cannot be imposed on the model at all, so a limit narrower than the movement recorded would truncate the trajectory the model is asked to follow.

Each limit was moved to a value taken from the clinical literature on the ROM of the hand and wrist, and only outward. Two of the widenings, the hyperextension of the finger MCP and DIP joints, rest on the clinical observation that a small hyperextension is common in the general population rather than on a specific measurement study. One is set to a mean instead of an end of a reported range: at the MCP joint of the thumb the spread between individuals is so wide, from 16 to 90° of flexion, that its upper end describes one participant instead of a normal range, so the mean of 59° was used instead.

The names of the two DOFs of the thumb carpometacarpal joint do not correspond to the clinical ones. Moving each of them in isolation and measuring the displacement of the thumb with respect to the plane of the palm shows that *cmc_abduction* produces a rotation almost entirely within that plane, 61° against 17° out of it, which is clinically a flexion–extension; whereas *cmc_flexion* produces a rotation almost entirely out of it, 51° against 2° within it, which is clinically the palmar abduction. The table names the joints as the model does and describes them as the clinic does, which is why the two appear to be exchanged.

Table A.4: Joint ranges of the MyoHand model, as distributed and as adopted here. Angles in degrees; the last column gives the basis of each widening.

Joint	Description	As distributed	Adopted	Basis
pro_sup	Forearm pronation–supination	[−90, 90]	[−90, 90]	unchanged
flexion	Wrist flexion–extension	[−45, 45]	[−70, 80]	[56, 57]
deviation	Wrist radial–ulnar deviation	[−10, 25]	[−33, 25]	[56, 57]
cmc_abduction	Thumb CMC flexion–extension	[−29, 45]	[−29, 45]	unchanged
cmc_flexion	Thumb CMC palmar abduction	[−45, 40]	[−45, 51]	[58]
mp_flexion	Thumb MCP flexion	[−45, 40]	[−45, 58]	[59]
ip_flexion	Thumb IP flexion	[−75, 25]	[−75, 90]	[60]
mcp2–5_flexion	Finger MCP flexion–extension	[0, 90]	[−20, 90]	clinical convention
mcp2–5_abduction	Finger MCP abduction	[−15, 15]	[−15, 15]	unchanged
pm2–5_flexion	Finger PIP flexion	[0, 90]	[0, 105]	[60]
md2–5_flexion	Finger DIP flexion–extension	[0, 90]	[−10, 90]	clinical convention

Settings of the parameter optimisation

Table A.5 reports the search bounds of the subject-specific parameter optimisation described in Sec. 3.1.10, and Table A.6 the settings of the optimisation algorithm. A separate optimisation is run for each participant and each task, so the parameter sets are not shared across tasks.

The same bounds and the same settings are used when the parameters are fitted again with the control reference present, for the configuration carried forward in

Sec. 4.3.1. The one exception is *RE*, whose range becomes 100 to 10000 and which is sampled logarithmically: the lower end is the bound of Sec. 3.3.3.

Table A.5: Search bounds of the subject-specific parameter optimisation. A separate optimisation is run for each participant and task.

Parameter	Meaning	Lower	Upper
KP	Proportional gain of the PD controller	1000	15000
KV	Derivative gain of the PD controller	10	1000
RR_WRIST	Weight of the torque replication term (wrist and abduction DOFs)	0	5000
RR_FINGERS	Weight of the torque replication term (finger flexion DOFs)	100	5000
RE	Weight of the control reference term	0.001	10
RS	Weight of the smoothness term	40	100

Table A.6: Settings of the differential evolution algorithm.

Setting	Value
Strategy	best1bin
Maximum number of generations	50
Population size multiplier	5
Convergence tolerance	0.01
Mutation range	(0.5, 1)
Recombination probability	0.7
Local polishing of the best solution	yes

Optimised parameter values

Table A.7 reports the values the optimisation converged to, for every participant and task that enters the model. The parameters are given for reproducibility and not for interpretation, since they are not physiological quantities but the coefficients an algorithm converged to.

One characteristic of the values obtained is that the value reported is a property of the range that was allowed as much as of the data, since several parameters settle at the edge of their admitted range, taking as the edge its outer one per cent. The proportional gain reaches its lower bound in 26 of the 45 fitted combinations and the smoothness weight reaches one of its bounds in 18; the error weight reaches its upper bound in four.

Table A.7: Parameter values returned by the optimisation, for every participant and task. The combinations excluded at the segmentation stage are omitted.

Subj.	Task	KP	KV	RR_WRIST	RR_FINGERS	RE	RS
S01	Wrist flex.	1034	42.53	1348	4135	1.41	47.01
S01	Wrist ext.	1023	28.81	1562	4914	7.53	40.00
S01	Radial dev.	1025	45.77	1336	4304	1.08	45.45
S01	Ulnar dev.	2176	75.85	2914	4944	3.69	41.51

Continued on the next page

Continued from the previous page

Subj.	Task	KP	KV	RR_WRIST	RR_FINGERS	RE	RS
S01	Finger flex.	1072	43.27	3428	4928	3.45	43.16
S01	Finger ext.	1186	48.07	3190	4254	3.83	40.14
S01	Fist	1026	38.23	1643	3034	1.31	99.58
S02	Wrist flex.	1014	16.27	859	4640	1.62	93.02
S02	Wrist ext.	1007	29.88	786	4407	10.00	40.00
S02	Radial dev.	1011	38.91	1937	4643	10.00	40.00
S02	Ulnar dev.	1353	77.37	2954	4849	1.78	48.55
S02	Finger flex.	1009	38.94	4142	2383	9.39	44.70
S02	Finger ext.	1384	76.16	4298	3978	9.62	40.00
S02	Fist	4244	151	433	3768	9.31	40.00
S03	Wrist flex.	1000	39.49	1641	4896	2.36	40.00
S03	Wrist ext.	1057	17.62	1219	4890	3.19	40.00
S03	Radial dev.	1237	19.88	771	4311	1.54	72.58
S03	Finger flex.	1038	21.50	3549	4348	3.32	71.24
S03	Finger ext.	1694	113	4085	3724	2.13	40.00
S03	Fist	7123	76.74	402	2624	10.00	40.00
S04	Wrist flex.	1008	34.56	914	4049	0.45	82.76
S04	Wrist ext.	1187	32.13	584	4160	0.17	46.90
S04	Radial dev.	1022	21.85	766	4794	0.50	77.00
S04	Finger flex.	1038	26.15	3079	3902	2.83	77.16
S04	Finger ext.	1172	29.58	4666	4468	4.83	76.88
S04	Fist	1032	32.94	984	4894	1.07	65.07
S05	Wrist flex.	1039	58.58	727	4680	4.11	100
S05	Wrist ext.	1238	34.60	142	686	9.43	41.02
S05	Radial dev.	1456	79.17	512	2505	5.67	54.24
S05	Ulnar dev.	1219	210	576	3604	5.64	65.93
S05	Finger flex.	1125	28.83	3776	4858	7.61	52.65
S05	Finger ext.	1489	113	1015	3213	9.01	70.89
S05	Fist	6373	75.15	766	3181	7.33	67.56
S06	Wrist flex.	1015	47.00	1878	4722	4.91	40.00
S06	Wrist ext.	1022	15.19	1573	4976	10.00	40.00
S06	Radial dev.	1077	30.96	560	3930	1.99	68.82
S06	Ulnar dev.	1757	73.25	3178	4816	3.40	45.46
S06	Finger flex.	1067	33.75	3283	4528	4.36	44.76
S06	Fist	1000	39.20	1703	4081	1.77	100
S07	Wrist flex.	1000	14.96	831	4933	1.36	40.00
S07	Wrist ext.	1199	14.07	1046	4939	2.54	40.00
S07	Radial dev.	1015	23.67	475	4294	0.22	55.76
S07	Finger flex.	1430	33.05	3168	4807	6.55	49.23
S07	Finger ext.	1065	53.86	4320	3675	1.46	65.21
S07	Fist	1249	31.62	577	4828	0.90	40.00

Table A.8 reports the same quantities for the configuration carried forward in Sec. 4.3.1, whose parameters are fitted again with the reference present as described in Sec. 3.3.4. That fit is run on the six tasks that reach the feasibility experiment, so fist closure is not included. *RE*, the weight of the control reference, which in the

unconstrained solution has no reference to weigh, here sits around its lower bound in every combination.

Table A.8: Parameter values returned by the optimisation for the configuration carried forward in Sec. 4.3.1, for every participant and task.

Subj.	Task	KP	KV	RR_WRIST	RR_FINGERS	RE	RS
S01	Wrist flex.	5613	14.71	1469	4951	100	40.00
S01	Wrist ext.	7147	30.54	574	4989	100	40.00
S01	Radial dev.	4358	10.00	1796	4793	100	42.10
S01	Ulnar dev.	13211	347	1738	4721	100	48.96
S01	Finger flex.	3876	20.12	4792	4861	100	90.89
S01	Finger ext.	14935	81.83	4066	4997	100	43.70
S02	Wrist flex.	4068	31.19	331	4988	100	40.96
S02	Wrist ext.	9862	51.66	409	4983	100	58.84
S02	Radial dev.	6049	267	124	4929	100	58.44
S02	Ulnar dev.	14977	426	1176	4969	100	59.50
S02	Finger flex.	3652	10.49	3883	4999	100	57.11
S02	Finger ext.	5004	306	4693	4957	100	74.50
S03	Wrist flex.	3269	74.64	1391	4926	100	40.00
S03	Wrist ext.	9710	37.02	464	4973	100	40.00
S03	Radial dev.	3440	14.27	1519	4864	100	40.07
S03	Finger flex.	4168	49.71	4772	4753	100	45.12
S03	Finger ext.	14823	228	4708	4909	100	63.42
S04	Wrist flex.	6310	22.29	615	4899	100	40.00
S04	Wrist ext.	14468	109	344	4963	100	40.00
S04	Radial dev.	10082	83.39	208	3855	100	91.44
S04	Finger flex.	4541	26.43	4815	4777	100	50.40
S04	Finger ext.	10590	25.76	2445	4917	105	59.40
S05	Wrist flex.	1870	56.13	1514	4924	100	44.46
S05	Wrist ext.	1548	85.03	729	4017	103	85.34
S05	Radial dev.	12626	68.31	320	4789	100	74.32
S05	Ulnar dev.	14843	418	2303	4947	100	40.00
S05	Finger flex.	2601	11.19	4567	4839	100	60.35
S05	Finger ext.	9203	31.52	4841	4928	100	61.13
S06	Wrist flex.	1181	73.96	2891	5000	100	52.18
S06	Wrist ext.	3723	40.64	403	4863	100	40.74
S06	Radial dev.	6784	102	169	4525	100	92.87
S06	Ulnar dev.	14933	429	1690	4977	100	44.41
S06	Finger flex.	4975	39.66	3838	4995	100	56.83
S07	Wrist flex.	10552	47.12	380	4910	100	48.31
S07	Wrist ext.	11396	68.92	445	4959	100	40.00
S07	Radial dev.	14960	88.79	297	4913	100	43.70
S07	Finger flex.	7192	25.78	3623	4990	100	74.47
S07	Finger ext.	14185	474	3801	4959	100	48.26

Appendix B

Supporting data

This appendix documents the material on which the results rest: how much data survives each stage of the pipeline, what the trials that were discarded look like, and how much the repetitions of a single trial differ from one another. The last section steps outside the six recorded channels and shows what the model does with all of its actuators. Each table or figure is referenced from the section that introduces the corresponding step

Repetitions available at each stage

Table B.1 reports, for every participant and task, the number of repetitions identified by the segmentation step and the number that reaches the model. The two counts differ because two automatic criteria are applied when the model inputs are built, both described in Sec. 3.1.7: a repetition is discarded if the joint that drives the task shows more than one prominent peak and it is discarded if its deviation from the average of the other repetitions exceeds three standard deviations.

Combinations marked *n.a.* do not enter the model at all: they are the four participant–task combinations whose repetitions could not be located on the kinematics, described in Sec. 3.1.5. Fist closure is treated differently, since its repetitions are simulated together with the others and are reported here with an asterisk, but the task is set aside after the tracking results of Sec. 4.2.1.

Table B.1: Number of repetitions available at three stages of the pipeline: after the segmentation step, after the automatic pruning applied when building the model inputs and after the exclusion of the fist task (as indicated by the asterisk).

Subj.	Task	Segmented	Model
S01	Wrist flex.	21	13
S01	Wrist ext.	21	21
S01	Radial dev.	21	21
S01	Ulnar dev.	21	21
S01	Finger flex.	21	19
S01	Finger ext.	20	18

Continued on the next page

Continued from the previous page

Subj.	Task	Segmented	Model
S01	Fist	21	19*
S02	Wrist flex.	21	13
S02	Wrist ext.	21	20
S02	Radial dev.	21	20
S02	Ulnar dev.	11	11
S02	Finger flex.	21	21
S02	Finger ext.	20	18
S02	Fist	21	20*
S03	Wrist flex.	21	20
S03	Wrist ext.	21	20
S03	Radial dev.	21	18
S03	Ulnar dev.	21	*n.a.*
S03	Finger flex.	21	19
S03	Finger ext.	18	17
S03	Fist	21	17*
S04	Wrist flex.	21	17
S04	Wrist ext.	19	18
S04	Radial dev.	18	16
S04	Ulnar dev.	21	*n.a.*
S04	Finger flex.	20	7
S04	Finger ext.	21	18
S04	Fist	19	4*
S05	Wrist flex.	21	21
S05	Wrist ext.	21	21
S05	Radial dev.	20	19
S05	Ulnar dev.	21	19
S05	Finger flex.	20	18
S05	Finger ext.	21	18
S05	Fist	21	21*
S06	Wrist flex.	21	20
S06	Wrist ext.	21	21
S06	Radial dev.	20	20
S06	Ulnar dev.	20	19
S06	Finger flex.	21	17
S06	Finger ext.	21	*n.a.*
S06	Fist	21	9*
S07	Wrist flex.	14	7
S07	Wrist ext.	11	11
S07	Radial dev.	20	15
S07	Ulnar dev.	21	*n.a.*
S07	Finger flex.	21	8
S07	Finger ext.	20	20
S07	Fist	21	1*
Total		**983**	**751**

∗ fist closure is simulated like the other tasks and is withdrawn as reported in Sec. 4.2.1 so 660 repetitions over 38 participant–task combinations reach the validation and the feasibility experiment.

Trials that could not be segmented

Fig. B.1 shows what those four excluded subject-task combinations look like. Both panels report the reference kinematic channel of the same task, ulnar deviation, for two different participants and on the same vertical scale. On the right the individual repetitions are unmistakable: each one is a deflection of at least ten degrees, well separated from the resting level, and the boundaries found by the segmentation coincide with them. On the left the same movement leaves no comparable trace: the channel fluctuates by a few degrees throughout the trial, with no excursion that can be told apart from the rest.

The participant did perform the task as confirmed by the sEMG; what failed is the extraction from the video of the wrist angle or of the middle finger MCP angle (for *SUBJ_06* in the finger extension task). This is why the repetitions of these trials were located on the sEMG envelope of the agonist instead, and also why the same trials cannot be used to drive the model.

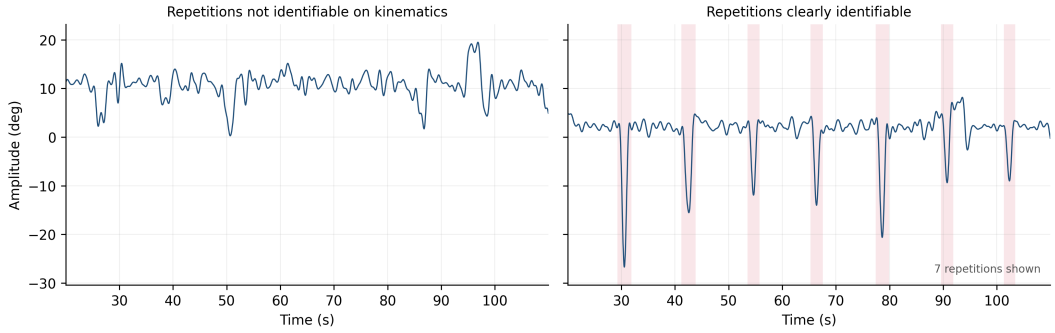


Figure B.1: Reference kinematic channel of the ulnar deviation task for two participants, over a 90 s window of the trial and on the same vertical scale. Right: a trial segmented on kinematics, with the repetitions found by the segmentation shaded. Left: one of the four combinations for which the same procedure finds nothing.

Variability between repetitions

Sec. 3.2.5 states that the metrics are computed on the individual repetitions rather than on the average of a participant's repetitions, and that the two need not agree. Fig. B.2 shows one case of each on real data.

In the left panel the repetitions of both signals are tightly grouped, their averages resemble one another and the two calculations give the same answer. In the right panel the modelled activations vary considerably from one repetition to the next, several of them placing their peak in a different part of the cycle from the others. Averaging them produces a curve that resembles none of them, and that resembles the measured envelope less than the individual repetitions do, as the correlations show (the correlation computed on the averaged curves is negative, while the correlation computed repetition by repetition is moderately positive).

The dispersion displayed in this figure is internal to one participant and one trial.

It is not the dispersion shown in the figures of the results, where the band around a curve describes the spread across participants, each of them already represented by the average of their own repetitions: the two are distinct quantities, and only the second is reported.

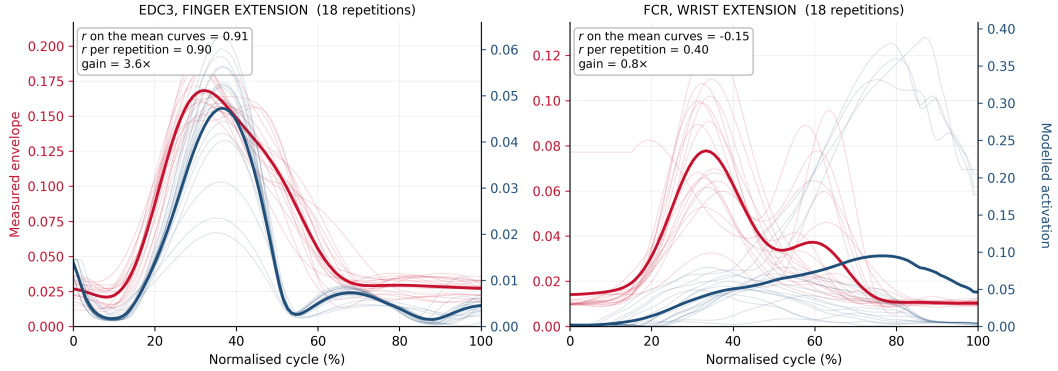


Figure B.2: Measured envelope and modelled activation over the normalised cycle, for one participant, one task and one muscle in each panel. Thin lines are the individual repetitions and thick lines their averages.

Controls of the whole model

Sec. 4.3.2 examines the six recorded channels, which are the only ones for which a reference exists and the remaining thirty-three actuators receive none, since no figure of that section shows what the model does with them. The six figures below fill that gap: for each task they show the group average control of every actuator of the model, in the unconstrained solution and in the EMG-informed one, with the reference traced on the six recorded channels.

The panels are arranged by anatomy: the first two rows hold the muscles that act on the wrist (the primary movers and then the two pronators), the third holds the thumb and the last four take one finger each, from the index to the little finger. Within a finger row the columns are, in order, the superficial flexor, the deep flexor, the extensor, and the three intrinsic muscles of that digit, followed where one exists by the extensor proper to it (which are Extensor Indicis Proprius on the index and Extensor Digiti Minimi on the little finger).

Within each figure the left axis is common to all thirty-nine panels, so the panels are comparable with one another, but it is not common across figures, since the tasks differ in amplitude by a factor of four. The reference (envelope) is drawn against the right axis, in red, and only on the six panels that have one (marked by a gold frame); its scale is likewise common within a figure. The constraint does not act only on the recorded muscles: in wrist flexion, for instance, EDC2, EDC4, EDC5 and the ulnar interosseous of the fourth finger change between the two conditions. Most of the thirty-three unrecorded actuators nonetheless stay close to zero in every task.

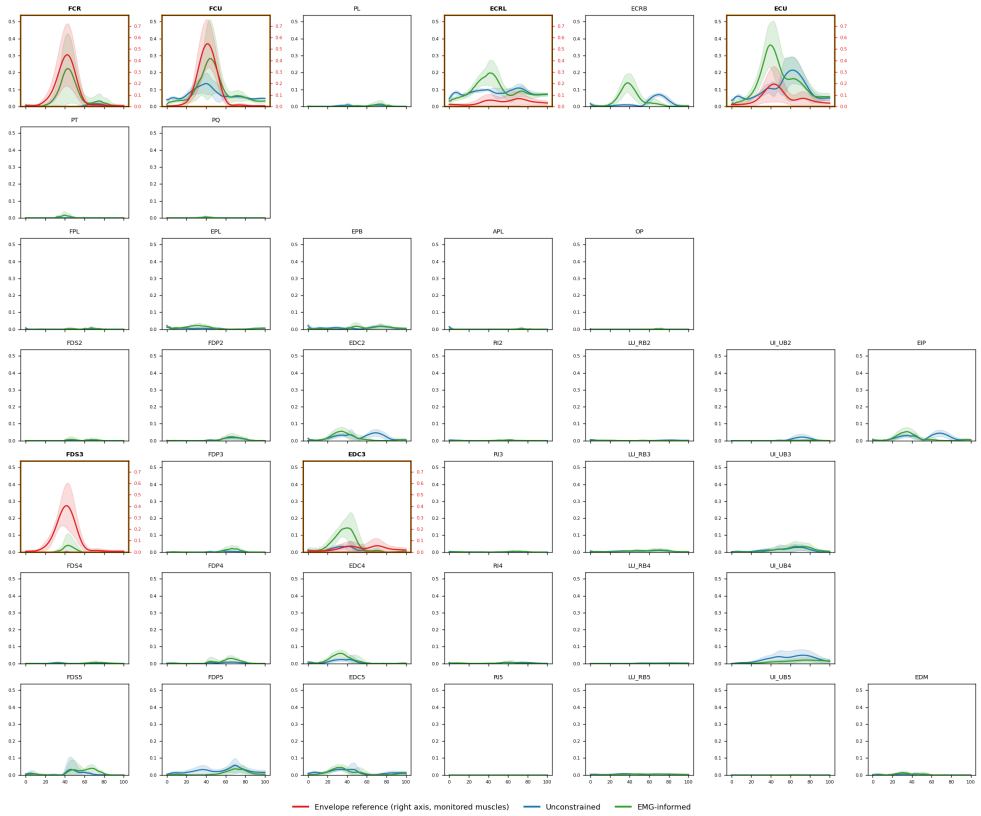


Figure B.3: Group average control of the thirty-nine actuators of the model in wrist flexion. Line: mean across participants; band: standard deviation.

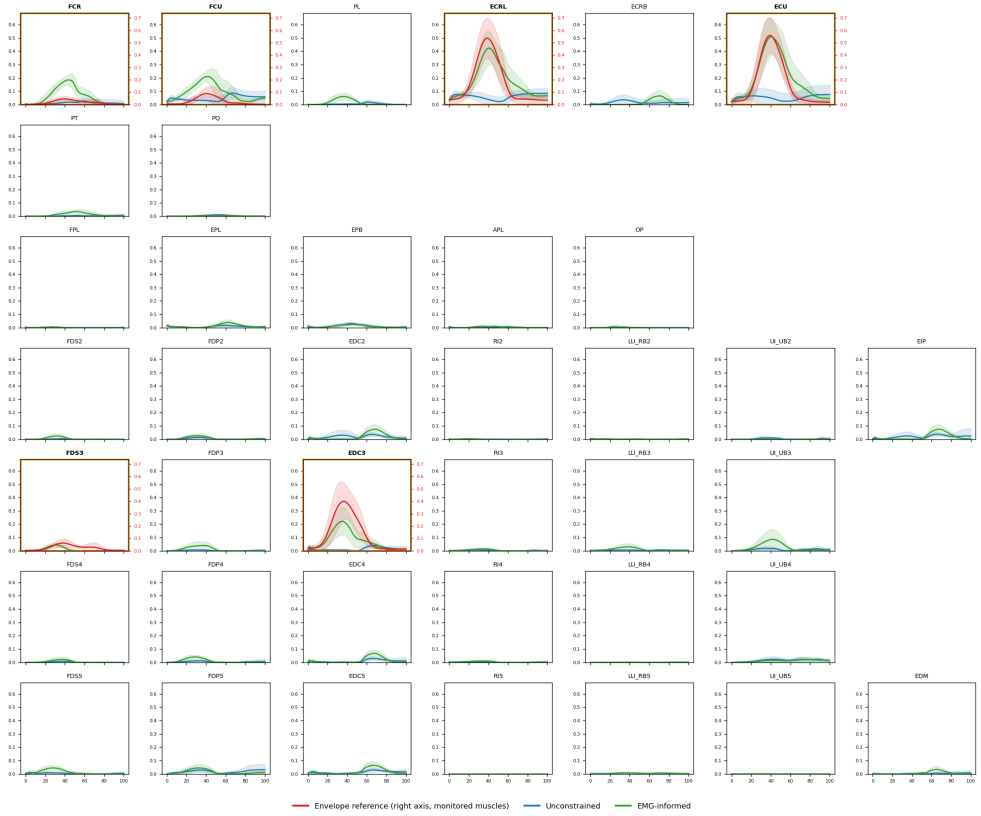


Figure B.4: The same for wrist extension. Axes, colours and bands as in Fig. B.3.

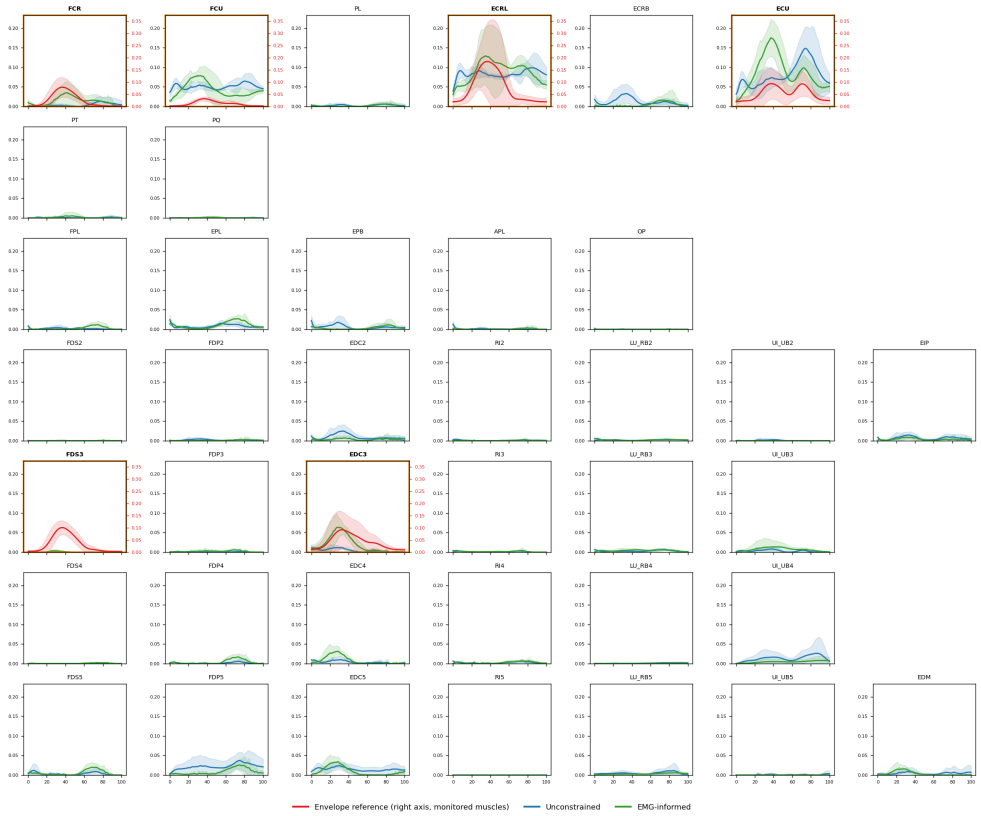


Figure B.5: The same for radial deviation. Axes, colours and bands as in Fig. B.3.

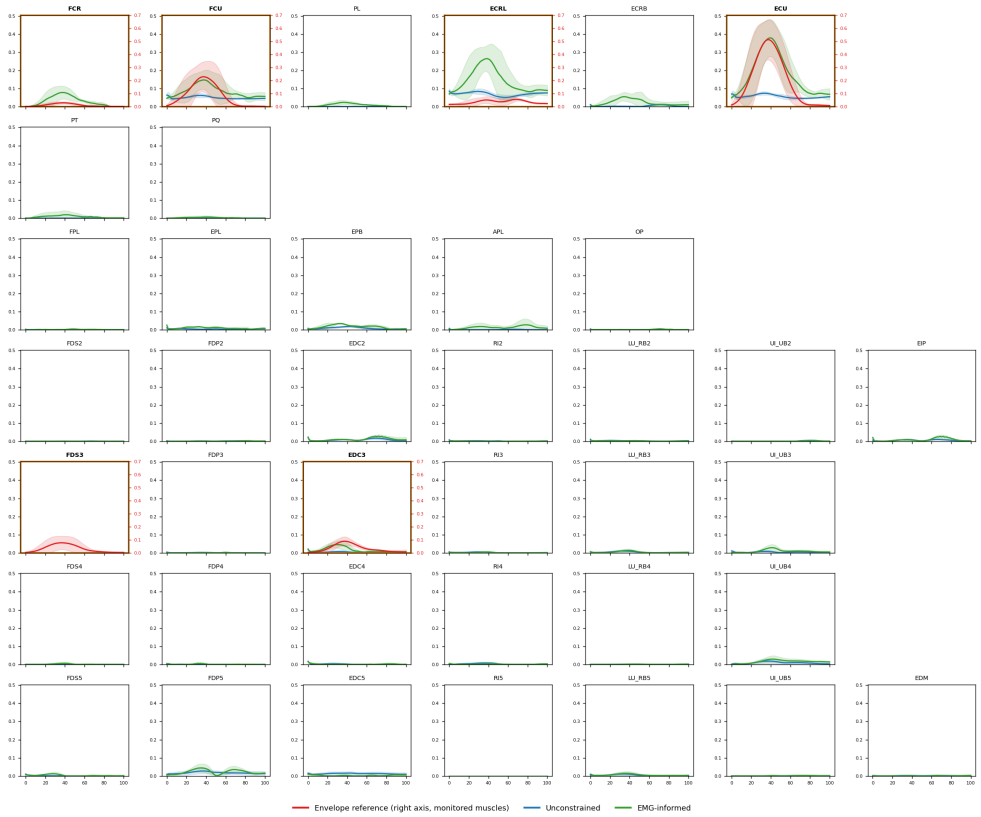


Figure B.6: The same for ulnar deviation. Axes, colours and bands as in Fig. B.3.

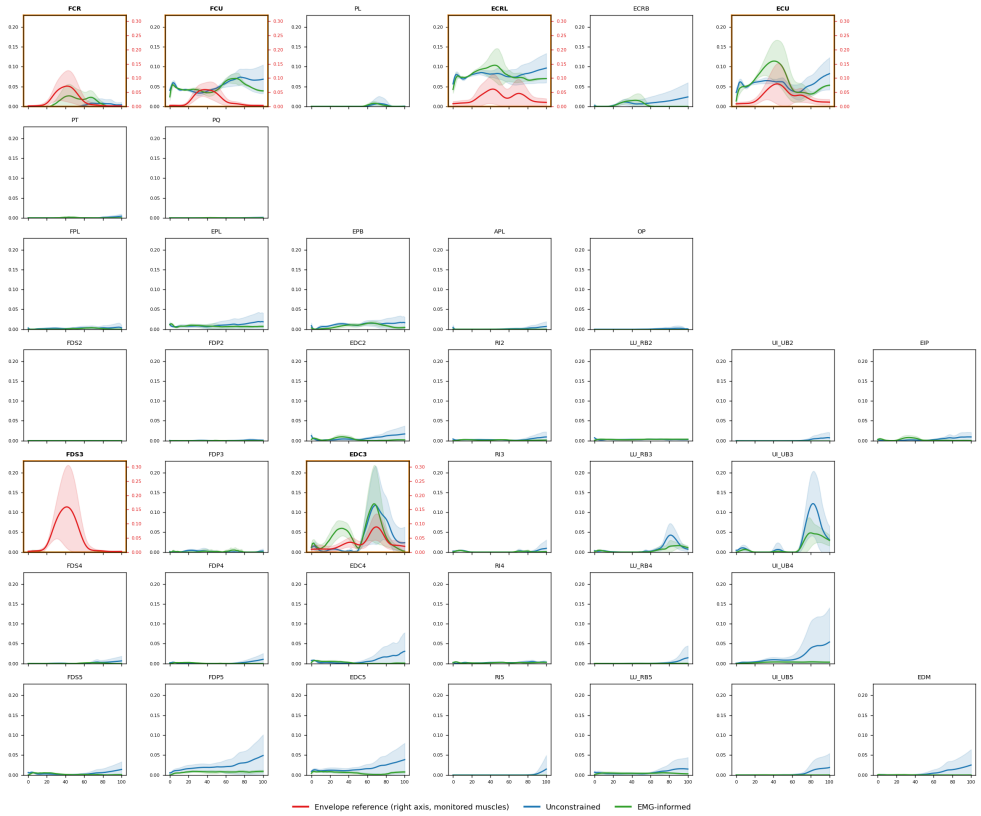


Figure B.7: The same for finger flexion. Axes, colours and bands as in Fig. B.3.

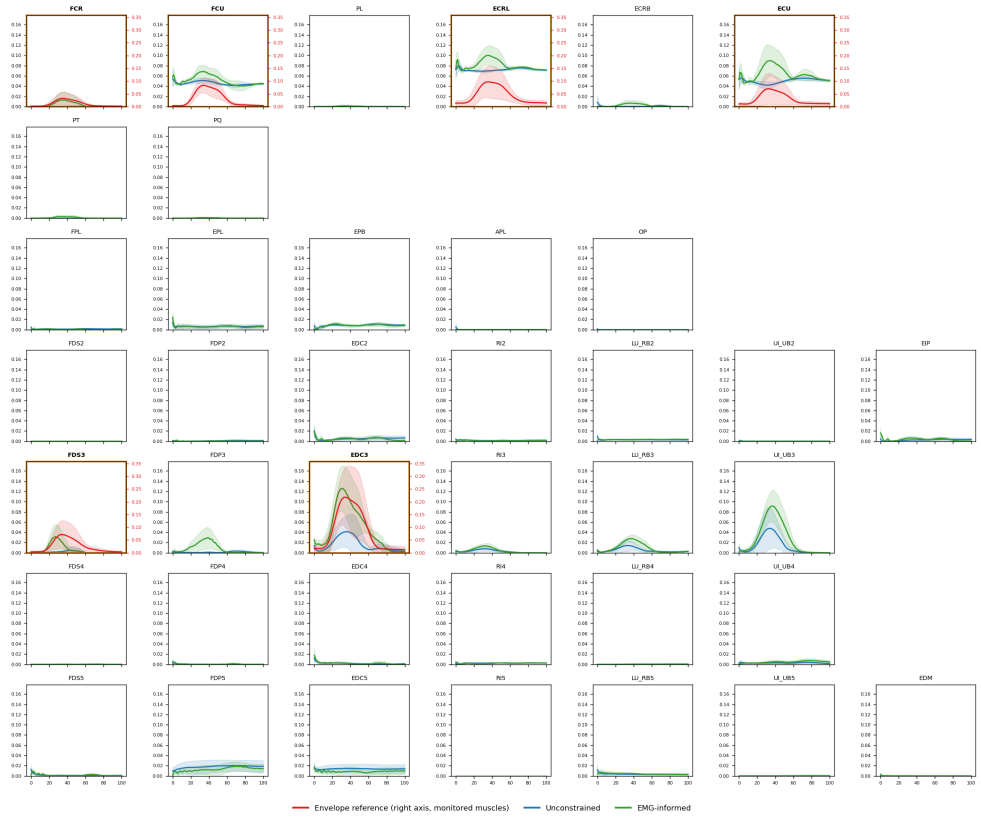


Figure B.8: The same for finger extension. Axes, colours and bands as in Fig. B.3.

Bibliography

[1] Xiao Li, Hong Zeng, Yongqiang Li, and Aiguo Song. "Quantitative assessment via multi-domain fusion of muscle synergy associated with upper-limb motor function for stroke rehabilitation". In: *IEEE Transactions on Biomedical Engineering* 71.5 (May 2024), pp. 1430–1441. DOI: 10.1109/TBME.2023.3339634 (cit. on p. 1).

[2] Aaron Fleming, Nicole Stafford, Stephanie Huang, Xiaogang Hu, Daniel P. Ferris, and He (Helen) Huang. "Myoelectric control of robotic lower limb prostheses: a review of electromyography interfaces, control paradigms, challenges and future directions". In: *Journal of Neural Engineering* 18.4 (July 2021), p. 041004. DOI: 10.1088/1741-2552/ac1176 (cit. on p. 1).

[3] J. M. Muggleton, R. Allen, and P. H. Chappell. "Hand and arm injuries associated with repetitive manual work in industry: a review of disorders, risk factors and preventive measures". In: *Ergonomics* 42.5 (May 1, 1999), pp. 714–739. DOI: 10.1080/001401399185405 (cit. on p. 1).

[4] Craig L. Taylor and Robert J. Schwarz. "The anatomy and mechanics of the human hand". In: *Artificial Limbs* 2 (1955), pp. 22–35 (cit. on pp. 1, 9).

[5] K. S. Türker. "Electromyography: some methodological problems and issues". In: *Physical Therapy* 73.10 (Oct. 1993), pp. 698–710. DOI: 10.1093/ptj/73.10.698 (cit. on pp. 1, 7).

[6] Logan Wade, Laurie Needham, Polly McGuigan, and James Bilzon. "Applications and limitations of current markerless motion capture methods for clinical gait biomechanics". In: *PeerJ* 10 (Feb. 25, 2022), e12995. DOI: 10.7717/peerj.12995 (cit. on pp. 1, 6).

[7] D. A. Winter, A. J. Fuglevand, and S. E. Archer. "Crosstalk in surface electromyography: theoretical and practical estimates". In: *Journal of Electromyography and Kinesiology* 4.1 (Jan. 1, 1994), pp. 15–26. DOI: 10.1016/1050-6411(94)90023-X (cit. on pp. 2, 7, 21).

[8] Felix E. Zajac. "Muscle and tendon: properties, models, scaling, and application to biomechanics and motor control". In: *Critical Reviews in Biomedical Engineering* 17.4 (1989), pp. 359–411 (cit. on pp. 2, 7, 8, 41).

[9] C. G. Burgar, F. J. Valero-Cuevas, and V. R. Hentz. "Fine-wire electromyographic recording during force generation. Application to index finger kinesiologic studies". In: *American Journal of Physical Medicine & Rehabilitation* 76.6 (1997), pp. 494–501. DOI: 10.1097/00002060-199711000-00012 (cit. on pp. 2, 8).

[10] Angelos Karatsidis, Moonki Jung, H. Martin Schepers, Giovanni Bellusci, Mark de Zee, Peter H. Veltink, and Michael Skipper Andersen. "Musculoskeletal model-based inverse dynamic analysis under ambulatory conditions using inertial motion capture". In: *Medical Engineering & Physics* 65 (Feb. 5, 2019), pp. 68–77. DOI: 10.1016/j.medengphy.2018.12.021 (cit. on pp. 2, 5, 10).

[11] Roy D. Crowninshield and Richard A. Brand. "A physiologically based criterion of muscle force prediction in locomotion". In: *Journal of Biomechanics* 14.11 (Jan. 1, 1981), pp. 793–801. DOI: 10.1016/0021-9290(81)90035-X (cit. on pp. 2, 10).

[12] Claudio Pizzolato, David G. Lloyd, Massimo Sartori, Elena Ceseracciu, Thor F. Besier, Benjamin J. Fregly, and Monica Reggiani. "CEINMS: a toolbox to investigate the influence of different neural control solutions on the prediction of muscle excitation and joint moments during dynamic motor tasks". In: *Journal of Biomechanics* 48.14 (Nov. 5, 2015), pp. 3929–3936. DOI: 10.1016/j.jbiomech.2015.09.021 (cit. on pp. 2, 12, 13, 58).

[13] Ajay Seth et al. "OpenSim: simulating musculoskeletal dynamics and neuromuscular control to study human and animal movement". In: *PLOS Computational Biology* 14.7 (2018), e1006223. DOI: 10.1371/journal.pcbi.1006223 (cit. on pp. 3, 9).

[14] Eline van der Kruk and Marco M. Reijne. "Accuracy of human motion capture systems for sport applications; state-of-the-art review". In: *European Journal of Sport Science* 18.6 (July 3, 2018), pp. 806–819. DOI: 10.1080/17461391.2018.1463397 (cit. on p. 5).

[15] Valentin Maggioni, Christine Azevedo-Coste, Sam Durand, and François Bailly. "Optimisation and comparison of markerless and marker-based motion capture methods for hand and finger movement analysis". In: *Sensors* 25.4 (Jan. 2025), p. 1079. DOI: 10.3390/s25041079 (cit. on p. 5).

[16] Laura Dipietro, Angelo M. Sabatini, and Paolo Dario. "A survey of glove-based systems and their applications". In: *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)* 38.4 (July 2008), pp. 461–482. DOI: 10.1109/TSMCC.2008.923862 (cit. on p. 5).

[17] Verónica Gracia-Ibáñez, Margarita Vergara, James H. Buffi, Wendy M. Murray, and Joaquín L. Sancho-Bru. "Across-subject calibration of an instrumented glove to measure hand movement for clinical purposes". In: *Computer Methods in Biomechanics and Biomedical Engineering* 20.6 (Apr. 26, 2017), pp. 587–597. DOI: 10.1080/10255842.2016.1265950 (cit. on p. 5).

[18] Zhe Cao, Gines Hidalgo, Tomas Simon, Shih-En Wei, and Yaser Sheikh. "OpenPose: realtime multi-person 2D pose estimation using part affinity fields". In: *IEEE Transactions on Pattern Analysis and Machine Intelligence* 43.1 (Jan. 2021), pp. 172–186. DOI: 10.1109/TPAMI.2019.2929257 (cit. on p. 6).

[19] Fan Zhang, Valentin Bazarevsky, Andrey Vakunov, Andrei Tkachenka, George Sung, Chuo-Ling Chang, and Matthias Grundmann. *MediaPipe Hands: on-device real-time hand tracking*. June 18, 2020. DOI: 10.48550/arXiv.2006.10214 (cit. on pp. 6, 23).

[20] Theocharis Chatzis, Andreas Stergioulas, Dimitrios Konstantinidis, Kosmas Dimitropoulos, and Petros Daras. "A comprehensive study on deep learning-based 3D hand pose estimation methods". In: *Applied Sciences* 10.19 (Jan. 2020), p. 6850. DOI: 10.3390/app10196850 (cit. on p. 6).

[21] Dario Farina, Roberto Merletti, and Roger M. Enoka. "The extraction of neural strategies from the surface EMG". In: *Journal of Applied Physiology* 96.4 (Apr. 2004), pp. 1486–1495. DOI: 10.1152/japplphysiol.01070.2003 (cit. on p. 6).

[22] R. Merletti and G. L. Cerone. "Tutorial. Surface EMG detection, conditioning and pre-processing: best practices". In: *Journal of Electromyography and Kinesiology* 54 (Oct. 1, 2020), p. 102440. DOI: 10.1016/j.jelekin.2020.102440 (cit. on pp. 6, 20, 21, 25).

[23] Manuela Besomi et al. "Consensus for experimental design in electromyography (CEDE) project: amplitude normalization matrix". In: *Journal of Electromyography and Kinesiology* 53 (Aug. 1, 2020), p. 102438. DOI: 10.1016/j.jelekin.2020.102438 (cit. on pp. 6, 19, 27).

[24] Kevin G. Keenan, Dario Farina, Katrina S. Maluf, Roberto Merletti, and Roger M. Enoka. "Influence of amplitude cancellation on the simulated surface electromyogram". In: *Journal of Applied Physiology* 98.1 (Jan. 2005), pp. 120–131. DOI: 10.1152/japplphysiol.00894.2004 (cit. on p. 7).

[25] Roberto Merletti, Aleš Holobar, and Dario Farina. "Analysis of motor units with high-density surface electromyography". In: *Journal of Electromyography and Kinesiology* 18.6 (Dec. 1, 2008), pp. 879–890. DOI: 10.1016/j.jelekin.2008.09.002 (cit. on p. 7).

[26] Nicholas Tacca, Collin Dunlap, Sean P. Donegan, James O. Hardin, Eric Meyers, Michael J. Darrow, Samuel Colachis IV, Andrew Gillman, and David A. Friedenberg. "Wearable high-density EMG sleeve for complex hand gesture classification and continuous joint angle estimation". In: *Scientific Reports* 14.1 (Aug. 9, 2024), p. 18564. DOI: 10.1038/s41598-024-64458-x (cit. on p. 7).

[27] Matthew C. Tresch, Vincent C. K. Cheung, and Andrea d'Avella. "Matrix factorization algorithms for the identification of muscle synergies: evaluation on simulated and experimental data sets". In: *Journal of Neurophysiology* 95.4 (Apr. 2006), pp. 2199–2212. DOI: 10.1152/jn.00222.2005 (cit. on p. 7).

[28] Ahmet Erdemir, Scott McLean, Walter Herzog, and Antonie J. van den Bogert. "Model-based estimation of muscle forces exerted during movements". In: *Clinical Biomechanics* 22.2 (Feb. 1, 2007), pp. 131–154. DOI: 10.1016/j.clinbiomech.2006.09.005 (cit. on p. 8).

[29] Vittorio Caggiano, Huawei Wang, Guillaume Durandau, Massimo Sartori, and Vikash Kumar. *MyoSuite – a contact-rich simulation suite for musculoskeletal motor control*. May 26, 2022. DOI: 10.48550/arXiv.2205.13600 (cit. on pp. 8, 9, 14, 40, 103).

[30] Michael Damsgaard, John Rasmussen, Søren Tørholm Christensen, Egidijus Surma, and Mark de Zee. "Analysis of musculoskeletal systems in the AnyBody Modeling System". In: *Simulation Modelling Practice and Theory* 14.8 (Nov. 1, 2006), pp. 1100–1111. DOI: 10.1016/j.simpat.2006.09.001 (cit. on p. 9).

[31] Emanuel Todorov, Tom Erez, and Yuval Tassa. "MuJoCo: a physics engine for model-based control". In: *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, Oct. 2012, pp. 5026–5033. DOI: 10.1109/IROS.2012.6386109 (cit. on pp. 9, 40, 103).

[32] K. N. An, Y. Ueba, E. Y. Chao, W. P. Cooney, and R. L. Linscheid. "Tendon excursion and moment arm of index finger muscles". In: *Journal of Biomechanics* 16.6 (Jan. 1, 1983), pp. 419–425. DOI: 10.1016/0021-9290(83)90074-X (cit. on p. 9).

[33] Agneta Gustus, Georg Stillfried, Judith Visser, Henrik Jörntell, and Patrick van der Smagt. "Human hand modelling: kinematics, dynamics, applications". In: *Biological Cybernetics* 106.11 (Dec. 1, 2012), pp. 741–755. DOI: 10.1007/s00422-012-0532-4 (cit. on p. 10).

[34] Florian Michaud, Mario Lamas, Urbano Lugrís, and Javier Cuadrado. "A fair and EMG-validated comparison of recruitment criteria, musculotendon models and muscle coordination strategies, for the inverse-dynamics based optimization of muscle forces during gait". In: *Journal of NeuroEngineering and Rehabilitation* 18.1 (Jan. 28, 2021), p. 17. DOI: 10.1186/s12984-021-00806-6 (cit. on p. 11).

[35] Darryl G. Thelen and Frank C. Anderson. "Using computed muscle control to generate forward dynamic simulations of human walking from experimental data". In: *Journal of Biomechanics* 39.6 (Jan. 1, 2006), pp. 1107–1115. DOI: 10.1016/j.jbiomech.2005.02.010 (cit. on p. 11).

[36] Francisco J. Valero-Cuevas, Heiko Hoffmann, Manish U. Kurse, Jason J. Kutch, and Evangelos A. Theodorou. "Computational models for neuromuscular function". In: *IEEE Reviews in Biomedical Engineering* 2 (2009), pp. 110–135. DOI: 10.1109/RBME.2009.2034981 (cit. on p. 11).

[37] David G. Lloyd and Thor F. Besier. "An EMG-driven musculoskeletal model to estimate muscle forces and knee joint moments in vivo". In: *Journal of Biomechanics* 36.6 (June 1, 2003), pp. 765–776. DOI: 10.1016/S0021-9290(03)00010-1 (cit. on pp. 12, 13, 26, 49, 53, 58).

[38] Di Ao, Marleny M. Vega, Mohammad S. Shourijeh, Carolynn Patten, and Benjamin J. Fregly. "EMG-driven musculoskeletal model calibration with estimation of unmeasured muscle excitations via synergy extrapolation". In: *Frontiers in Bioengineering and Biotechnology* 10 (Sept. 7, 2022), p. 962959. DOI: 10.3389/fbioe.2022.962959 (cit. on pp. 12, 13).

[39] Hermie J. Hermens, Bart Freriks, Catherine Disselhorst-Klug, and Günter Rau. "Development of recommendations for SEMG sensors and sensor placement procedures". In: *Journal of Electromyography and Kinesiology* 10.5 (Oct. 1, 2000), pp. 361–374. DOI: 10.1016/S1050-6411(00)00027-4 (cit. on p. 21).

[40] Paul H. C. Eilers and Hans F. M. Boelens. *Baseline correction with asymmetric least squares smoothing*. Leiden University Medical Centre, Oct. 21, 2005 (cit. on p. 29).

[41] D. Farina and R. Merletti. "A novel approach for estimating muscle fiber conduction velocity by spatial and temporal filtering of surface EMG signals". In: *IEEE Transactions on Biomedical Engineering* 50.12 (Dec. 2003), pp. 1340–1351. DOI: 10.1109/TBME.2003.819847 (cit. on p. 33).

[42] Julius O. Smith. *Linear interpolation frequency response*. 2010. URL: https://ccrma.stanford.edu/~jos/pasp/ (cit. on p. 33).

[43] Davis A. Forman, Garrick N. Forman, Edwin J. Avila-Mireles, Maddalena Mugnosso, Jacopo Zenzeri, Bernadette Murphy, and Michael W. R. Holmes. "Characterizing forearm muscle activity in university-aged males during dynamic radial-ulnar deviation of the wrist using a wrist robot". In: *Journal of Biomechanics* 108 (July 17, 2020), p. 109897. DOI: 10.1016/j.jbiomech.2020.109897 (cit. on p. 35).

[44] Leon Alexander. "Practical applications of tenodesis in hand surgery". In: *Journal of Musculoskeletal Surgery and Research* 5.4 (Nov. 13, 2021), pp. 312–314. DOI: 10.25259/JMSR_67_2021 (cit. on p. 37).

[45] M. H. Schieber. "Muscular production of individuated finger movements: the roles of extrinsic finger muscles". In: *Journal of Neuroscience* 15.1 (Jan. 1, 1995), pp. 284–297. DOI: 10.1523/JNEUROSCI.15-01-00284.1995 (cit. on p. 37).

[46] J. N. A. L. Leijnse, P. M. Quesada, and C. W. Spoor. "Kinematic evaluation of the finger's interphalangeal joints coupling mechanism—variability, flexion–extension differences, triggers, locking swanneck deformities, anthropometric correlations". In: *Journal of Biomechanics* 43.12 (Aug. 26, 2010), pp. 2381–2393. DOI: 10.1016/j.jbiomech.2010.04.021 (cit. on p. 38).

[47] M. Mentzel, A. Benlic, N. J. Wachter, D. Gulkin, S. Bauknecht, and J. Gülke. "[The dynamics of motion sequences of the finger joints during fist closure]". In: *Handchirurgie, Mikrochirurgie, Plastische Chirurgie* 43.3 (June 2011), pp. 147–154. DOI: 10.1055/s-0031-1271804 (cit. on p. 38).

[48] Andrea Prestia. "Muscle-driven solutions for rehabilitation and control". PhD thesis. Politecnico di Torino, Apr. 3, 2025. DOI: 10.13121/polito/porto/2999500 (cit. on pp. 42–44, 94).

[49] Bartolomeo Stellato, Goran Banjac, Paul Goulart, Alberto Bemporad, and Stephen Boyd. "OSQP: an operator splitting solver for quadratic programs". In: *Mathematical Programming Computation* 12.4 (Dec. 1, 2020), pp. 637–672. DOI: 10.1007/s12532-020-00179-2 (cit. on p. 44).

[50] Rainer Storn and Kenneth Price. "Differential evolution – a simple and efficient heuristic for global optimization over continuous spaces". In: *Journal of Global Optimization* 11.4 (Dec. 1, 1997), pp. 341–359. DOI: 10.1023/A:1008202821328 (cit. on p. 45).

[51] Azadeh Kian, Claudio Pizzolato, Mark Halaki, Karen Ginn, David Lloyd, Darren Reed, and David Ackland. "Static optimization underestimates antagonist muscle activity at the glenohumeral joint: a musculoskeletal modeling study". In: *Journal of Biomechanics* 97 (Dec. 3, 2019), p. 109348. DOI: 10.1016/j.jbiomech.2019.109348 (cit. on pp. 48, 98).

[52] William P. Dunlap, Marshall B. Jones, and Alvah C. Bittner. "Average correlations vs. correlated averages". In: *Bulletin of the Psychonomic Society* 21.3 (Mar. 1, 1983), pp. 213–216. DOI: 10.3758/BF03334690 (cit. on p. 51).

[53] Thomas A. Lang and Douglas G. Altman. "Basic statistical reporting for articles published in biomedical journals: the "Statistical Analyses and Methods in the Published Literature" or the SAMPL guidelines". In: *International Journal of Nursing Studies* 52.1 (Jan. 1, 2015), pp. 5–9. DOI: 10.1016/j.ijnurstu.2014.09.006 (cit. on p. 52).

[54] K. Falconer and D. A. Winter. "Quantitative assessment of co-contraction at the ankle joint in walking". In: *Electromyography and Clinical Neurophysiology* 25.2 (1985), pp. 135–149 (cit. on p. 61).

[55] Claire V. Hammond et al. "The Neuromusculoskeletal Modeling Pipeline: MATLAB-based model personalization and treatment optimization functionality for OpenSim". In: *Journal of NeuroEngineering and Rehabilitation* 22.1 (May 19, 2025), p. 112. DOI: 10.1186/s12984-025-01629-5 (cit. on p. 102).

[56] American Academy of Orthopaedic Surgeons. *Joint motion: method of measuring and recording*. Chicago: American Academy of Orthopaedic Surgeons, 1965 (cit. on p. 108).

[57] D. C. Boone and S. P. Azen. "Normal range of motion of joints in male subjects". In: *The Journal of Bone & Joint Surgery* 61.5 (1979), pp. 756–759. DOI: 10.2106/00004623-197961050-00017 (cit. on p. 108).

[58] Joyce White, Lisa Coppola, Andrea Skomurski, and Emma Mitchell-Rekrut. "Influence of age and gender on normative passive range of motion values of the carpometacarpal joint of the thumb". In: *Journal of Hand Therapy* 31.3 (July 1, 2018), pp. 390–397. DOI: 10.1016/j.jht.2017.02.011 (cit. on p. 108).

[59] Tatsuki Ebata, Ikuo Nakai, Akira Kogure, and Kenichiro Goto. "Normal range of motion of thumb metacarpophalangeal joint". In: *HAND* 11.1 (Sept. 1, 2016), 78S–79S. DOI: 10.1177/1558944716660555eq (cit. on p. 108).

[60] Mary C. Hume, Harris Gellman, Harry McKellop, and Robert H. Brumfield. "Functional range of motion of the joints of the hand". In: *The Journal of Hand Surgery* 15.2 (Mar. 1, 1990), pp. 240–243. DOI: 10.1016/0363-5023(90)90102-W (cit. on p. 108).