UC Berkeley

UC Berkeley Electronic Theses and Dissertations

Title

A Systematic Study of the Muscle Force–Deformation Relationship at the Human Elbow: Toward Physiology-Aware Assistive Device Control and Noninvasive Muscle Force Sensing

Permalink

https://escholarship.org/uc/item/7dn0v17d

Author

Hallock, Laura

Publication Date

2021

Peer reviewed|Thesis/dissertation

eScholarship.org Powered by the California Digital Library
University of California

A Systematic Study of the Muscle Force–Deformation Relationship at the Human Elbow: Toward Physiology-Aware Assistive Device Control and Noninvasive Muscle Force Sensing

by

Laura Hallock

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Engineering — Electrical Engineering and Computer Sciences

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Ruzena Bajcsy, Chair
Professor Hannah Stuart
Professor Robert Full
Professor Neville Hogan
Professor Claire Tomlin

Summer 2021

A Systematic Study of the Muscle Force–Deformation Relationship at the Human Elbow: Toward Physiology-Aware Assistive Device Control and Noninvasive Muscle Force Sensing

Copyright 2021
by
Laura Hallock

Abstract

A Systematic Study of the Muscle Force–Deformation Relationship at the Human Elbow: Toward Physiology-Aware Assistive Device Control and Noninvasive Muscle Force Sensing

by

Laura Hallock

Doctor of Philosophy in Engineering — Electrical Engineering and Computer Sciences

University of California, Berkeley

Professor Ruzena Bajcsy, Chair

While there exist a number of mechanically sophisticated exoskeletons, prostheses, and assistive robots, with articulations similar to those of the intact human arm and hand, these devices remain limited in their ability to augment human dexterity and safely interact with human users and collaborators. In particular, due to the limits of conventional sensing, robots remain locked in industrial cages, prosthesis users can often modulate only a single degree of freedom, and when human–device interactions do occur, we have almost no understanding of the resulting physiological impacts on the user's musculoskeletal system.

This thesis begins to address these dual problems of device capability and safety by leveraging a novel signal class — muscle deformation, as measured via ultrasound — to probe individual muscle forces, which cannot in general be measured noninvasively but are key to understanding musculoskeletal dynamics during dexterous motion. Using the elbow flexion motion as a case study, this document describes first, preliminary characterization of this deformation and its relationship to output force; second, proof-of-concept use of this signal for a trajectory tracking task; third, collection of a comprehensive 3D deformation data set to systematically examine the force–deformation relationship; fourth, initial geometric analysis of this data set and a proposed framework for future modeling; and lastly, a discussion of the open-source release of code and data associated with this project, preliminary work on alternative sensing modalities, and speculation on future applications in robotics and medicine. This dissertation work — collectively released as the OpenArm project for collaborative research use at https://simtk.org/projects/openarm — paves the way for future investigation on both the extraction of multiple independent signals for high-dimensional device control and enhanced overall understanding of the joint human–machine dynamical system.

To my father, my mother, my sister, and my love.

Contents

Contents ii

List of Figures v

List of Tables xv

1 Introduction 1

2 Motivation & Related Work 4
2.1 The Case for Improved Human Muscle Force Modeling . . . 4
2.2 Our Focus: Net Torque at the Human Elbow . . . 7
2.3 Existing Musculoskeletal Modeling Frameworks . . . 8
2.4 Existing (Noninvasive) Musculoskeletal Sensing Paradigms . . . 12
2.5 Deformation as a Measure of Musculoskeletal Forces . . . 13
2.6 Ultrasound as a Sensing & Control Modality . . . 15
2.7 Summary . . . 16

I Exploratory Data Analysis 17

3 Exploratory Force–Deformation Correlation 18
3.1 Study Scope . . . 19
3.2 Data Set Collection . . . 19
3.3 Data Processing & Definition of Deformation Signals . . . 21
3.4 Correlation & Evaluation . . . 23
3.5 Summary & Limitations . . . 28

4 Refined Force–Deformation Correlation 29
4.1 Study Scope . . . 30
4.2 Real-Time Deformation & Activation Tracking Platform . . . 30
4.3 Data Set Collection . . . 34
4.4 Correlation & Evaluation . . . 35
4.5 Summary & Limitations . . . 39

II Tracking & Control **41**

5 Trajectory Following & Control **42**
5.1 Study Scope . . . 42
5.2 Data Collection . . . 43
5.3 Quantitative Tracking Performance . . . 43
5.4 Quantitative & Qualitative Tracking Preferences . . . 45
5.5 Summary & Limitations . . . 48

6 Expanded Deformation Tracking **49**
6.1 Study Scope . . . 49
6.2 Candidate Tracking Algorithms . . . 50
6.3 Tracking Error Metrics . . . 52
6.4 Tracking Performance . . . 52
6.5 Summary & Limitations . . . 54

III Comprehensive Data Set **58**

7 3D Data Set Collection **59**
7.1 Data Scope . . . 60
7.2 3D Imaging Modality: Localized 2D Ultrasound . . . 60
7.3 Prototype Data Set (OpenArm 1.0) . . . 61
7.4 Refined, Autosegmented, Multi-Subject Data Set (OpenArm 2.0) . . . 65
7.5 Summary & Limitations . . . 69

8 Automated Tissue Segmentation **71**
8.1 Segmentation Targets & Scope . . . 71
8.2 Segmentation Approach: Convolutional Neural Networks . . . 72
8.3 Candidate Segmentation Architectures & Modifications . . . 72
8.4 Ground-Truth Data Set Generation . . . 74
8.5 Segmentation Performance . . . 74
8.6 Summary & Limitations . . . 79

IV Systematic Modeling **80**

9 Exploratory Geometric Analysis **81**
9.1 Study Scope . . . 81
9.2 Definition of Deformation Signals . . . 82
9.3 Spatial Deformation Analysis . . . 82
9.4 Summary & Limitations . . . 88

10 Systematic Force–Deformation Modeling **89**
10.1 Model Scope . . . 89
10.2 Example Modeling Framework . . . 90
10.3 Thoughts on Model Selection . . . 95
10.4 Summary & Limitations . . . 95

V Data, Codebases, Alternative Sensing, & Conclusions **96**

11 The OpenArm Codebase **97**
11.1 OpenArm Multisensor 1.0 . . . 97
11.2 OpenArm Multisensor 2.0 . . . 97
11.3 OpenArm 1.0 . . . 98
11.4 OpenArm 2.0 . . . 98
11.5 Annotation Source Code . . . 99

12 Alternative Sensing: Force–Vibration Correlation via Acoustic Myography **100**
12.1 Study Scope . . . 100
12.2 Existing Muscle Force–Vibration Measurement & Models . . . 101
12.3 Proof-of-Concept Model . . . 102
12.4 Summary & Limitations . . . 103

13 Conclusions & Future Directions **105**
13.1 Safe & Capable Device Control . . . 105
13.2 Neuromuscular Dynamics Modeling . . . 108
13.3 Quantifying, Tracking, & Rehabilitating Pathology . . . 110
13.4 Additional Fields of Interest . . . 111
13.5 Final Thoughts . . . 111

Bibliography **112**

List of Figures

1.1 Human motion results from complex neuromuscular contraction dynamics that convert a neurological activation signal into muscle output force. This process is poorly understood and difficult to measure, greatly limiting our ability to replicate such motions in robots and assistive devices from neurological sensor data like surface electromyography (sEMG, *red*). This dissertation work (*green*), in contrast, infers motion dynamics more directly by measuring two classes of signals intrinsically coupled to force output: deformation (via ultrasound) and vibration (via acoustic myography, or AMG). These signals can be used directly to control safe and dexterous assistive devices without modeling these complex neuromuscular contraction dynamics, and the resulting models can yield new insights into healthy and pathological human motion. Select icons made by hunotika from the Noun Project and smalllikeart from Flaticon. 2

2.1 Musculoskeletal anatomy of the human arm. *Left*: The human arm is complex, with tens of muscles actuating multi-degree-of-freedom (DoF) joints in synergistic and antagonistic ways. *Right*: The elbow provides a comparatively simple modeling testbed, with only three flexors (the biceps brachii, brachialis, and brachioradialis) and a single extensor (the triceps brachii) actuating a single-DoF joint. Left and center images from [42]; right image adapted from [154]. 7

2.2 Basic Hill-type muscle model [59], a ubiquitous model of muscle force–length and force–velocity relationships, including nonlinear contractile element CE, series elastic element SE, and parallel elastic element PE. Image from [54]. 11

2.3 Toy example illustrating limitations of musculoskeletal simulation frameworks, depicting a simple sagittal model of the elbow with one biceps-like flexor exerting force F_{biceps} and one triceps-like extensor exerting force F_{triceps}. Both left and right images depict muscle forces consistent with observed output force F_{ext}; in the absence of muscle force measurement, simulation frameworks are forced to choose a solution. Most frameworks will likely choose the solution depicted on the left if optimizing for minimal exertion, though the image on the right is likely more consistent with stabilized, higher-impedance motion. Select images adapted from [67]. 12

3.1 Experimental setup for the collection of time series force, surface electromyography (sEMG), and ultrasound data under multiple elbow angles. Setup includes sEMG electrodes (a); ultrasound probe (b); UR5 robot (c) with attached handle, through which subject may transmit force to force plate (d); and wrist brace (e) for elbow isolation. ©2020 IEEE. 20
3.2 Example cross section of the brachioradialis muscle, as collected via ultrasound (probe to left of frame, scanning right), with illustrated shape measures of interest: cross-sectional area (*CSA*, *cyan*), thickness (*T*, *blue*) and aspect ratio (*AR*, *dark blue*, the ratio of thickness T to width W). Each measure is calculated for each ultrasound frame during trials of isometric elbow flexion to generate time series deformation data. 22
3.3 *Top*: Correlation of cross-sectional area (*CSA*, *cyan*), thickness (*T*, *blue*) and aspect ratio (*AR*, *dark blue*) of the brachioradialis cross section, both raw (*solid*) and detrended (*dashed*), with elbow output force, alongside baseline force correlations with biceps (*light orange*) and brachioradialis (*dark orange*) surface electromyography (*sEMG*) data. Correlations were computed across flexion angles ranging from near full extension (25°) to near maximum flexion (97°). *CSA*, *T*, and *AR* deformation signals, especially when detrended, correlate well with elbow output force for most elbow angles, but this correlation collapses near full flexion. *Bottom*: Example ultrasound frames with annotated brachioradialis contours depicting no force (*top row*) and high output force (*bottom row*) at each examined flexion angle reflect the changing presentation of muscle deformation with changes in elbow angle. ©2020 IEEE. 26
3.4 *Top*: Correlation of cross-sectional area (*CSA*, *cyan*), thickness (*T*, *blue*) and aspect ratio (*AR*, *dark blue*) of the brachioradialis cross section, both raw (*solid*) and detrended (*hashed*), with elbow output force, across various subjects, collected at ∼69° elbow flexion. *CSA*, *T*, and *AR* deformation signals, especially when detrended, correlate well with output force, though the magnitude and direction of these correlations varies widely across subjects. *Bottom*: Example ultrasound frames with annotated brachioradialis contours depicting no force (*top row*) and high output force (*bottom row*) for each subject reflect substantial morphological variation. ©2020 IEEE. 27

4.1 Experimental setup for the collection of time series force, ultrasound, and surface electromyography (sEMG) data during constrained isometric elbow flexion. Setup includes ultrasound probe (a) attached securely to user's forearm with cuff (b); sEMG electrodes (c); wrist brace (d) through which subject transmits force to attached to force-torque sensor (e), in turn held stable by KUKA robot (f); elbow stabilizing jig (g); goal and sensor trajectory display (h) for real time visual feedback for subject self-assessment; and real-time ultrasound thickness tracking data (i) for continuous experimenter system status monitoring. This system allows subjects to precisely follow a specified force trajectory to enable study of force–deformation correlation under varied trajectory types (section 4.3) and to perform trajectory tracking tasks using experimental deformation- and activation-based signals (chapter 5). ©2021 IEEE. 31
4.2 Still frame of optical-flow-based brachioradialis thickness tracking system. Points were tracked along the superficial (*red*) and deep (*blue*) fascial surfaces of the brachioradialis, and thickness was reported as the vertical (superficial-to-deep, *green*-to-*green*) distance between the center (mean) of each cluster. A line connecting each cluster center was also displayed (*magenta*) to allow for easy observation of particle drift. ©2021 IEEE. 33
4.3 Example time series data collected from subject *Sub1* for force–deformation and force–activation correlation analysis, including output force (*black*, *solid*), alongside specified goal trajectory (*black*, *dotted*), deformation (i.e., brachioradialis thickness change as tracked via optical flow, *blue*), and activation (as measured via sEMG, *orange*). Subjects were able to track the specified force trajectory with little error, enabling controlled observation of a variety of sustained and quickly varying force conditions, and both deformation and activation were shown to be highly correlated with output force during all portions of the trajectory. Signal values are reported — as they were displayed — as a fraction of measured maximal value, as described in section 4.2. ©2021 IEEE. 36
4.4 Correlation of muscle deformation (*blue*) and activation (*orange*) signals with elbow output force across all subjects. Despite substantial differences in morphology (illustrated below in Fig. 4.5), most subjects — aside from *Sub7* and *Sub9*, who displayed morphological quirks that resulted in poor signal quality, as discussed in section 4.4 — showed strong correlation between deformation and output force. ©2021 IEEE. 37
4.5 Example ultrasound frames from an illustrative subset of subjects, with tracked and annotated brachioradialis thickness, depicting no force (*top row*) and high output force (*bottom row*) for each subject. While subjects' morphology varies significantly, most subjects (like the pictured *Sub1*, *Sub6*, *Sub8*, and *Sub10*) display a reliable thickness increase with output force, while several (like *Sub7*) primarily display lateral motion, leaving thickness uncorrelated with force output. ©2021 IEEE. 38

4.6 Correlation of muscle deformation (*blue*) and activation (*orange*) signals with elbow output force across various trajectory types and in aggregate, with noted standard deviation across subjects. Deformation remained strongly correlated with output force for all examined trajectory types, with slightly lower and more variable correlation during *sustained* and *sine* trajectories that is likely the result of limitations in the optical flow tracking system, as discussed in section 4.4. ©2021 IEEE. 39

5.1 Example time series data collected from subject *Sub1* during separate deformation (*blue*, *top*) and activation (*orange*, *bottom*) tracking trials, alongside specified goal trajectories (*black*, *dotted*). Most subjects were able to complete the tracking task qualitatively well using both signal traces, despite drift in the deformation tracking system that sometimes made returning the system to baseline difficult (as is evident to the right of this trajectory), and uncontrolled, high-frequency oscillations in the activation signal. ©2021 IEEE. 44

5.2 Tracking error during muscle deformation (*blue*) and activation (*orange*) trajectory tracking tasks across all subjects. With the exception of *Sub7* — whose irregular anatomy prevented almost any deformation-based control — subjects were largely able to achieve better, or at worst comparable, performance when performing deformation-based control, as compared with our baseline activation-based system. ©2021 IEEE. 45

5.3 Tracking error during muscle deformation (*blue*) and activation (*orange*) trajectory tracking tasks across various trajectory types and in aggregate, with noted standard deviation across subjects. During all but *sine* trajectories — which we again theorize were impacted by the drift in our deformation tracking system — subjects were able to consistently achieve lower tracking error using the deformation signal. ©2021 IEEE. 46

5.4 Subject preferences when asked to evaluate deformation- and activation-based tracking tasks separately (*top*) and in head-to-head comparisons (*bottom*). Subjects largely found the deformation-based tracking task easier and perceived the deformation signal to better match their output force, though they rated the activation-based tracker as more responsive, and most preferred the deformation-based tracker overall. Full survey questions and responses are included with the open-source data release. ©2021 IEEE. 47

6.1 Example ultrasound frames from *Sub1* 69° data series and their respective tracked contour points using each candidate algorithm, at (relaxed) initialization (*top*) and approximately 91 s into tracking during force exertion (*bottom*). *Center left*: Complete set of contour points (*red*) tracked via naive Lucas–Kanade (*LK*) generally describes ground truth muscle shape (*GT*, *left*), but fails to track more extreme deviations and exhibits significant drift error. *Center*: Refined set of contour points based on feature quality (*red*) and tracked via Lucas–Kanade (*FRLK*) describe muscle shape only mildly more accurately than naive *LK* approach. *Center right*: Refined contour points tracked after aggressive (*red*) and less aggressive (*yellow*) bilateral filtering (*BFLK*) experience slightly less drift error, but still fail to capture deviation in the top right quadrant of the contour where fascia is narrow. *Right*: Combining Lucas–Kanade-tracked refined feature points (*red*) with contour points (*green, closed*) predicted based on supporter point locations (*green, open*) (*SBLK*) results in reasonable tracking of the full muscle contour as long as supporter point motion is well-correlated with contour motion. ©2020 IEEE. 51

6.2 Mean tracking error and standard deviation of *LK* (*pink*), *FRLK* (*brown*), *BFLK* (*blue*), and *SBLK* (*purple*) algorithms, both tuned (*hashed*) and untuned (*solid*), computed as Jaccard distance, across various subjects, collected at ∼69° elbow flexion. Both the top performing algorithm and the level of tracking accuracy achieved, as well as the extent to which algorithm tuning mitigated errors, varied substantially by subject, likely due to variations in both morphology and motion qualia. Note that for *Sub1*, whose tuned values formed the generic baseline for all tracking, tuned and untuned tracking algorithms are the same. ©2020 IEEE. 55

6.3 Overall tracking efficacy of *SBLK-T* algorithm for all aggregated time series, expressed as fractional *CSA* (*cyan*), *T* (*blue*), and *AR* (*dark blue*) error and Jaccard distance (*JD*, *red*). Histogram weights indicate that thickness *T* can be tracked most reliably, as compared with *CSA* and *AR*, making it the most promising control signal using our examined tracking algorithms. 56

6.4 Example *CSA* (*cyan*), *T* (*blue*), and *AR* (*dark blue*) trajectories, both ground truth (*solid*) and tracked via *SBLK-T* algorithm (*dashed*), alongside time series Jaccard distance error (*JD*, *red*), for *Sub3* data series. For most data series, thickness *T* is tracked most reliably, while aspect ratio *AR* (and, to a lesser extent, cross-sectional area *CSA*) drift substantially. 57

7.1 Experimental setup for the collection of OpenArm 1.0 full-arm upper-limb morphology data under multiple elbow angles and loading conditions, enabling a factorial study of the sources of muscle deformation (shown here at a 60° angle of elbow flexion under HF (*left*) and FS (*right*) loading conditions). Setup includes ultrasound scanner (a) and probe (b) (with attached active motion capture markers (c) used for spatial tracking); weight bands (d) used to load the elbow flexors (during LF, MF, and HF trials); mechanical jig (e) used to support the elbow (during all trials, *left*) and the forearm (during FS trials, *right*); and the phantom devices (f) required to calibrate the coordinate transformation between the motion capture world frame and the measured ultrasound scans. ©2018 IEEE. 62

7.2 Tissue structures from the OpenArm 1.0 data set release, including the (*magenta*) biceps brachii and surrounding muscles (*turquoise* brachialis, *purple* brachioradialis, and *gray* deltoid) under multiple loading conditions and kinematic configurations, as segmented from volumetric reconstruction of ultrasound data. Locations of the coronal cross-sectional scans shown are noted by lines transecting the associated sagittal scan in each inset. Volumetric changes across both kinematic configurations and loading conditions are readily observable, confirming the necessity of modeling both signal sources when employing muscle deformation as a device control signal and informing our analyses in part IV. ©2018 IEEE. . . . 66

7.3 Experimental setup for the collection of OpenArm 2.0 full-arm upper-limb morphology data under multiple elbow angles and loading conditions (shown here at a 30° angle of elbow flexion under FS (*left*) and P1 (*right*) loading conditions). Setup includes ultrasound probe (a) (with attached active motion capture markers (b) used for spatial tracking); force-torque sensor (c), held statically in place by KUKA LBR iiwa 14 R820 robot (d) and used by the subject via real-time visual feedback (e) to maintain constant force output during loaded trials; mechanical jig (f) used to support the elbow (during all trials) and the forearm (during FS trials, *left*); and real-time ultrasound and motion capture data (g) for continuous system status monitoring. ©2019 IEEE. 68

8.1 Baseline U-Net architecture for segmenting 2D slices of 3D volumetric scans, which are then compiled to generate full 3D segmentation. Architecture is almost exactly that of the original U-Net [124], with only a single additional concatenation block (corresponding to four additional 3x3 convolutions, one additional 2x2 max pooling operation, and one additional "up-convolution"). 73

8.2 Training error for all reported network architectures over 40 epochs, over which networks reliably — if sometimes messily — converged. Accuracy values reported in Table 8.1 and predictions shown in Fig. 8.3 are computed from the minimum-loss epoch network for each respective architecture. ©2019 IEEE. 76

8.3 Exemplar volumetric data, as segmented manually (*Ground Truth*), via optimized classical image registration (RANR), and via neural network, unmodified (U-NET), using elastic deformation data augmentation (U-NET+EA), and using an augmented multi-subject data set (Multi-Subject U-NET+EA). Data used for RANR were ground truth values of (*Sub1*, 30°, FS); data for training the optimized neural networks are those described in Section 8.5 and Table 8.1. Data were trained, tested, and predicted only on the upper part of the arm, above the elbow; raw lower-arm intensity maps are provided for context. Although superficially smooth and well-formed, RANR segmentation poorly localizes biceps and humerus (a), resulting in poor segmentation accuracy; in contrast, neural network methods perform reliably along the middle section of the biceps (b) but segment more poorly near the ends of the structure (c). Adding elastic deformation augmentation data generally helps smooth the data and improve accuracy (d), though many artifacts remain. Note that *Sub2* Multi-Subject U-NET prediction (e) shows a scan used in network training; its high level of accuracy thus represents network memorization, and the scan is presented for completeness only. ©2019 IEEE. 77

9.1 Visualization of deformation measures calculated for spatial geometric analysis at each elbow angle θ, loading condition LC, and location x along the length of the biceps brachii on a single representative coronal cross section. Cross-sectional area $CSA_{\theta,LC}(x)$ was directly extracted from the segmented volumetric biceps volume, and thickness $T_{\theta,LC}(x)$ was computed as the mean of the measure shown about a 1 cm region surrounding the centroid of the computed area. Eccentricity $E_{\theta,LC}(x)$ was computed as the major-to-minor-axis ratio of the best fit ellipse to the cross-sectional spatial data values in the least-squares sense. The data collected suggest that CSA changes are consistently associated with muscle output force, and that thickness measurements T, when combined with a model of eccentricity E, could be used to estimate the CSA signal using cheaper 1D A-mode ultrasound sensors. ©2018 IEEE. 83

9.2 Variation in cross-sectional area *CSA* (*top*), thickness T (*middle*), and eccentricity E (*bottom*) along the length of the biceps brachii from shoulder to elbow of the prototype OpenArm 1.0 subject, under multiple elbow flexion angles with the forearm fully supported (*CSA* and T) and at all tested conditions (E). The location of maximum CSA/T (as measured from the corresponding quadratic regression models shown in overlay) was not shown to vary with angle, but the changes in width of the fitted *CSA* quadratics reflect the compression of the muscle as elbow flexion increases, a preliminary and intuitive insight that suggests that building low-dimensional predictive models of *CSA* change may be possible. The steep increase in eccentricity E near the muscle's ends is reflected in the cross-sectional images shown in Fig. 7.2, and the consistent shape of the eccentricity map across all tested conditions indicates that a spatial eccentricity map could be of use in developing a predictive model of muscle *CSA* from 1D thickness T data. ©2018 IEEE. 85

9.3 Spatial variation of change in biceps brachii cross-sectional area (*CSA*) from that of the fully-supported (FS) volume under low (LF, *dashed*) and high (HF, *solid*) loading conditions. Significant variation is consistently observed in a range centered approximately 2.6 cm distal from the location of maximum absolute *CSA* and is larger under higher loading at each configuration. These observed $\Delta\,CSA$ values thus indicate a candidate location from which to extract a spatially robust assistive device or biomechanical modeling control signal. ©2018 IEEE. 86

10.1 Fully empirical "black box" baseline model mapping elbow angle θ and arbitrarily-parameterized muscle deformation D_{biceps} to elbow torque τ_{ext} through arbitrary function $f_0(\cdot,\cdot)$. This model enables data-driven prediction of output torque but lends little to no insight into underlying physiology. Image adapted from [67]. . 91

10.2 Imposing multi-muscle structure on the baseline model in Fig. 10.1 enables extraction of biceps force F_{biceps} and is more representative of real physiology, but requires knowledge of either the configuration-varying percentage of elbow force attributable to biceps exertion $c(\cdot)$ (*top*) or actual torque contributions of other actuating muscles τ_{brach}, τ_{brachrad}, and τ_{triceps} and mass contribution τ_{mg} (*bottom*). Using the currently available OpenArm 2.0 data set, these values must be assumed from literature, but as data sets expand, they too could be determined from corresponding muscle deformation signals. Image adapted from [67]. 92

10.3 The Fig. 10.2 model above can be further refined by incorporating dynamics of the muscle–tendon unit (MTU), assuming springlike behavior of the tendon with spring constant k. This model class requires additional assumption of both k and kinematically varying MTU length $\ell_{\text{MTU}}(\cdot)$ but can be partially validated through measurement of muscle length ℓ_{m} on the OpenArm 2.0 data set. Image adapted from [67]. 93

10.4 Isometric ellipsoidal muscle kinematics can be imposed on the Fig. 10.3 model above to generate a fully "white box" model requiring little to no parameter fitting. This additional structure enforces that deformation is specifically parameterized as cross-sectional area A_{m} and requires knowledge of (assumed constant) muscle volume V_{m}, which can be extracted from the OpenArm 2.0 data set. Image adapted from [67]. 94

12.1 Preliminary acoustic myography (AMG) data of the biceps and triceps show substantial correlation with muscle output force. *Left*: Simplified sagittal model of the elbow used in data analysis. *Right, top*: $A_1\nu_1$ of the biceps is highly correlated with output output torque τ ($r = 0.9, p < 10^{-6}$). *Right, bottom*: Example $A_1\nu_1$ and $A_2\nu_2$ trajectories (of the biceps and triceps, respectively) during random elbow stiffness modulation, showing significant correlation between the two data series ($r = 0.6, p < 10^{-7}$), consistent with maintaining constant output torque. ©2018 IEEE. 103

13.1 Idealized example of using deformation-based muscle force inference to enable simultaneous position, force, and stiffness control of a single joint in a teleoperation scenario. *Top*: Classical force-based control scenario in which human user's electrogoniometer-measured elbow angle θ_{u} is mapped to robot elbow angle θ_{r} through PID control loop with constants k_{P}, k_{I}, and k_{D}, with goal force trajectory F_{des} specified proportional to differential biceps–triceps surface electromyography (sEMG) signal α_{diff} and added to feedback-loop-generated position-correcting force F_{dy} to generate robot elbow force F_{r}. *Bottom*: By introducing ultrasound-measured muscle deformation D_{biceps} — a better analog for the user's absolute output force than the differential sEMG signal — we can configure the robot to track this output force more directly by specifying F_{des} as proportional to D_{biceps} and remap α_{diff} to generate controller stiffness constants consistent with the user's own elbow stiffness. This enables enhanced control in scenarios requiring precise independent modulation of position, force, and stiffness parameters (e.g., manipulation of cane-like weight-bearing tools, catching delicate objects). Note that this prototype example assumes ideal controllable-torque actuators; implementation would require careful actuator selection and likely innovation. Select images adapted from [67]. 106

13.2 Example usage of deformation-based force inference in inverse and forward musculoskeletal dynamics modeling. *Top*: Current musculoskeletal simulation frameworks (e.g., OpenSim [22], AnyBody [19]) convert joint angles and velocities $\{\theta, \dot{\theta}\}$ to joint torques $\{\tau\}$ via inverse dynamics, then convert these torques to muscle forces $\{F_{\text{m}}\}$ by optimizing over cost functions (e.g., minimum metabolic energy expenditure, gait cycle or sEMG matching) that may or may not represent true human dynamics. By measuring even a subset of muscle forces directly (e.g. F_{biceps} at the biceps brachii), we can begin to evaluate the validity of these cost functions, determine under what conditions they hold, and develop new ones. *Bottom*: Direct, simultaneous, deformation-based measurement of multiple muscle forces $\{F_{\text{m}}\}$ would enable the creation of forward dynamics models of the human musculoskeletal system, a class of models that is not yet widely used in the biomechanics community, as these individual muscle forces are not often otherwise available even during invasive study. 109

13.3 By providing mechanical, output-force-related deformation and vibration signals, both ultrasound and acoustic myography (AMG), respectively, "close" the neuromuscular control loop in terms of available sensor measurements, enabling enhanced understanding of neurological signals measured (via nerve cuff, EEG, ECoG, etc., *blue*) throughout the brain, spine, and peripheral nervous system (PNS), as well as a more holistic view of the way in which pathologies (*red*) throughout the system impact physical capabilities and behaviors. Image adapted from [67]. 110

List of Tables

3.1 Example Time Series Data 24

6.1 Example Tracking Error 53

7.1 Volumetric Data Collection Conditions, OpenArm 1.0 63
7.2 Volumetric Data Collection Conditions, OpenArm 2.0 69

8.1 Segmentation Accuracy 75

Acknowledgments

People often say a PhD is like a marathon. Having never run one, I can't speak directly to the validity of this analogy, but while I imagine both successful graduate students and long-distance runners exhibit similar levels of dedication and self-motivation, this metaphor makes grad school seem like an individual sport. Instead, I think of the PhD process as more like an escape room: everyone joins because they're excited to solve (scientific) puzzles, and brings with them a particular skill set, but it ultimately takes everyone's collective expertise to help each individual succeed.[1] I've been fortunate to benefit immensely from those around me during my own PhD journey — far more people than I can possibly mention here, but I'd like to thank as many as I can.

First and foremost, I'd like to thank my advisor, Ruzena Bajcsy, both for sharing her (terrifyingly extensive and always accessible) expertise and for giving me freedom and ownership over my own research path — and making me unafraid to "zig" when everyone else is "zagging," as she's been doing for decades. I've benefited equally from her years of perspective on the robotics field and from the warm community she's helped build as everyone's "academic grandmother."[2] I thank her for both keeping me honest (I'll never forget that there's "nothing new" and "no free lunch") and for building me up, and trusting in my own academic potential — earlier and more completely than I ever deserved.

Next, I'd like to thank the other members of my thesis committee, who have given me excellent and wide-ranging advice on what has become a highly interdisciplinary project. Hannah Stuart has granted me not only her expertise in technical design and manipulation science, but also great advice on building a research program and on the young faculty experience in general. Bob Full has given me an indispensable biologist's perspective on my modeling work, and I've always looked forward to his unrelenting (and aspirational!) positivity at each of our meetings, whether we're discussing my own research, cockroaches falling from the ceiling, or the comically ill-designed mammalian laryngeal nerve. Neville Hogan (who has been generous in sharing his time from across the country) has given me invaluable perspective — and lots of concrete technical guidance — on doing research at the intersection of neuroscience, biomechanics, robotics, and control, and I've benefited greatly from his candid and specific advice. Claire Tomlin, in addition to providing a valuable control theoretic perspective on my work, is the best mentor I know, and I hope to emulate her mentorship and teaching methods when I one day have students of my own.

Perhaps even more so than by faculty, my graduate experience has been shaped by my peers — who were both my best source of technical know-how and key to maintaining my sanity. To start, a big thank you to my fellow Human Assistive Robotic Technologies (HART) Lab members. In particular, thank you to Robert Matthew, Sarah Seko, Joel Loeza, Eric Mica, Waiman Meinhold, Patrick Ciccone, and Drew McPherson for welcoming

[1] This analogy also permits an appealingly variable level of cynicism that I find true to the PhD experience — we all signed up because we genuinely enjoy the challenge, but sometimes you just want to *get out*.

[2] I'll especially miss all her wonderful parties up in the hills, which were critical in helping me get to know both professors and other students and to feel like part of the team when I first arrived in Berkeley.

me onto the human modeling team, especially when I first arrived and was still getting my feet under me; to Aaron Bestick, Katie Driggs-Campbell, Daniel Aranki, and Victor Shia for being amazing role models, and always giving great technical and general survival advice; to Daisuke Kaneishi and Ben McInroe, for helping me broaden my expertise beyond the EECS department; to Akira Kato, for being an incredible co-author on my first real paper, with whom I had the whiteboard conversation that led to this dissertation's core insight; to Carolyn Matl, Isabella Huang, and Zoe Cohen for always picking up the slack when I dropped it, and sharing their time, skills, and kittens when I needed them most; and to Gregorij Kurillo, without whom I would never have been able to create such elaborate experimental platforms, who quietly kept the lab running like clockwork and always smiled at my jokes. Thanks also to my many undergraduate students, including Jeffrey Zhang, Aaron Sy, Shivani Sharma, Nicole Peternel, Irene Kim, Daniel Ho, Ian McDonald, Sathvik Nair, Logan Howard, Yannan Tuo, Yoni Nozik, Nandita Iyer, David Wang, Sachiko Matsumoto, Sai Mandava, Thomas Li, Kireet Agrawal, Michelle He, Evan Shu, Jason Liu, Hayden Sheung, Prerana Kiran, Jaeyun Stella Seo, Chris Mitchell, Akash Velu, Amanda Schwartz, Bhavna Sud, Varun Bhatia, Eric Hu, and Fayyaz Ahamed, for their tireless system building efforts, infectious energy, and technical competence that often surpassed my own.

Beyond the HART Lab, thanks to Justin Yim, David and Sara Fridovich-Keil, Mike Estrada, Frank Chiu, Margaret Chapman, Jaime Fernández Fisac, Joe Menke, Dapo Afolabi, Dexter Scobee, Sylvia Herbert, Tyler Westenbroek, Anusha Nagabandi, Andrea Bajcsy, Vijay Govindarajan, Eric Mazumdar, David McPherson, Vicenç Rubies Royo, Victoria Tuck, Valmik Prabhu, and the other denizens of 337/307 Cory and SDH 7, from whom I osmosed most of what I know about control theory and who provided me with both productive and unproductive distractions that were really the core of my grad school experience. Thanks also to the many wonderful humans I met (and free food I consumed) through WICSE and the EE/CSGSA.

A core part of my graduate experience has been in the classroom, as both a student and an instructor, and the expertise I gained there was instrumental to writing this thesis. Thanks to my own wonderful instructors in robotics, optimization, control, and statistics, the teams I taught with, and my students, from whom I learned just as much. Thanks in particular to Gireeja Ranade, who served as an amazing mentor and role model during a superlatively chaotic semester and who gave me the opportunity to develop content of my own, and to Amir Mousavi, Erik Rogers, and the rest of the (chronically underappreciated) ESG team for teaching me how to manage robot hardware and software of all kinds.

No one would ever graduate without the phenomenal work of the Berkeley EECS admins. Thanks especially to Shirley Salanio and Susanne Kauer, who know everything, and whose offices and help were always open to me; to Jessica Gamble, who herds cats better than anyone I've ever met, and who has saved me variously from administrative headaches, malfunctioning printers, and potentially moldy thai food; and to Leslie Goldstein, whom I'm not sure I've ever met but has handled the Bajcsy group finances with both extreme competence and admirable patience.

Graduating during COVID has presented exceptional challenges, both logistically and

mentally. A special thanks to my "accountabilibubble" — David McPherson, Victoria Tuck, Jaimie Swartz, and Kristen Cotner — for keeping me on track not just with writing, but living my life. Thanks also to all the amazing instructors and students at Cal Yongmudo — and in particular, my mentors David Commins and Elaine Chao, and our fearless leader Norman Link — for giving me not only martial arts skills, but a community beyond the EECS department, structure during the pandemic, and a reason to take care of myself physically in an environment that's not always conducive to doing so.

I couldn't imagine a better housing situation than New House West (aka New New House), where one of our first collective purchases was a whiteboard and our most heated argument was over whether spiders could be moved outdoors in our communal tupperware.[3] Thanks to all housemates (and de facto housemates) past and present — Kathryn Materna, Dan Borgnia, Aaron Brookner, Sophie Weber, Kristen Cotner, Marie Pelagie Una Elimbi Moudio, Sami Harper, Marc Lim, Tianjiao Zhang, Bennett Ng, Tobias Sjölander, Ana McTaggart, David Xiao, Krista Ann Lim, and Saul Wilson — for giving me a true home, and for humoring me when I attempted repair jobs instead of calling in a professional. Thanks also to Laura and Carolyn Peckham and to Audrey Zeis — the friends I wish I saw more often, and who keep me grounded when I do.

Above all, I owe my grad school career to the unwavering support of my family, both Hallock and Tilson (and, newly, Vega-Brown). A million thanks to my father, Gary — the first Dr. Hallock — for his decades of encouragement and instruction in all things STEM, and for empowering me to build and fix things myself; to my mother, Tricia Hallock, for giving me a love of reading and narrative that shapes my approach to research, and for wisdom that I quote more often than she knows; and to my sister, Deborah Hallock, the funniest and most compassionate person I know, and whose opinion I value most. I thank Will Vega-Brown more than I can say, for his unshakable emotional and practical support, commitment to collective problem solving, and for giving me many adventures to look forward to.

Lastly, thanks to (my real sources of grad school sanity) Aurora, Roxie, Ruby, Sasha, Abby, Elton, Pixel, Nike, Emmy, Hannah, Sibo, Bodie, Bodie, Mashi, Harold, Hobbs, Megan, Zoe, Zelda, Zeena, Molly, Boots, Nala, Brandy, Millie (and Rainbow and Jessie), Florence, Otis and her compatriot, several random corgis, kittens, and beehives, and especially King Midas, the best thesis-writing partner I could ever ask for.

[3] I *can* imagine a better housing situation than Granville, but I can imagine no better humans to share it with. :)

Statement of Co-Authorship & Prior Publication

This dissertation contains content adapted from the following publications (as noted at the start of each relevent chapter):

Laura A. Hallock, Akira Kato, and Ruzena Bajcsy. "Empirical Quantification and Modeling of Muscle Deformation: Toward Ultrasound-Driven Assistive Device Control". In: *IEEE International Conference on Robotics and Automation (ICRA)*. IEEE. 2018, pp. 1825–1832. DOI: 10.1109/ICRA.2018.8462887. [49]

Yonatan Nozik*, Laura A. Hallock*, Daniel Ho, Sai Mandava, Chris Mitchell, Thomas Hui Li, and Ruzena Bajcsy. "OpenArm 2.0: Automated Segmentation of 3D Tissue Structures for Multi-Subject Study of Muscle Deformation Dynamics". In: *International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*. IEEE. 2019, pp. 982–988. DOI: 10.1109/EMBC.2019.8857669. *Equal contribution. [113]

Laura A. Hallock, Akash Velu, Amanda Schwartz, and Ruzena Bajcsy. "Muscle deformation correlates with output force during isometric contraction". In: *IEEE RAS/EMBS International Conference on Biomedical Robotics and Biomechatronics (BioRob)*. IEEE. 2020, pp. 1188–1195. DOI: 10.1109/BioRob49111.2020.9224391. [51]

Laura A. Hallock, Bhavna Sud, Chris Mitchell, Eric Hu, Fayyaz Ahamed, Akash Velu, Amanda Schwartz, and Ruzena Bajcsy. "Toward Real-Time Muscle Force Inference and Device Control via Optical-Flow-Tracked Muscle Deformation". In: *IEEE Transactions on Neural Systems and Rehabilitation Engineering (TNSRE)* (2021, under review). [50]

Laura A. Hallock and Ruzena Bajcsy. "A Preliminary Evaluation of Acoustic Myography for Real-Time Muscle Force Inference". In: *International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*. IEEE. 2018. [48]

In addition to the co-authors above, we acknowledge the contributions of Gregorij Kurillo, Robert Matthew, Daisuke Kaneishi, Sarah Seko, Carolyn Matl, Isabella Huang, Jaeyun Stella Seo, Varun Bhatia, David Wang, Sachiko Matsumoto, Nandita Iyer, Ian McDonald, Jeffrey Zhang, Kireet Agrawal, Aaron Sy, Sathvik Nair, Logan Howard, Yannan Tuo, Vijay Govindarajan, Prerana Kiran, Sara Fridovich-Keil, Justin Yim, Shivani Sharma, Michelle He, Evan Shu, Jason Liu, Hayden Sheung, Nicole Peternel, and Irene Kim in hardware, software, and theory development and in data collection.

We also acknowledge the sponsorship of the NSF National Robotics Initiative (award no. 81774), Siemens Healthcare (85993), the NVIDIA Corporation GPU Grant Program, eZono AG, the Berkeley Artificial Intelligence Research Lab (BAIR), and the NSF Graduate Research Fellowship Program.

Chapter 1

Introduction

Despite the prevalence of musculoskeletal injury and pathology — from spinal cord injury, to muscular dystrophy, to amputation, to stroke — upper-limb assistive devices (including prosthetics, exoskeletons, and robots) remain severely limited in both overall dexterity and capacity to aid in tasks of daily living. While a number of sophisticated, anthropomorphic devices exist, with all the right joints in all the right places, none approach the speed and functional flexibility of the intact human arm, even when operated by a healthy human user.

To understand why these devices remain limited, consider the healthy human neuromusculoskeletal system. To perform a dexterous maneuver, the brain maps high-level intent to a complicated feedback loop involving visual, proprioceptive, and contact sensing, transmitting complex control signals to a multitude of motor units distributed across tens of muscles. At the macro level, there are no good models for this neurological feedback system; in particular, we have a poor understanding of how these neurological signals translate into muscle force, and even if this understanding were improved, the standard way to measure them noninvasively (surface electromyography, or sEMG) is noisy, nonspecific, and unreliable. This means that even the most straightforward possible control scheme, collecting data from a healthy human user and mapping it to a joint-for-joint robotic replica, remains beyond the means of existing systems — let alone controlling prostheses that replace missing limbs, or robots that replace lost motor capabilities.

My work provides a path toward high-degree-of-freedom, dextrous control by leveraging a single fundamental insight, illustrated in Figure 1.1: *for devices in which we wish to cede a high degree of control authority to the human user (prostheses, teleoperated robots, etc.), we can use mechanical measures of muscle output force as a control signal, rather than neurological activation or poorly-modeled "intent."* This approach allows us to make predictions about human dynamics without directly modeling the neurological control loop — until we want to explicitly study it.

Perhaps surprisingly, estimating these individual muscle forces remains an open problem in the biomechanics community; no sensing modality exists that can measure such forces directly and noninvasively in vivo. Addressing this force inference challenge is the primary contribution of this thesis work, which introduces two novel, promising signal classes from

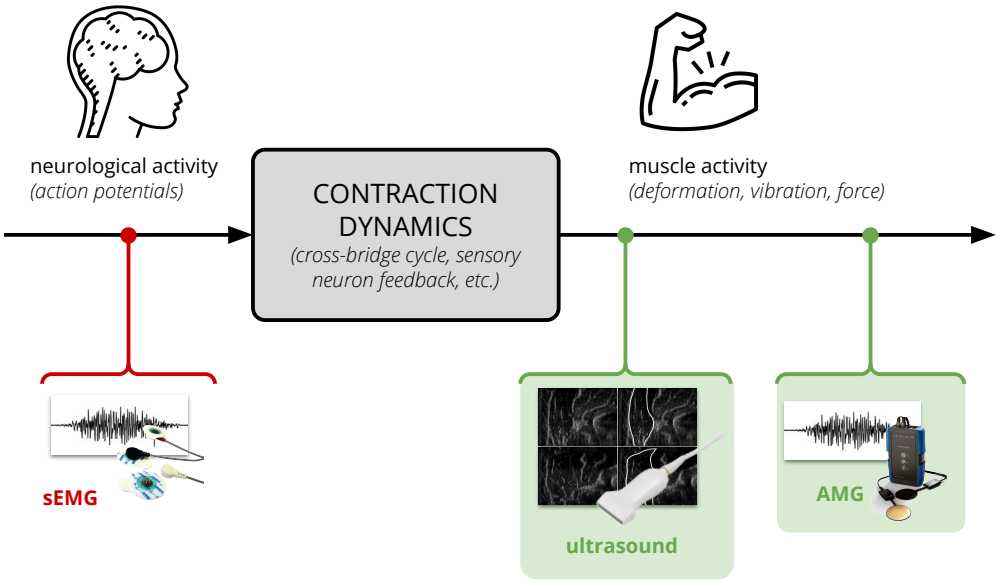

Figure 1.1: Human motion results from complex neuromuscular contraction dynamics that convert a neurological activation signal into muscle output force. This process is poorly understood and difficult to measure, greatly limiting our ability to replicate such motions in robots and assistive devices from neurological sensor data like surface electromyography (sEMG, *red*). This dissertation work (*green*), in contrast, infers motion dynamics more directly by measuring two classes of signals intrinsically coupled to force output: deformation (via ultrasound) and vibration (via acoustic myography, or AMG). These signals can be used directly to control safe and dexterous assistive devices without modeling these complex neuromuscular contraction dynamics, and the resulting models can yield new insights into healthy and pathological human motion. Select icons made by hunotika from the Noun Project and smalllikeart from Flaticon.

which this force can likely be inferred: muscle deformation (as measured via ultrasound) and muscle vibration (as measured via acoustic myography, or AMG). In particular, this work presents a thorough examination, and proof-of-concept control use application, of deformation, which — unlike sEMG or AMG — can be measured in a highly localizable manner, allowing for straightforward extraction of multiple independent signals.

At the cellular level, the relationship between muscle deformation and output force is well understood: the cross-bridge cycle ratchets muscle fibers together, resulting in a shape change that pulls the connected tendon and imparts force to the skeleton. At the same time, prior to this work, there existed no macro-level, in vivo characterizations of this force–deformation relationship during natural motion, nor even any data with which to begin study, let alone models for how to leverage these deformation signals for control. This dissertation thus presents preliminary work in two complementary directions: first, prototyping the use

of simple deformation signals to perform a proof-of-concept control task (i.e., trajectory tracking), and second, building a comprehensive body of volumetric data, as well as an analysis framework, to characterize this force–deformation relationship and ascertain the most correlated signals.

Specifically, the contributions of this dissertation, as organized by part and chapter, are as follows:

- a preliminary time series correlation analysis of simple, 2D muscle deformation signals with output force at the elbow (part I), including an exploration of candidate signals under multiple kinematic configurations (chapter 3) and a refined analysis of select signals using a novel, automated tracking framework (chapter 4);
- proof-of-concept use of a simple deformation signal (i.e., thickness change) for control (part II), including a comparative study of deformation- and activation-based trajectory tracking via elbow flexion (chapter 5) and a preliminary exploration of additional tracking methods and signals (chapter 6);
- generation and open-source release of two multi-subject volumetric data sets of the arm, under multiple kinematic configurations and force conditions, to enable comprehensive study of the force–deformation relationship (part III), including data set collection (chapter 7) and neural-network-based tissue segmentation (chapter 8); and
- preliminary modeling analyses leveraging these 3D data sets (part IV), including initial geometric analyses (chapter 9) and a proposed systematic modeling framework (chapter 10) for future exploration of the force–deformation relationship.

Also included are a discussion of related work contextualizing and motivating this document's contributions (chapter 2), as well as (in part V) information on released code and data (chapter 11), a brief exploration of an alternative muscle force sensing modality (i.e., AMG, chapter 12), and some final thoughts on limitations and future applications of this work in control, biomechanics, and rehabilitative medicine (chapter 13).

Ultimately, this work aims to form the basis for future exploration of the force–deformation relationship for a wide range of applications, enabling both fundamental technical advances (in assistive technology, human–robot collaboration, and diagnosis and rehabilitation of pathology) and theoretical breakthroughs (not only in human movement science, but control theory, robotics, machine learning, and system identification more broadly).

Chapter 2

Motivation & Related Work

In this chapter, drawing on existing biomechanics, robotics, and assistive technology literature, we argue first, that understanding particular aspects of human musculoskeletal dynamics — specifically, individual muscle forces — is critical for many applications in assistive device control, physical human–robot interaction, and rehabilitation of pathology; second, that the human elbow provides a relevant and tractable target for preliminary force modeling; third, that current modeling frameworks do not provide sufficient force inference accuracy to enable the desired safety and capability of these systems; fourth, that existing sensing modalities cannot fully close this modeling gap; fifth, that muscle deformation is a promising signal from which to infer detailed muscle forces; and lastly, that ultrasound is an ideal modality — one that is safe and commercially well-positioned for future incorporation into real-world devices — with which to measure these deformation signals and use them for biomechanical study and control.

Together, these claims motivate the two key contributions of the remainder of this dissertation: the proof-of-concept use of ultrasound-measured deformation for control, and the development of a comprehensive data set and analysis framework to enable future, more expressive force–deformation models for robotic and therapeutic applications.

2.1 The Case for Improved Human Muscle Force Modeling

Despite decades of study, noninvasive, in vivo, real-time measurement (or even accurate prediction) of musculoskeletal dynamics remains an open problem. In particular, without good models of muscle force output during natural movement, our understanding of how humans execute dexterous motions — in terms of position, force output, and impedance — is fundamentally limited, as is our ability to safely modify or replace this execution using assistive devices and to accurately characterize and treat musculoskeletal pathology. In this section, we first clarify the relationship between muscle forces and the dynamics of human motion to argue for the importance of individual muscle force inference, then highlight three

key application areas in which we see existing modeling frameworks as particularly insufficient and improvements to these models as especially critical.

Skeletal Muscle: The Human Actuator

The human musculoskeletal system is capable of a wide range of motions comprising changes in position (and its derivatives), output force, and impedance — all of which are critical to how we interact with the world, from locomotion to tool manipulation. Healthy humans are, in fact, experts at modulating each of these aspects of motion: we can precisely control foot position to walk on uneven ground; we can pick strawberries without crushing them; we can use a cane or a power drill.

At the same time, with current sensing technology, our ability to measure and model these aspects of motion is limited. At the joint level, changes in position can be measured comparatively easily (e.g., via motion capture [102], inertial measurement units [129], or electrogoniometry [126]), and changes in force can be monitored in more limited settings via sensorized environments (e.g., using force plates or contact sensors [125, 94], or exoskeleton-inbuilt torque or strain sensors [78, 95]), but changes in impedance (including stiffness, inertia, and damping properties) can generally be measured only under tightly controlled conditions [109] (e.g., during oscillation [46] or impact perturbation [33] or "creep tests" [121]) and not during arbitrary natural movement, restricting their modeling usage to stereotyped motions (e.g., the walking gait [89], reaching and grasping [111]).

While each of these aspects of motion is influenced by a number of factors — including external load, limb masses and inertias, and the frictional properties of tissue structures — humans modulate each of them using the same actuators: muscles. Specifically, for a given set of environmental conditions, a change in muscle force (i.e., tension, as discussed below in section 2.3) can result in a change in position (as the muscle–tendon unit, or MTU, shortens and pulls along the skeleton), a change in force (as MTU tension increases to compensate for external load), and/or a change in impedance (as multiple MTUs exert tension in different, often antagonistic directions about a joint [109], and muscle material properties change). Thus, measuring individual muscle forces would enable substantial advances in our ability to model each of these aspects of human dynamics — and in particular, force and impedance trajectories during natural motion, which are largely inaccessible using current technology, as discussed above.

Key Application Areas

In this subsection, we highlight three areas in which improved human dynamics modeling — specifically, greater understanding of force and impedance properties, as enabled by measurement of individual muscle forces — could greatly enhance state-of-the-art approaches.

Assistive Device Control & Human–Robot Collaboration

When constructing assistive devices that interact physically with the human body, such as exoskeletons and prostheses, it is crucial to understand not only what overall joint moments are being applied to the human user, but what muscle forces the user is induced to exert based on this interaction. This information is critical foremost for device safety analysis — i.e., determining whether the device induces forces that will result in acute or long-term injury — but also, in cases of rehabilitation, assessment of whether a patient is actively performing desired motions, or simply relying on the device. Additionally, these devices remain far less dexterous than the limbs they are meant to aid or replace, a limitation that improved musculoskeletal modeling and control signal extraction can address.

More broadly, enhanced human models are necessary in the wider field of physical human–robot collaboration, in which robots must intuit human collaborators' movements, intents, and capabilities to safely accomplish tasks (for example, helping a factory worker with complex part assembly, or aiding an elderly individual in performing lifting tasks). While current robotic systems sometimes leverage basic cognitive and ergonomic models to evaluate user comfort (as discussed below in section 2.3), both safety and efficient allocation of tasks across human and robotic collaborators requires a deeper understanding of human users' physical capabilities.

Human Neuromuscular Science

Beyond applications to device control, finding a noninvasive measure of individual muscle forces is recognized by the biomechanics community as a core challenge inhibiting our understanding of human motion [56]. Specifically, a measure of muscle output force would greatly advance scientific understanding of the neuromusculoskeletal system in two antipodal directions: with a view toward total-body modeling, individual muscle measures would allow for study of force balancing across synergists and agonists/antagonists, relationships that are central to understanding capabilities like balance and smoothness of motion, without the explicit need for measurement of neurological signals; conversely, a mechanical muscle force measure constitutes a novel output signal corresponding to this neurological input, enabling new insights when interpreting surface electromyography (sEMG), electroencephalogram (EEG), or electrocorticograph (ECoG) data.

Diagnosis & Rehabilitation of Musculoskeletal Pathology

These kinds of advances in device control and motor science are perhaps most impactful when they can be leveraged to address musculoskeletal pathology. Many such pathologies — including muscular dystrophy, spinal cord injury, stroke, and Parkinson's disease — often present with pathologies in force distribution (e.g., disrupted balance) and lack quantitative metrics (beyond, e.g., the Brooke scale [35]) to evaluate disease progression and treatment efficacy. Muscle force measures could provide much-needed insights regarding both capability and behavior after injury or disease onset and throughout rehabilitation, and to evaluate the

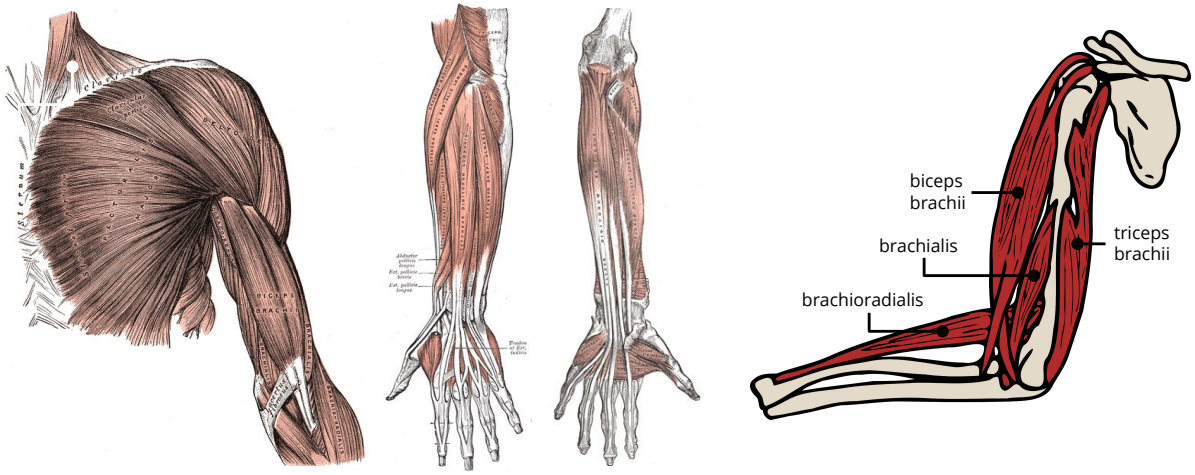

Figure 2.1: Musculoskeletal anatomy of the human arm. *Left*: The human arm is complex, with tens of muscles actuating multi-degree-of-freedom (DoF) joints in synergistic and antagonistic ways. *Right*: The elbow provides a comparatively simple modeling testbed, with only three flexors (the biceps brachii, brachialis, and brachioradialis) and a single extensor (the triceps brachii) actuating a single-DoF joint. Left and center images from [42]; right image adapted from [154].

efficacy of treatments, whether they are rehabilitative devices, pharmaceuticals, or lifestyle changes.

2.2 Our Focus: Net Torque at the Human Elbow

In light of the target application areas above, we focus our modeling efforts in this thesis on the human arm, examining signals and modeling schema appropriate for use in upper-limb (manipulation-type) device control. While the techniques we present could likely be used to model a number of other joints and motions (e.g., gait at the legs, balance at the torso), we view novel models at the arm to be especially important, as existing assistive device control paradigms remain particularly limited when compared with healthy human motion.

Below, we motivate our specific model target — the human elbow — and outline the scope of our (as yet limited) modeling capabilities.

Modeling Target: The Human Elbow

As illustrated in the left and center images of Fig. 2.1, the human arm in its entirety is incredibly complex, with tens of muscles each actuating one or more joints, many of which have multiple degrees of freedom (DoF). In this dissertation, we focus on the elbow as a case

study for our proposed sensing and modeling methods, a comparatively simple, single-DoF joint with only a few actuating muscle groups (three flexors and a single extensor, as shown in the right image of Fig. 2.1) that is nevertheless critical to performing manipulation tasks requiring substantial force exertion.

Modeling Scope: Net Joint Torque Inference

The primary sensing method we explore in this thesis — specifically, the use of ultrasound to measure muscle deformation — was selected specifically for its ability to attribute signals to individual muscles, enabling individual muscle force–deformation modeling: as discussed below in section 2.5, muscle deformation is an inherent aspect of the muscle force generation mechanism, and this deformation can be measured directly for an individual muscle (or, indeed, any spatial location) via ultrasound. At the same time, this work is valuable precisely because there exist no standard methods for noninvasive, in vivo individual muscle force measurement, which leaves us unable to validate force–deformation models at the level of individual muscles.

Instead, in this document, we limit our analyses (aside from some speculative modeling techniques in chapter 10) to the relationship between muscle deformation and net output joint torque, a measure that is easily accessible via contact force-torque sensing. Nevertheless, we feel that the data we present, alongside our understanding of the physiological relationship between deformation and output force, constitute a powerful case that this signal is a promising candidate for individual muscle force measurement. In the future, we aim to validate these assertions using both invasive and noninvasive sensing, as discussed in section 13.2.

2.3 Existing Musculoskeletal Modeling Frameworks

In this section, we enumerate existing musculoskeletal modeling frameworks — focusing on those most relevant to device control and macro-level biomechanical modeling — and remark on their utility and limitations when applied to the domains of interest above.

Micro-Scale Models: Neurological Control & the Sliding Filament Theory

At the cellular level, the mechanism by which skeletal muscles generate force is well-characterized: the nervous system (central or peripheral) sends an electrochemical signal to a given motor neuron, which then sends an action potential to ennervate a "motor unit" of a few or more muscle fibers via chemical depolarization, filaments of which are induced to slide relative to one another via the actin-myosin cross-bridge cycle [65, 66] and stabilized by titin and other elastic proteins [57], causing a length (and thus tension) change and imparting force. Increased force in a given muscle results from increased force contributions of a single

motor unit (induced by more frequent action potentials) and from the recruitment of larger motor units in increasing order of size (following the Henneman size principle [55]).

While these mechanisms cannot be tractably sensed and modeled for direct full-muscle in vivo analysis, they highlight several phenomena that must be considered when modeling macro-scale musculoskeletal dynamics. First, the force generation signal may originate in the brain or as a peripheral reflex; second, motor units (which come in many sizes [15]), and not full muscles or individual fibers, are the quanta by which muscles are activated; and third, each motor unit may be activated by only a single action potential ("twitch"), multiple summed action potentials ("summation"), or to maximum ("tetanus").

Additional Complexities: Contraction Types, Fiber Types, Hysteresis, & Fatigue

In addition to the intricacies of the mechanism described above, muscle contraction dynamics vary with a number of factors, including contraction type, relative directions of motion and exertion, fiber type composition, stretch reflex prominence, and fatigue. In this subsection, we delineate modeling considerations necessitated by these phenomena.

Contraction types. While muscles exhibit time-varying force, length, and velocity behavior during natural motion, contraction behaviors are often simplified as being isometric (changing in force, but with no associated length change or velocity), isokinetic (moving at constant velocity, with varying force and length) or isotonic (consistent in force, but changing in length and velocity), and muscle force generation capabilities vary substantially with contraction type [80, 60]. In this dissertation, we primarily target isometric contraction for simplified initial modeling, though we ultimately aim to build a system that accounts for all combinations of these contraction modes.

Contraction directions. While muscles can only execute unidirectional force (i.e., can only pull, not push), they can pull with (concentric) or against (eccentric) the direction of joint motion, shortening or extending during contraction, respectively. Concentric contraction enables overall skeletal movement; eccentric contraction is critical for stabilizing and smoothing motions, balancing, and as a braking force to prevent joint damage [31, 112]. Contraction dynamics exhibit considerable differences in these two regimes: muscles are substantially stronger and require less energy expenditure when contracting eccentrically [112], and eccentric training is generally considered to be more effective at increasing force generation capability, but also more damaging [123].

Fiber type. Muscle fibers exhibit different temporal force generation properties (e.g., time for tension to rise to peak value during twitch) depending on their primary metabolic energy source. In particular, three fiber types have been identified [103]: slow oxidative (SO, powered primarily by aerobic respiration), fast oxidative-glycolytic (FOG, powered by both aerobic respiration and anaerobic glycolysis), and fast glycolytic (FG, powered primarily by anaerobic glycolysis).[1] Muscles generally exhibit a mix of fiber types, as each produces move-

[1]Interestingly, these fiber types also differentiate light and dark meat, which correspond to fast and

ment advantages (SO for fatigue resistance, FG for speed, and FOG for strategic flexibility), and this balance can be modified with appropriate training and exercise [23].

Myotatic ("stretch") reflex & hysteresis. In addition to active, central-nervous-system-controlled output forces, muscles exert reflexive contraction forces when passively stretched (as sensed by the muscle spindle, a stretch receptor within the muscle body that conveys information back to the nervous system via afferent nerve fibers), acting to keep the muscle at constant length [25]. This reflex is critical to maintaining balance and stability, but also exhibits substantial hysteresis depending on current and previous movement conditions that make modeling difficult [82].

Fatigue. Muscle fatigue, a decrease in maximal force output and power production during contractile activity, both acute and chronic, can occur due to exertion-related changes in any aspect of the contractile mechanism, including neurological activation, ion concentrations, and metabolic energy systems [149]. In general, these changes are poorly characterized and rarely accounted for in real-time musculoskeletal modeling frameworks.

Hill-Type Muscle Models

Although there exist models extending the sliding filament theory to muscle-level dynamics [157], the complexities described above have meant that the most successful macro-scale dynamics models are *phenomenological* — i.e., models largely constructed from data using system identification techniques rather than grounded in detailed biological processes. The most widely-adopted example is the Hill model [59], one common variation of which is shown in Fig. 2.2, a lumped parameter model containing a nonlinear contractile element (CE) and nonlinear parallel (PE) and series (SE) elastic spring elements, corresponding roughly to the activating muscle, elastic elements surrounding and throughout the muscle (fascia, etc.), and the tendon, respectively. The force–velocity relationship of CE is then described by equation

$$(F + a)(v + b) = (F_{\max} + a)b$$

for muscle tensile force F, muscle shortening velocity v, maximum (tetanic) isometric force $F_{\max}$, constant coefficient of shortening heat a, and constant $b = av_{\max}/F_{\max}$ for maximum shortening velocity $v_{\max}$. This force–velocity model is often paired with force–length models derived from the sliding filament theory [40].

This Hill model is advantageous in that it requires fitting only a few parameters to a given muscle ($F_{\max}$, a, and b above, along with a few tendon stiffness parameters) and easily supports a number of structural extensions [151]. Due to this simplicity, it is widely employed for musculoskeletal modeling and simulation, including in the state-of-the-art optimization systems described below [22, 19]. At the same time, the model has many limitations often ignored by its users: in particular, the model was only ever designed to describe maximal, tetanic contractions, and often poorly describes submaximal contraction during natural motion [151]. In addition, fitting even these few parameters becomes an underdetermined and challenging problem when simulating many muscles at a time, as discussed below.

(myoglobin-rich) slow twitch fibers, respectively.

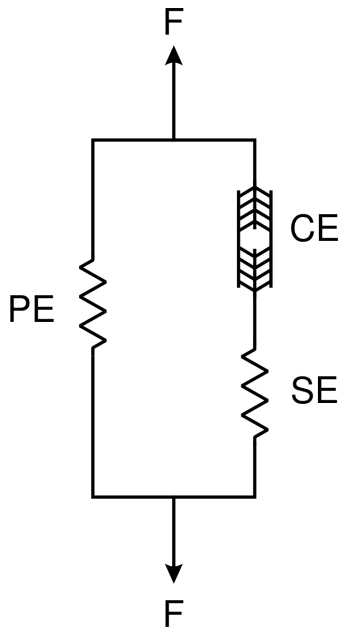

Figure 2.2: Basic Hill-type muscle model [59], a ubiquitous model of muscle force–length and force–velocity relationships, including nonlinear contractile element CE, series elastic element SE, and parallel elastic element PE. Image from [54].

Optimization-Based Models: Full-Body Simulation

Drawing on the Hill-type muscle models described above, as well as extensive literature in tissue morphology, several simulation frameworks have been developed to enable full-body musculoskeletal simulation, including OpenSim [22] and AnyBody [19]. These frameworks have been applied to a wide array of human modeling domains, including gait simulation [153, 143], sports training [87], rehabilitation [84], and device evaluation [85].

At the same time, these detailed simulations intrinsically make many morphological assumptions and offer limited customization parameters, relying on sparse and aggregate data to fit models to individuals and thus generating substantial modeling error. These systems also have to contend with a significant "biological nullspace" [137] (induced by redundant muscles) of possible muscle force values for any given joint force trajectory. As illustrated in Fig. 2.3's toy example, even simple motions admit a range of possible solutions for muscle output force values, forcing these frameworks to rely on optimizations (e.g., least possible force output, matching a gait cycle) known to be inaccurate when modeling many phenomena of interest (e.g., balance, co-contracted or stabilized motion) [20] — and especially when measuring pathological motions (e.g., stroke-induced plegia, antagonistic co-contraction due to Parkinsons' disease) — and thus precluding any safety guarantees when used to inform

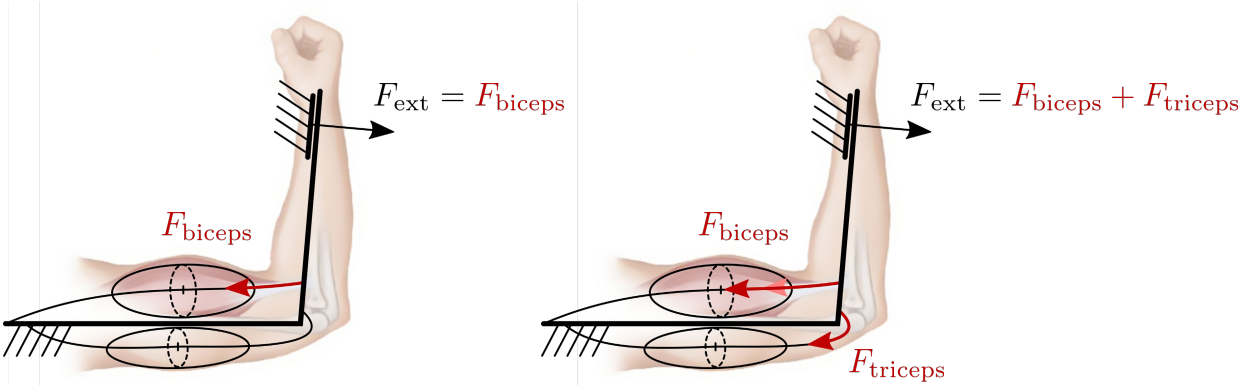

Figure 2.3: Toy example illustrating limitations of musculoskeletal simulation frameworks, depicting a simple sagittal model of the elbow with one biceps-like flexor exerting force F_{biceps} and one triceps-like extensor exerting force F_{triceps}. Both left and right images depict muscle forces consistent with observed output force F_{ext}; in the absence of muscle force measurement, simulation frameworks are forced to choose a solution. Most frameworks will likely choose the solution depicted on the left if optimizing for minimal exertion, though the image on the right is likely more consistent with stabilized, higher-impedance motion. Select images adapted from [67].

device control.

Higher-Level Models: Cognition & Ergonomics

Despite the advances above, the use of musculoskeletal modeling remains relatively uncommon in robotics and device control fields. Instead, many human–robot systems leverage simpler human models grounded in ergonomics [132, 11], assumptions of intended trajectory characteristics [34], cognitive intent inference [68, 52, 93], or simple rigid body dynamics [133]. Considering these frameworks is mostly beyond the scope of this dissertation, as they model fundamentally different attributes of the human system, and we see them as complementary to our musculoskeletal dynamics inference efforts.

2.4 Existing (Noninvasive) Musculoskeletal Sensing Paradigms

As discussed above in section 2.1, while there exist a number of sensing modalities with which to measure real-time human kinematics (including motion capture and electrogoniometry), as well as ground reaction and contact forces (e.g., via force plates and load cells), sensing of individual muscle forces — especially noninvasively — remains limited. Surface electromyo-

graphy (sEMG) — which measures (via surface electrodes) the electrical signals generated throughout the muscle by the motor unit action potentials (MUAPs) described above in section 2.3 — is by far the most ubiquitous modality from which muscle force is inferred. A rich body of literature, using a wide variety of electrode configurations (from simple differential to complex arrays [105]), has shown success in decomposing the various MUAP contributions to the sEMG signal to measure which muscles are synergistically activated during a variety of motions and with what relative timing [108]; however, while individual sEMG signals are generally correlated with output force, they cannot be used directly for force inference, as signal magnitudes are highly dependent on electrode locations, subject morphology, and many other factors, and often vary substantially from day to day or session to session [21]. Nevertheless, enabled by additional modeling assumptions, sEMG signals are often used for force inference [21], including via the optimization-based frameworks discussed above [130, 61, 119], and we use these signals as a baseline for comparison with our own deformation-based force inference methods throughout this thesis.

A second, relatively recent force inference method involves measuring tendons' vibrational behavior to infer strain — specifically, using a shear wave tensiometer that exerts micron-scale taps on the tendon of interest, the resulting wave speeds of which are then measured by skin-mounted accelerometers [99]. This method has been used successfully at both the knee and ankle joints (at the achilles and patellar tendons [26], for example), and is a promising technology with which to validate much of this thesis work, as discussed in section 13.2. At the same time, this technique is limited in two important ways: first, it relies on tendon access near the skin surface, which is not the case for many muscles of interest (e.g., in this thesis, the brachialis), and second, it assumes a roughly elastic "beamlike" tendon to compute expected shear wave propagation, a quality that is likely unsatisfied by many (geometrically complex) tendons and aponeuroses. Nevertheless, for MTU structures that meet these criteria, this "tendon tapping" method is arguably the most direct extant noninvasive method of individual muscle force measurement.

A final, less common muscle force sensing paradigm is to measure contraction-associated muscle vibrations using microphones or inertial measurement units, in a method known variously as acoustic myography (AMG), vibromyography (VMG), or mechanomyography (MMG). While these signals remain poorly characterized — in fact, the mechanistic sources of the vibration signal are still debated — they present an additional source of contractile dynamics information, and our own preliminary study evaluating vibration-based force inference is included in chapter 12.

2.5 Deformation as a Measure of Musculoskeletal Forces

Deformation provides a force inference signal that addresses a number of limitations of the signals discussed above. First, like tapping-induced tendon shear waves and vibration, it is

mechanically, rather than *neurologically* (as sEMG-measured activation) coupled to muscle dynamics, such that it provides a measure representative of actual force production rather than neurological intent. Second, when measured via ultrasound (and unlike vibration or activation), it can be precisely localized to specific structures of interest, rather than a nonspecific area around the transducer, theoretically allowing for the extraction of multiple independent signals from different muscles (e.g., to allow control of different device degrees of freedom, or simply to enhance biomechanical understanding). Lastly, unlike any of the methods above, the deformation signal can be extracted equally well from both deep and superficial muscles, enabling more complete consideration of the full musculoskeletal system.

Despite the intricacies described above in section 2.3, the core mechanism relating muscle deformation and output force is straightforward: as muscle fibers ratchet together in the cross-bridge cycle described above, they induce a length change along the line of action, which pulls the attached (roughly elastic) tendons, which then impart the force to the skeleton, a phenomenon that has been observed and characterized empirically for select muscles at select (e.g., tetanic) conditions [97]. Under the (mild) assumption that muscles are isovolumetric [139, 103], this length change by definition induces further shape changes, or deformation, in the activated muscle. A number of isolated studies have observed this deformation empirically under limited conditions and established a correlation with muscle activation [104, 131, 159, 135, 72, 45].

At the same time, the human musculoskeletal system is highly complex, geometrically irregular, and dominated by the physics of various nonlinear materials — as illustrated in the above Fig. 2.1 — and morphology often varies substantially across individuals in ways that significantly impact system dynamics [64, 27].[2] While muscle deformation can be observed in one or two dimensions (e.g., muscle thickness or cross-sectional area changes, respectively), there do not exist models that can translate these observations into clinically-relevant quantities like stiffness, activation, and force output, and there exist few macro-scale force–deformation models relevant to assistive device control or system-level biomechanics; fewer still are (even cursorily) validated with experimental data. Within the field of computer graphics, a number of algorithms exist to generate plausible muscle deformations given a virtual model's kinematic configuration, dynamics, and/or skin surface deformation [117, 90, 91, 163], but these algorithms are generally evaluated on their output's visual appeal rather than a physiological or data-driven error metric. Although some such models are rooted in biomechanical data and are intended for real-world medical applications [163], they largely ignore muscles' fundamental nature as actuators, instead treating them as passive elements whose deformation is purely a function of configuration and applied external force, without regard to force distribution across muscles. Lastly, most of these models rely on finite element analysis, which makes them too computationally intensive for direct application to

[2] In addition, it is an open question whether this deformation can be observed in muscles without well-defined series tendons and/or with otherwise complex connective geometry (e.g., the deltoid), in which the one-dimensional length change mechanism described above is likely to break down; as the underlying cross-bridge-cycle mechanism of force generation is the same, we anticipate that there exists observable deformation of some kind, but that characterizing it may prove more challenging.

a real-time control scheme.

At the other end of the muscle deformation modeling spectrum are micro- and mezzo-scale analyses of (one to a few) individual muscle fibers [83, 44]. While these models offer significant insight into the biological sources of deformation, they are not readily extensible to macro-scale analysis. Indeed, there is significant evidence not only that fibers behave quite differently in vivo and in isolation (due to contact dynamics and other factors), but that passive tissues like tendons and aponeuroses significantly impact the deformation behavior of the muscle-tendon unit when modeled at this level of detail [63]. This evidence makes it difficult or impossible to conjecture how deformations scale across model resolution.

Another deformation-based force inference approach is to measure the force-generating muscle–tendon unit length change directly, in either the muscle or the tendon [81, 5, 162]. These approaches are complementary, but limited, in a similar manner to the "tendon tapping" methods described above in section 2.4: several muscles of interest may be attached to the same tendon, and these attached tendons — or, in many cases, aponeuroses — may have complex geometry that makes establishing a single "length" value difficult or impossible.

While these existing studies provide evidence that deformation can be used for force inference, there exist no principled, generalizable frameworks for modeling the force–deformation relationship in a widely applicable, real-time manner. Ultimately, the field of muscle deformation modeling could greatly benefit from its own version of the Hill model [59] — a ubiquitous, cleanly parameterized formulation that could be applied to a wide variety of musculoskeletal modeling endeavors. While the creation of such a model will doubtless require years of rigorous system identification, both in vivo and in vitro, this dissertation presents the first steps toward this kind of phenomenological modeling, first validating the correlation results above on a manipulation-relevant joint (the elbow), then testing their use in a real-time control application, and lastly, building a data set and analysis framework to allow for comprehensive examination of many possible models.

2.6 Ultrasound as a Sensing & Control Modality

There exist a number of sensing modalities with which we could measure this force-associated muscle deformation, from magnetic resonance imaging (MRI) to scan full muscle geometry [12] and strain [161], to force myography (FMG) [18] and tactile myography (TMG) [70] to measure shape and force changes at the surface of the skin. When used for real-time inference of muscle dynamics, each of these modalities presents both advantages and limitations: MRI provides a detailed view of muscle dynamics, but is expensive, slow and/or repetitive to collect, requires collection within a limited bore, and is not immediately applicable to device control outside lab settings; FMG and TMG are readily wearable, but cannot yield insight into the complex deformation relationships between deep and surface muscles that characterize many joints' actuation. (Elbow joint flexion, for example, the motion we explore in this thesis, is generated not only by the biceps brachii and brachioradialis, both surface muscles, but also by the — deep — brachialis.) In this dissertation, given our aspirations

toward device control and real-time modeling of natural motion, we present ultrasound as a deformation measurement technology that combines the strengths of the above sensors, yielding images of both shallow and deep muscle geometry while maintaining a reasonably portable form factor.

While commercial ultrasound-driven assistive devices are not yet available, diverse studies have demonstrated the use of 2D brightness mode (B-mode) ultrasound to measure a wide variety of muscle architectural parameters, including muscle thickness, pennation angle, and fascicle length [62, 118], as well as to execute preliminary ("sonomyographic") control on both simple single-degree-of-freedom prostheses [16], and on more complex devices by leveraging learning-based image processing and gesture classification [150, 24, 2, 127, 29]. A number of studies have specifically targeted muscle thickness, both as a feasible signal for prosthesis control [160] and as a measure of muscle activation (as correlated with EMG) [104, 76] and fatigue [135]. Additionally, wearable 1D amplitude mode (A-mode) ultrasonic sensors — which could potentially be applied to acquire a single muscle thickness measurement — have been used to continually monitor the mechanical properties of plantar soft tissue in diabetic patients [142] and knee torque [144] and are much closer to commercial integration with assistive devices than 2D B-mode scanners.

At the same time, these ultrasound-based control techniques have not seen wide adoption, are rarely evaluated against state-of-the-art sEMG control systems, and are inhibited by poor understanding of the underlying musculoskeletal dynamics. In this thesis, we begin to address these limitations by developing enhanced force–deformation models and ultrasound-based control paradigms in tandem and alongside corresponding sEMG data.

2.7 Summary

In this chapter, we introduced the necessity of enhanced musculoskeletal modeling in robotic and human motion science applications and argued that ultrasound-based extraction of deformation signals represents a promising path toward generating the expressive models we seek. In the following chapters, we leverage this insight to both prototype deformation-based control at the elbow joint and develop a systematic force–deformation analysis framework to enable expanded modeling in the future.

Part I

Exploratory Data Analysis

Chapter 3

Exploratory Force–Deformation Correlation

This chapter is adapted in part from [51].[1]

As discussed at length in chapter 2, muscle deformation is a rich, complex signal class, and there exist no standard parameterizations or control-ready models relating it to output force. In this chapter, we begin our systematic analysis of the force–deformation relationship at the elbow with a simple correlation study, examining several deformation signals of the brachioradialis alongside surface electromyography (sEMG) and contact force data during varied isometric contraction.

In the sections below, we first outline the scope of this preliminary study and detail our subject cohort and collection procedure, then present example time series data alongside qualitative and quantitative analysis of the force–deformation relationship. Lastly, we discuss the implications and limitations of our findings — namely, that several measures of brachioradialis deformation indeed correlate with output force and are consistent with simultaneous sEMG data, validating prior correlation results in the literature and informing our choice of deformation-based signals for both control and detailed study in the remaining chapters.

All data and analysis code have been released to the OpenArm project in the OpenArm Multisensor 1.0 package (see section 11.1).

[1] ©2020 IEEE. Reprinted, with permission, from Laura A. Hallock, Akash Velu, Amanda Schwartz, and Ruzena Bajcsy. "Muscle deformation correlates with output force during isometric contraction". In: *IEEE RAS/EMBS International Conference on Biomedical Robotics and Biomechatronics (BioRob)*. IEEE. 2020, pp. 1188–1195. DOI: 10.1109/BioRob49111.2020.9224391.

3.1 Study Scope

In this study, we target the cross section of the brachioradialis (the smallest of the elbow flexors, and thus the most visible in a single time-varying ultrasound frame) as a source of deformation signals, analyzing cross-sectional area, thickness, and aspect ratio changes — signals we expect to correlate well with muscle force given the basic force–deformation mechanism described in section 2.5 — during time-varying isometric elbow flexion. We collect these signals alongside sEMG-measured activation of the biceps brachii and brachioradialis and output joint force at the elbow, during a variety of pulsed and sustained flexion pulses, under multiple kinematic configurations (i.e., elbow angles) and for multiple subjects. We examine correlations between these deformation and activation signals and output force and how they vary by subject and kinematic configuration.

3.2 Data Set Collection

As an exploratory data set, we collected simultaneous sEMG, ultrasound, and output force data from a single subject under multiple kinematic configurations during a sequence of voluntary flexion pulses, in a protocol similar to that used in [47] and illustrated in Fig. 3.1. To explore consistency across individuals, we collected ultrasound and force data from an additional 4 subjects at a single representative configuration. Details of this subject cohort and collection procedure are outlined below.

Subject Biometric Data & Consent

Data were collected from the right arms of 5 subjects (2 female, 3 male). Of these, 4 subjects were within age 18–24, while the last subject was age 83 and used to investigate generalizability across age groups. All subjects were healthy, with a wide variety of exercise regimens and body types. The study protocol was approved by the University of California Institutional Review Board for human protection and privacy under Protocol ID 2016-01-8261 (first approved 4 April 2016) and written informed consent was obtained from each subject.

Data Collection

Each experimental trial consisted of a pulsed sequence of isometric elbow flexions, executed while the subject was instrumented with a lab-developed sEMG system [74] and a 3–12 MHz linear ultrasound transducer (L3–12 NGS, eZono AG, Jena, Germany) attached to its corresponding ultrasound unit (eZono 4000, eZono AG, Jena, Germany). Surface EMG electrodes were placed in a differential configuration on both the upper and lower arm, targeting the biceps and brachioradialis, respectively. The ultrasound transducer was placed perpendicular to the lower arm (i.e., roughly perpendicular to the brachioradialis) at a consistent

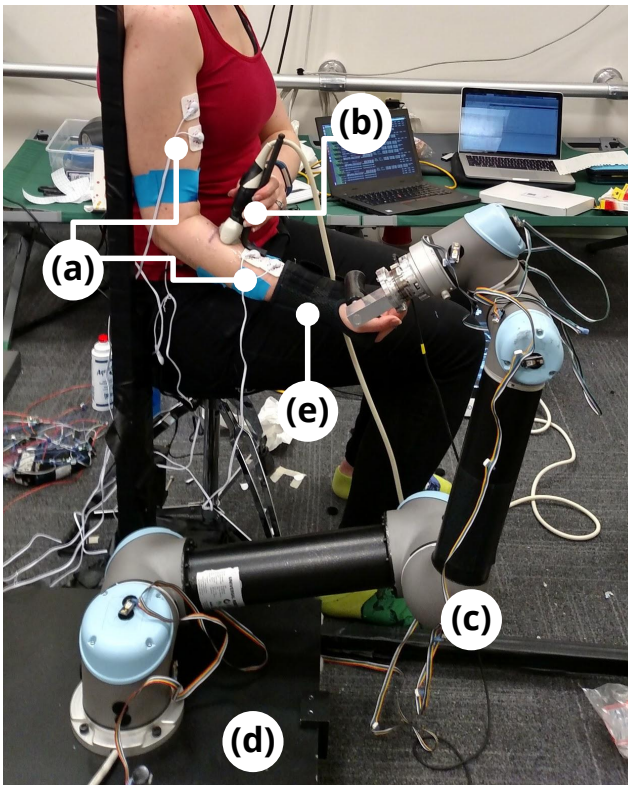

Figure 3.1: Experimental setup for the collection of time series force, surface electromyography (sEMG), and ultrasound data under multiple elbow angles. Setup includes sEMG electrodes (a); ultrasound probe (b); UR5 robot (c) with attached handle, through which subject may transmit force to force plate (d); and wrist brace (e) for elbow isolation. ©2020 IEEE.

marked location and lightly stabilized in the sagittal plane via an L-shaped support. When exerting force, subjects pressed upward on a handle mounted to the end effector of a 6-degree-of-freedom robot arm (UR5, Universal Robots, Odense, Denmark).[2] This robot remained static for each trial, but changing its configuration between trials served as an easy manner of repositioning the handle in space based on subject physiology and desired elbow angle during measurement. The robot was in turn mounted to a 6-channel force plate (Optima-HPS, Advanced Mechanical Technology, Inc., Watertown, MA, USA), which was used to measure the output force exerted by the subject. Additionally, each subject wore a brace to immobilize the wrist.

[2]Note that subjects did not actually grasp this handle, but simply pressed upward via elbow flexion; the handle was used to allow subjects to exert a force vector exactly in the elbow flexion direction (i.e., tangent to the rotational flexion motion, with no radial component) without necessitating precise placement of the flat surface of the robot end effector in space for each elbow angle and subject.

During data collection, the subject sat comfortably upright, feet planted, with the back of the upper arm supported and elbow flexed to the prescribed angle, as shown in Fig. 3.1. Subjects were instructed to press upward in the sagittal plane while only exerting force at the elbow (and otherwise relaxing, to minimize co-contraction) and keeping the forearm completely supinated. For each experimental trial, the subject was guided in a series of 21 force pulses, each of 2 s duration with 1 s rest in between, via a visual display that instructed the subject to press with "low", "medium", or "high" intensity at each interval. The first 3 intervals were always performed in order of increasing intensity, while the remaining 18 were randomized. For select trials, this pulsed sequence was followed by a set of 3 sustained force presses (each 5 s in duration, with 2 s rest in between) in increasing order of intensity. Note that this protocol was intended primarily to generate a wide variety of force conditions, not to prescribe specific force values; none of our analyses rely on subjects following this protocol exactly, and there is high variance in the extent to which this sequence is readily observable in the collected force data.

Data Scope

Using the protocol above, simultaneous force, sEMG, and ultrasound data were collected from a single subject (denoted *Sub1*) at elbow angles of 25°, 44°, 69°, 82°, and 97°, as measured from full extension. Simultaneous force and ultrasound data were collected from an additional 4 subjects (denoted *Sub2*–*Sub5*) at a self-selected, comfortable elbow angle near 69° for comparison with the primary subject's 69° trial.

3.3 Data Processing & Definition of Deformation Signals

Prior to our correlation analyses, we preprocessed the raw data streams in the following manner to extract force, neural activation, and deformation measures that are readily comparable. All preprocessing and analysis details are included with code release, and an exemplar of the finalized trial data can be seen in Table 3.1.

Force and sEMG Preprocessing

To generate a single output force magnitude from the collected 6-channel (i.e., 3-channel force and 3-channel torque) force plate data, the measured wrench was first transformed into the handle's frame of reference based on the robot's kinematic configuration. The magnitude of the subject's output force was then calculated as the magnitude of the force component of the resultant wrench.

To generate muscle activation data from raw sEMG values, ambient noise was first removed with a 60 Hz notch filter. Our final activation measure is the absolute value of this denoised signal, smoothed by an exponential moving average filter. These sEMG values were

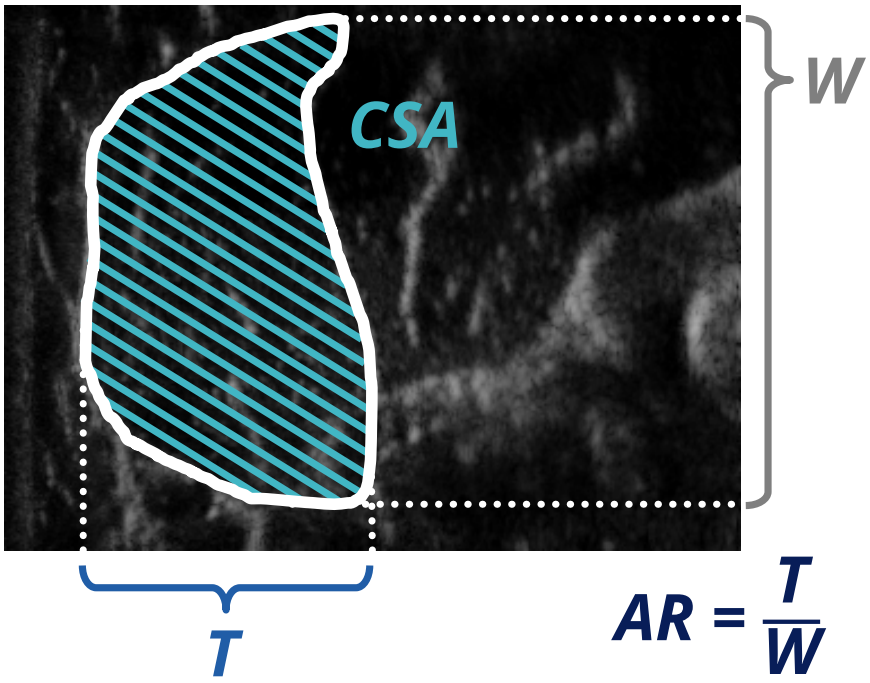

Figure 3.2: Example cross section of the brachioradialis muscle, as collected via ultrasound (probe to left of frame, scanning right), with illustrated shape measures of interest: cross-sectional area (*CSA*, *cyan*), thickness (T, *blue*) and aspect ratio (AR, *dark blue*, the ratio of thickness T to width W). Each measure is calculated for each ultrasound frame during trials of isometric elbow flexion to generate time series deformation data.

collected from both the brachioradialis and the biceps brachii, both of which aid in elbow flexion. We present both in the analyses below — the first because it is more closely related to the observed brachioradialis deformation, and the second because it is a mildly cleaner signal and thus constitutes a more competitive baseline for our comparison with deformation measures.

Extraction of Deformation Measures

Given our time series data of ultrasound images, we can parameterize muscle deformation in a multitude of ways, from shifts in individual pixel values to changes in muscle dimensions. For this preliminary study, we examine the following three shape measures, which are also illustrated in Fig. 3.2.

Cross-sectional area (CSA). We define *CSA* as the total area of the brachioradialis cross section visible on a given ultrasound scan.

Thickness (T). We define T as the maximum extent of the brachioradialis cross section measured from the surface of the skin perpendicularly away from the transducer.

Aspect ratio (AR). We define AR as the ratio between the thickness T of the brachioradialis cross section and its maximum extent in the perpendicular dimension.

Each of these 3 metrics is calculated by first manually masking the brachioradialis cross section at each time point, then extracting each above-defined value from the mask. We note two salient sources of measurement error. First, the manual masking process relies on human annotation that is inconsistent and time consuming. Second, the full extent of the brachioradialis cross section often extends beyond the measured ultrasound frame, impacting primarily the CSA and AR measures. These limitations are addressed in future chapters, through our development of automated cross section tracking methods (chapter 6) and studies of full 3D muscle deformation (parts III and IV), respectively.

Alignment of Data Streams

Due to the technical peculiarities of each sensor, a fully time-synced data collection system remained under development at the time of data collection; we therefore align the force, sEMG, and ultrasound data streams based on the temporal location of the first force pulse, which is manually noted in each data stream, then linearly interpolate each series to attain a consistent frequency. In this preliminary study, we make no claims on the precise temporal relationships between force, activation, and deformation signals. This is one of the first limitations addressed in the refined follow-up correlation study in chapter 4, which uses a more precisely time-synced system.

Ultrasound Drift Compensation

Although the ultrasound probe was stabilized by a support structure during collection, the precise location of the probe relative to the brachioradialis muscle tends to drift over time, primarily due to subjects' skin shifting and the probe rotating relative to the arm. To compensate for this drift, we fit a third-order polynomial to each of the CSA, T, and AR data series, measured only at time points at which force was near zero (i.e., between prescribed pulses). We then subtract the value of this polynomial at each point in the original series to generate "detrended" data. We examine both the original deformation data streams and their detrended counterparts in the analyses below.

As with data alignment, this issue is addressed in the refined chapter 4 study, which employs improved ultrasound attachment hardware.

3.4 Correlation & Evaluation

In the following analyses, we use the Pearson correlation to assess the viability of using our candidate deformation measures to infer output force, alongside or as an alternative to

Table 3.1: Example Time Series Data

SIGNAL	Symbol	Pearson Correlation Coefficient raw, *detrended* $CC(\cdot, \mathrm{f})$	Time Series Signal
force (*N*)	f	1.00	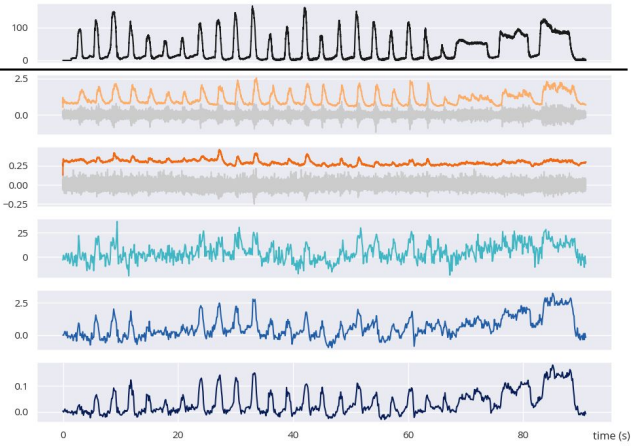
sEMG, biceps (*mV* ***raw,*** *0.1mV* ***processed***)	sEMG-BIC	0.38	
sEMG, brachioradialis (*mV* ***raw,*** *0.1mV* ***processed***)	sEMG-BRD	0.27	
cross-sectional area (mm^2)	CSA(-DT)	0.44, *0.48*	
thickness (*mm*)	T(-DT)	0.59, *0.72*	
aspect ratio	AR(-DT)	0.70, *0.78*	

Example time series data (collected from *Sub1* at 25° elbow flexion) of elbow output force (*black*) in which the subject executed a series of prescribed force pulses and sustained exertions. Data streams include biceps (*light orange*) and brachioradialis (*dark orange*) surface electromyography (*sEMG*) data, as well as detrended ultrasound-measured cross-sectional area (*CSA*, *cyan*), thickness (*T*, *blue*) and aspect ratio (*AR*, *dark blue*) of the brachioradialis cross section. In general, *CSA*, *T*, and *AR* deformation measures correlate well with output force, especially when detrended. Moreover, these force–deformation correlations are of comparable strength to those between force and more standard *sEMG* measures. Note that raw (*gray*) and processed (*orange*) *sEMG* traces are plotted at different scales for clarity as noted, and reported correlation coefficients are for processed data. ©2020 IEEE.

sEMG. We first examine an illustrative time series, then discuss how our assertions translate across changes in kinematic configuration and across subjects.

An Illustrative Time Series

Table 3.1 shows representative pulse data from a single trial (specifically, that of *Sub1* collected at 25°) alongside each data stream's correlation coefficient as measured against force data. In this series, and in general, our deformation measures *CSA*, *T*, and *AR* correlate comparably with sEMG from both the biceps and brachioradialis; in this case, in fact, they are substantially more correlated (though this varies with both elbow angle and subject, as discussed below).

This exemplar also illustrates two other trends largely consistent throughout the data.

First, thickness T and aspect ratio AR are somewhat more correlated with output force than cross-sectional area CSA, perhaps because an increase in one dimension is often accompanied by a decrease in the other due to muscles' isovolumetric nature. Second, detrending the data using the method detailed in section 3.3 generally improves the observed correlation, providing evidence that sensor drift (or some biological source of drift, e.g., fatigue) indeed occurred.

Deformation and Elbow Angle

Fig. 3.3 illustrates the manner in which correlation strength varies with angle for each data stream on our exemplar subject, *Sub1*. For flexion angles less than approximately 90°, deformation measures remain reasonably correlated with force, with detrended data performing slightly better as in our illustrative data series. The relative performance of CSA, T, and AR measures, however, varies in a complex manner that is difficult to characterize but is reflected qualitatively in the accompanying illustrative frames, which show significant angle-dependent variation in shape and size. This variation is perhaps most obvious in the highest-flexion 97° data series, in which the cross section changes shape entirely and our defined deformation metrics largely fail, perhaps because the belly of the muscle has shifted significantly away from the ultrasound probe location. Interestingly, sEMG correlations behave in the opposite manner, with the largest correlations occurring at the highest flexion angles. While this may simply be a subject-specific phenomenon, it suggests that hybrid sensor approaches — in which ultrasound-measured deformation and sEMG-measured activation are employed, depending on kinematic configuration — may hold more promise for robust device control than either technology alone.

Multi-Subject Deformation

Fig. 3.4 shows the strength of each deformation signal's correlation with force on a representative trial (near elbow angle 69°) for each of our 5 subjects. All measures CSA, T, and AR are consistently correlated with a magnitude of around 0.5 or higher even though the subjects vary significantly in terms of muscle morphology, suggesting that the underlying biological mechanism generating the signal is common to all subjects. As before, detrending the data generally provides a mild performance boost if a significant trend was removed, or makes little difference if no trend was observed. (The one exception, in which detrending resulted in decreased correlation, is *Sub4*, for whom a truly meaningful trendline could not be computed due to lack of zero-force baseline data throughout the data series.)

Perhaps the most interesting subject data is that of *Sub5*, our single elderly subject and the only subject physically small enough that virtually the entire brachioradialis cross section was visible in every data frame. Like other subjects' data, the magnitude of each deformation signal's correlation was substantial — in fact, T and AR signals were more correlated with force than for any other subject, perhaps due to the completeness of the cross-sectional image. Oddly, the subject's CSA data, while strongly correlated with force, was negatively

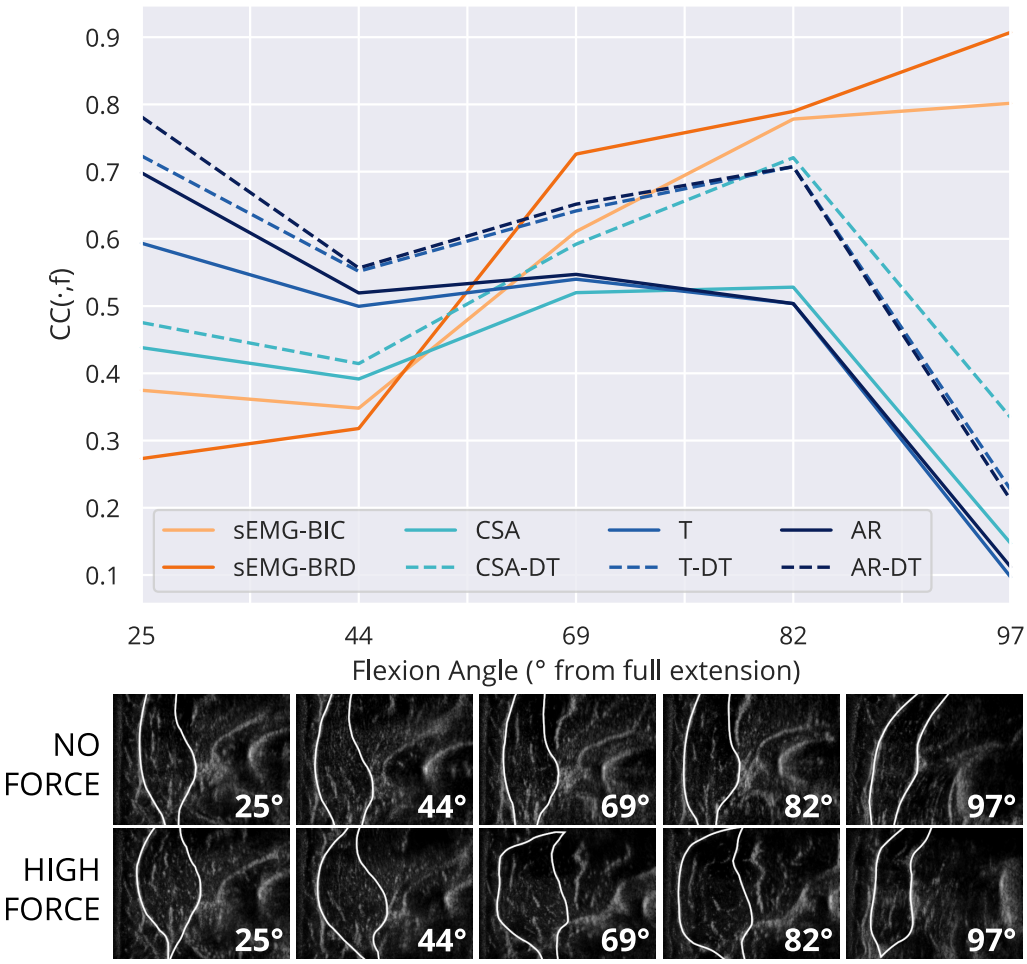

Figure 3.3: *Top*: Correlation of cross-sectional area (*CSA*, *cyan*), thickness (*T*, *blue*) and aspect ratio (*AR*, *dark blue*) of the brachioradialis cross section, both raw (*solid*) and detrended (*dashed*), with elbow output force, alongside baseline force correlations with biceps (*light orange*) and brachioradialis (*dark orange*) surface electromyography (*sEMG*) data. Correlations were computed across flexion angles ranging from near full extension (25°) to near maximum flexion (97°). *CSA*, *T*, and *AR* deformation signals, especially when detrended, correlate well with elbow output force for most elbow angles, but this correlation collapses near full flexion. *Bottom*: Example ultrasound frames with annotated brachioradialis contours depicting no force (*top row*) and high output force (*bottom row*) at each examined flexion angle reflect the changing presentation of muscle deformation with changes in elbow angle. ©2020 IEEE.

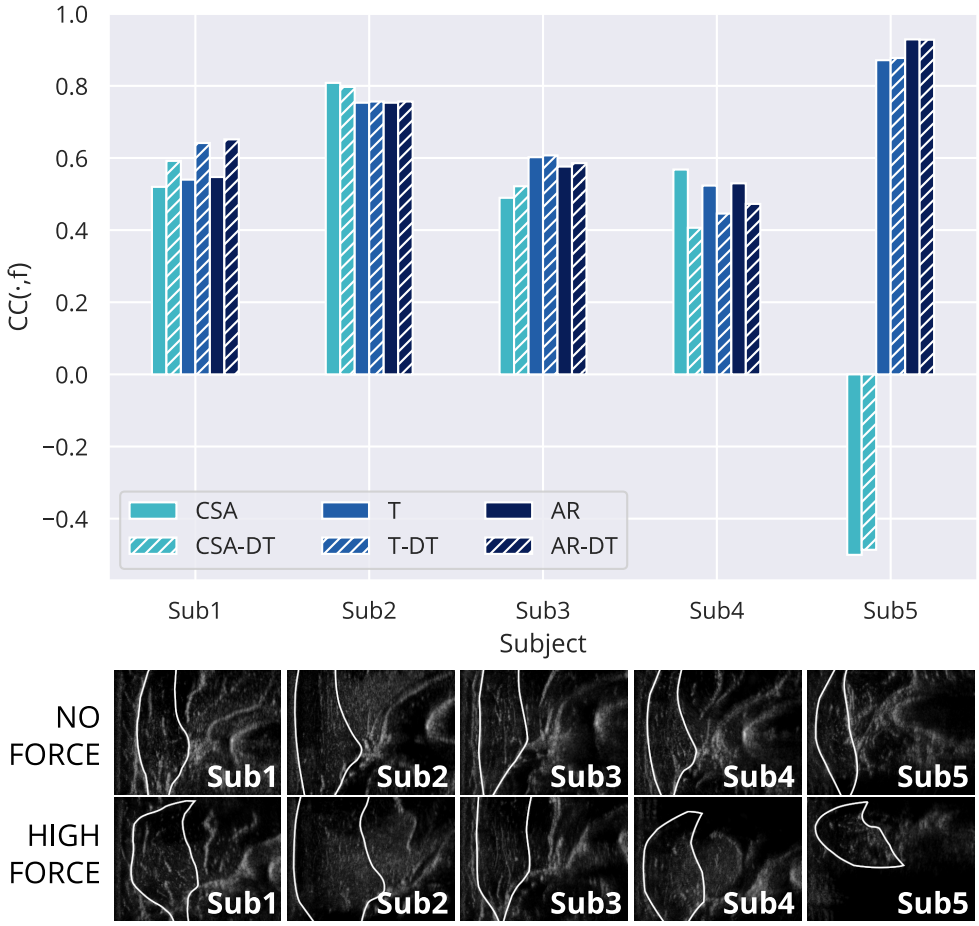

Figure 3.4: *Top*: Correlation of cross-sectional area (*CSA*, *cyan*), thickness (*T*, *blue*) and aspect ratio (*AR*, *dark blue*) of the brachioradialis cross section, both raw (*solid*) and detrended (*hashed*), with elbow output force, across various subjects, collected at ∼69° elbow flexion. *CSA*, *T*, and *AR* deformation signals, especially when detrended, correlate well with output force, though the magnitude and direction of these correlations varies widely across subjects. *Bottom*: Example ultrasound frames with annotated brachioradialis contours depicting no force (*top row*) and high output force (*bottom row*) for each subject reflect substantial morphological variation. ©2020 IEEE.

correlated (i.e., *CSA* decreased with the application of force), in what could be a quirk of morphology, a function of sensor placement, or a property of aging musculature. We see this question in particular as meriting future study, as aging individuals could especially benefit from advances in muscle force measurement technology. We present this preliminary data as evidence both that measuring deformation data in elderly subjects is possible and that the data collected is largely consistent with that of younger subjects in terms of correlation magnitude.

3.5 Summary & Limitations

This study provides novel evidence that several simple measures of ultrasound-measured muscle deformation correlate well with output force and are consistent with sEMG-measured activation at the elbow joint during isometric contraction, showing initial feasibility of leveraging such signals for device control. In the next few chapters, we expand on these positive results, first building a real-time deformation tracking system (with improved sensing and hardware) to expand the correlation analyses above to a wider cohort of subjects and refined set of force conditions (chapter 4), then employing this tracked signal for control (chapter 5).

Chapter 4

Refined Force–Deformation Correlation

This chapter is adapted in part from [50].[1]

In the previous chapter, we showed preliminary evidence that simple measures of deformation, as measured via ultrasound, correlate with output force, and thus represent promising device control signals; however, these analyses were limited by the need for extensive manual annotation and disparate sensing infrastructure. In this chapter, we present a novel experimental platform, consisting of both networked hardware and real-time optical-flow-based muscle contour tracking software, that enables the simultaneous, time-synced collection and display of joint force, muscle activation, and muscle deformation data, facilitating force–deformation–activation correlation analysis at the elbow on an expanded cohort of subjects. (This platform is also employed in the deformation-based control studies in chapter 5.)

In the sections below, we outline the scope of this refined study and detail this novel data collection platform, then present an enhanced correlation analysis enabled by this new system, and lastly, discuss both the platform's limitations and its direct applicability to the study of ultrasound-based control. As before, we find that simple measures of deformation correlate well with output force, and newly, that we can successfully track one measure of this deformation (i.e., muscle thickness) in real-time, enabling its use as a control signal in future chapters.

All data and analysis code have been released to the OpenArm project in the OpenArm Multisensor 2.0 package (see section 11.2).

[1] ©2021 IEEE. Reprinted, with permission, from Laura A. Hallock, Bhavna Sud, Chris Mitchell, Eric Hu, Fayyaz Ahamed, Akash Velu, Amanda Schwartz, and Ruzena Bajcsy. "Toward Real-Time Muscle Force Inference and Device Control via Optical-Flow-Tracked Muscle Deformation". In: *IEEE Transactions on Neural Systems and Rehabilitation Engineering (TNSRE)* (2021, under review).

4.1 Study Scope

As with our preliminary study in chapter 3, we target the brachioradialis cross section as a source of the deformation signal and examine its relationship to output force and surface electromyography (sEMG)-measured activation at the elbow during isometric contraction, this time during precise force trajectories enabled by real-time visual feedback. Based on these previous correlation results, as well as the tracking capabilities of our novel system, we focus specifically on thickness change as our measure of deformation, and use sEMG of only the biceps brachii as our activation for comparison. We also examine these signals at only a single kinematic configuration, as we explore kinematic-associated deformation in our 3D analyses in future chapters (parts III and IV).

4.2 Real-Time Deformation & Activation Tracking Platform

To enable improved evaluation of force–deformation correlation during isometric contraction at the elbow, we developed a novel experimental system to enable simultaneous collection and optional real-time display of joint force, muscle thickness (as measured via ultrasound), and muscle activation (as measured via sEMG) during varied isometric elbow flexion.[2]

This section details, first, hardware aspects of this system that allow for collection of these data under consistent kinematic conditions, and second, the signal processing software used to extract, calibrate, display, and record these signals over time.

Hardware Setup

The hardware platform was designed for data collection from the right arm of a subject seated comfortably upright, feet planted, right upper arm comfortably adducted (vertical), elbow flexed 90°, forearm fully supinated, with the elbow supported by a static jig from below, as shown in Fig. 4.1. The right wrist was firmly strapped into a brace mounted to a 6-channel force-torque sensor (ATI Mini45, ATI Industrial Automation, Apex, NC, USA), which was in turn mounted to the end effector of a 7-degree-of-freedom robot arm (KUKA LBR iiwa 14 R820, KUKA AG, Augsburg, Germany).[3] Subjects pressed upward on this sensor to generate a measure of "output joint force" at the elbow; while this measured contact force was inherently confounded by contributions from other linked joints, subjects were instructed to exert force using only elbow flexion motion, and the wrist was immobilized, with the force

[2]Note that the software aspects of this platform — contained in the OpenArm open-source release — are agnostic to the particular muscles and joints observed and could readily be adapted for study of other joints.

[3]Note that this robot remained static throughout all data collection, but changing its configuration between subjects served as an easy manner of re-placing the force-torque sensor in space to maintain consistent elbow angle across differing subject physiology.

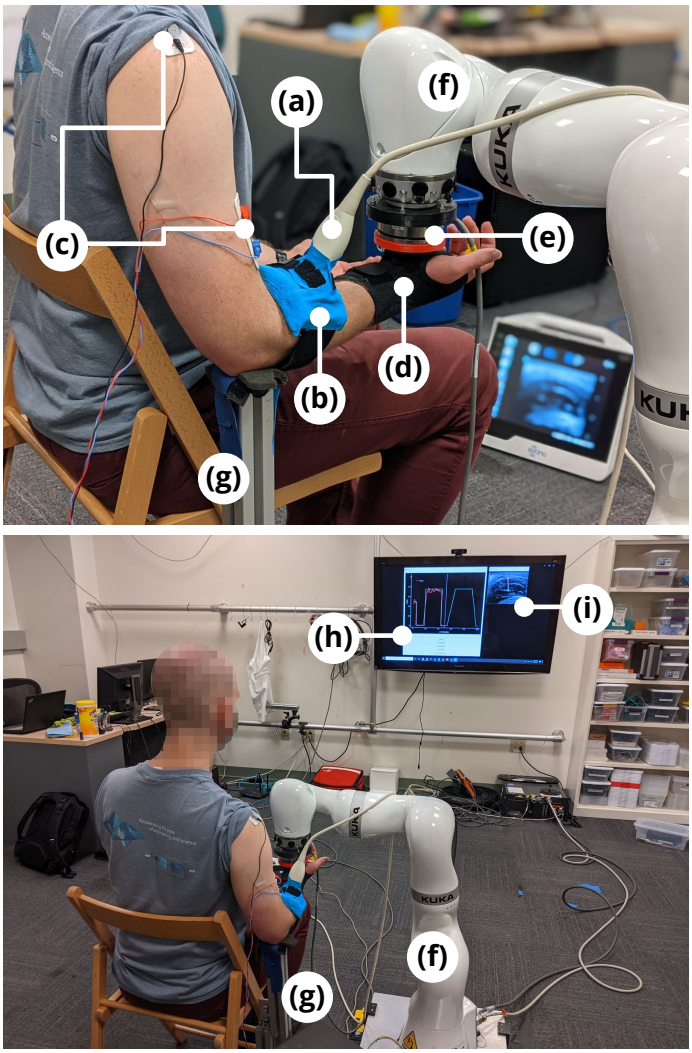

Figure 4.1: Experimental setup for the collection of time series force, ultrasound, and surface electromyography (sEMG) data during constrained isometric elbow flexion. Setup includes ultrasound probe (a) attached securely to user's forearm with cuff (b); sEMG electrodes (c); wrist brace (d) through which subject transmits force to attached to force-torque sensor (e), in turn held stable by KUKA robot (f); elbow stabilizing jig (g); goal and sensor trajectory display (h) for real time visual feedback for subject self-assessment; and real-time ultrasound thickness tracking data (i) for continuous experimenter system status monitoring. This system allows subjects to precisely follow a specified force trajectory to enable study of force–deformation correlation under varied trajectory types (section 4.3) and to perform trajectory tracking tasks using experimental deformation- and activation-based signals (chapter 5). ©2021 IEEE.

sensor attached near the proximal edge of the palm, to isolate the elbow joint as much as possible.

To gather muscle activation and deformation data, the subject's right arm was instrumented with a MyoWare sEMG system (Advancer Technologies, Raleigh, NC, USA) and a 3–12 MHz linear ultrasound transducer (L3–12 NGS, eZono AG, Jena, Germany) attached to its corresponding ultrasound unit (eZono 4000, eZono AG, Jena, Germany). Surface EMG electrodes (Red Dot 2560, 3M, St. Paul, MN, USA) were placed in a differential configuration on the biceps brachii, with the two signaling electrodes placed with 40 mm proximal–distal center-to-center separation (adhesive edges abutting) and the top electrode roughly centered on the lateral belly of the muscle, and the grounding electrode placed on the acromion. The ultrasound transducer was manually placed perpendicular to the lower arm such that the brachioradialis cross section was maximally in frame and held in place by an adjustable foam and neoprene cuff.

Note that ultrasound and sEMG sensors were placed to target different muscles: the brachioradialis and the biceps, respectively. While this inherently limits any correlation insights we might make between deformation and activation values, both are active during elbow flexion, and this configuration was chosen to allow simultaneous collection while preventing sensors from physically interfering with one another. Earlier work (chapter 3) also found stronger sEMG signals from the biceps than the brachioradialis, such that this configuration provides a more competitive baseline against which to evaluate deformation data.

Signal Tracking & Display

During data collection, the subject faced a large monitor, which displayed two or more signal streams: a time series goal trajectory, and either force (to enable consistent force output for the correlation analysis described in this chapter), and/or ultrasound-extracted deformation/thickness or sEMG-extracted activation (to compare deformation- and activation-based trajectory tracking in the study discussed in chapter 5), each normalized to the subject's strength capabilities (via calibration to their maximum voluntary contraction force, as detailed below). Our methods for signal extraction, calibration, and display are outlined below, and full details can be found in the released codebase.

Tracking Muscle Thickness via Optical Flow

Brachioradialis thickness was tracked over time via the standard iterative Lucas–Kanade method of optical flow estimation [96] as implemented in the OpenCV Python library [14]. Specifically, at the start of each trial, while the subject was instructed to remain still, 10 points were manually selected along both the top and bottom (i.e., superficial and deep) surfaces of the brachioradialis muscle fascia, forming two clusters of points. These points were used to define the vertices of two polygons from which contours were extracted [138]; these contour points, as shown in Fig. 4.2, were then tracked over time (on a bilaterally-

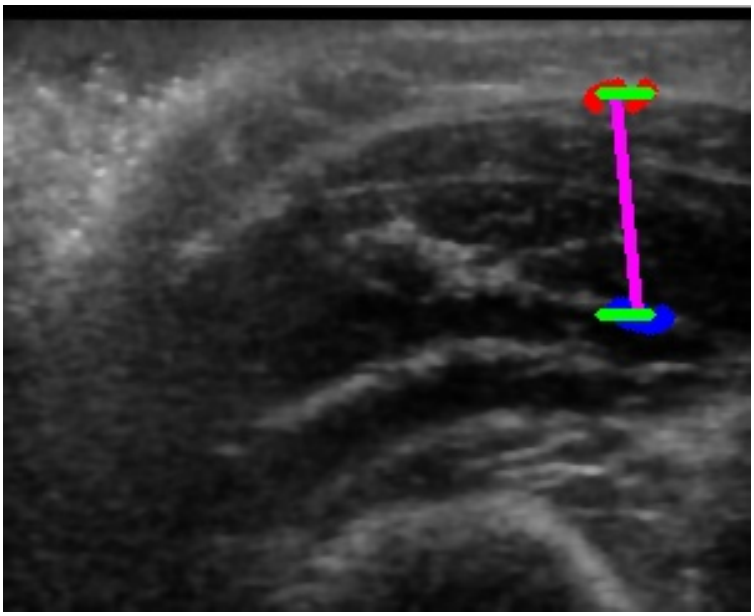

Figure 4.2: Still frame of optical-flow-based brachioradialis thickness tracking system. Points were tracked along the superficial (*red*) and deep (*blue*) fascial surfaces of the brachioradialis, and thickness was reported as the vertical (superficial-to-deep, *green*-to-*green*) distance between the center (mean) of each cluster. A line connecting each cluster center was also displayed (*magenta*) to allow for easy observation of particle drift. ©2021 IEEE.

filtered version of each image [141] to suppress speckle noise) at a best-effort frame rate[4] of 1 kHz, and the thickness value calculated at each iteration as the vertical (superficial-to-deep) distance between the mean location of tracked points in each cluster.

To prevent points from drifting away from the selected surfaces (a frequent challenge in optical flow estimation), we leveraged the knowledge that the fascial segment selected by each point cluster should remain intact — i.e., points within a cluster should remain at similar positions relative to each other. Thus, during each iteration, if the average squared distance from tracked points to their cluster center exceeded a specified distance (here, 200 px^2, or approximately 3 mm^2), all tracked point locations were reset to their initial locations. In practice, these resets happened rarely for most subjects and are noted in released data.

Signal Processing & Calibration

Force, ultrasound-extracted thickness, and sEMG-measured activation were sampled for both recording and display at a best-effort rate of 1 kHz. Activation values were smoothed via a 250-point (0.25 s) moving average filter.

[4]In practice, this frame rate was slightly slower due to delay associated with image recording.

To normalize force, deformation, and activation signal traces $s(t)$ to a subject's strength, at the beginning of each trial, subjects were first instructed to remain still for several seconds, then to press upward with maximum possible contraction force for several seconds; mean minimum and maximum signal values were then calculated over the final 200 samples (approximately 0.2 s) at each condition as $s_{\min}$ and $s_{\max}$, respectively.

Normalized signal traces $\bar{s}(t)$ were then calculated for both display and recording as

$$\bar{s}(t) = \frac{s(t) - s_{\min}}{s_{\max} - s_{\min}}.$$

During each trial, one or more of these normalized traces $\bar{s}(t)$ was then displayed alongside a goal trajectory, which the subject was then instructed to track by modulating elbow flexion force, as detailed in sections 4.3 and 5.2. Note that regardless of display, all goal, force, deformation, and activation signals, both raw and processed, were collected for all trials.

4.3 Data Set Collection

As a refined exploratory data set, simultaneous force, deformation, and activation signals were collected from a preliminary cohort of subjects using the platform described above as they tracked a specified elbow flexion force trajectory with visual feedback. We first outline our subject cohort and collection procedure, then present preliminary time series data alongside qualitative and quantitative analysis of the force–deformation relationship.

Subject Biometric Data & Consent

Data were collected from the right arm of 10 subjects (7 female, 3 male, 9 right-handed, 1 left-handed, age 25.6 ± 0.966, mass 61.7 ± 10.5 kg, height 1.69 ± 0.072 m, body mass index 21.5 ± 2.89), hereafter denoted *Sub1–Sub10*.[5] All subjects were healthy, with a wide variety of exercise regimes, body types, and familiarity with nonstandard computer interfaces. The study protocol was approved by the University of California Institutional Review Board for human protection and privacy under Protocol ID 2016-01-8261 (first approved 4 April 2016) and written informed consent was obtained from each subject.

Trial Specification

After being strapped into the data collection system outlined in section 4.2 and instrumented with all relevant sensors, each subject performed three 90 s tracking trials: one unstructured trial to familiarize them with the system, followed by two trials intended for correlation analysis.

[5]Statistics are reported as mean ± standard deviation. For additional demographic data, broken down by subject, see the full open-source data release. Note that the subjects examined here and their identifying numbers bear no relation to those examined in the previous chapter 3 study.

The initial familiarization trial — from which no actual data was analyzed for publication — was designed to both familiarize subjects with the system and assure investigators that all sensors were behaving as expected. After initializing the thickness tracking system and performing min/max calibration (as outlined in section 4.2), the subject's monitor was set to display all three normalized force, thickness, and activation traces, as well as the goal trajectory to be used in future trials. Subjects were then instructed to freely modulate elbow flexion to get a sense for how the various signals changed with force output, though they were not yet asked to perform the tracking task.

During the two data collection trials, the same tracking initialization and min/max calibration was performed, but this time, the subject's monitor displayed only the goal and normalized force trajectories. The subject was instructed to match the force trace to the goal trace — scaled to the subject's force generation capability and detailed below — by modulating isometric elbow flexion force. The second trial was used in correlation analysis, except where noted in the corresponding data release (e.g., due to sensor malfunction).

Goal Trajectory Elements

To evaluate correlation across various types of force exertion, both sustained and quickly varying, the same 90 s goal trajectory containing the following varied elements and shown in Fig. 4.3 was used for all trials. Specifically, the trajectory contained the following elements in sequence:

- *sustained* flexion at 0.25, 0.5, and 0.75 of maximum force capability, each for 8 s, with intervening 4 s rest phases;
- a slow 6 s *ramp* in flexion from 0 to 0.75 of maximum force capability, 5 s sustained flexion at this 0.75 level, then a slow 6 s ramp back to 0;
- arbitrary, quickly varying 2 s *step* changes in flexion (specifically, to 0.3, 0.5, 0.2, 0.8, 0.4, and 0.6 of maximum force capability); and
- a 0.3 Hz, 10 s *sine* wave ranging from 0.25 to 0.75 of maximum force capability.

Force–deformation and force–activation correlation were evaluated across both the entire trajectory and individual elements, as detailed below.

4.4 Correlation & Evaluation

In the following analyses, we use the Pearson correlation, applied directly to the synchronized data streams, to assess the viability of using our candidate deformation measures to infer output force, alongside or as an alternative to sEMG. We first examine an illustrative time series, then discuss how our assertions translate across across subjects and trajectory types.

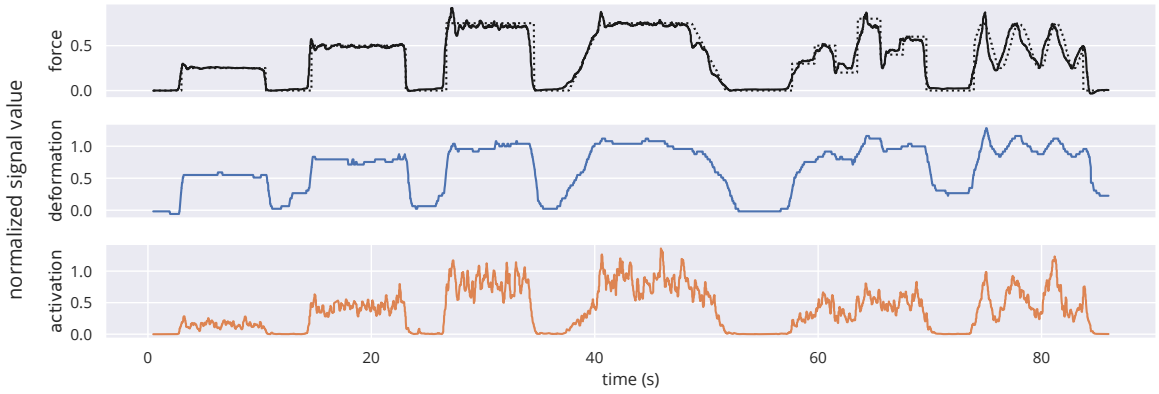

Figure 4.3: Example time series data collected from subject *Sub1* for force–deformation and force–activation correlation analysis, including output force (*black*, *solid*), alongside specified goal trajectory (*black*, *dotted*), deformation (i.e., brachioradialis thickness change as tracked via optical flow, *blue*), and activation (as measured via sEMG, *orange*). Subjects were able to track the specified force trajectory with little error, enabling controlled observation of a variety of sustained and quickly varying force conditions, and both deformation and activation were shown to be highly correlated with output force during all portions of the trajectory. Signal values are reported — as they were displayed — as a fraction of measured maximal value, as described in section 4.2. ©2021 IEEE.

An Illustrative Time Series

Fig. 4.3 shows representative trajectory data from a single trial (specifically, that of *Sub1*). In this series, and in general, brachioradialis thickness deformation correlates comparably with sEMG activation (though this varies with both subject and trajectory type, as discussed below).

This exemplar also illustrates two important data set qualities: first, that subjects proved impressively skilled at following goal trajectories given visual feedback, yielding both the sustained and quickly varying signal types we sought for comparison; and second, that while our optical-flow-based tracking system exhibited some drift over time, thickness was largely well-tracked for most subjects throughout the full 90 s duration of each trial.

Correlation by Subject

Fig. 4.4 shows the strength of deformation and activation correlation with force for each of our 10 subjects. Both deformation and activation are consistently correlated, with most subjects showing a correlation magnitude of around 0.6 or higher — an even higher magnitude than that found in our chapter 3 study — even though subjects vary significantly in terms of muscle morphology, as illustrated in Fig. 4.5, suggesting (as in chapter 3) a common underlying biological mechanism.

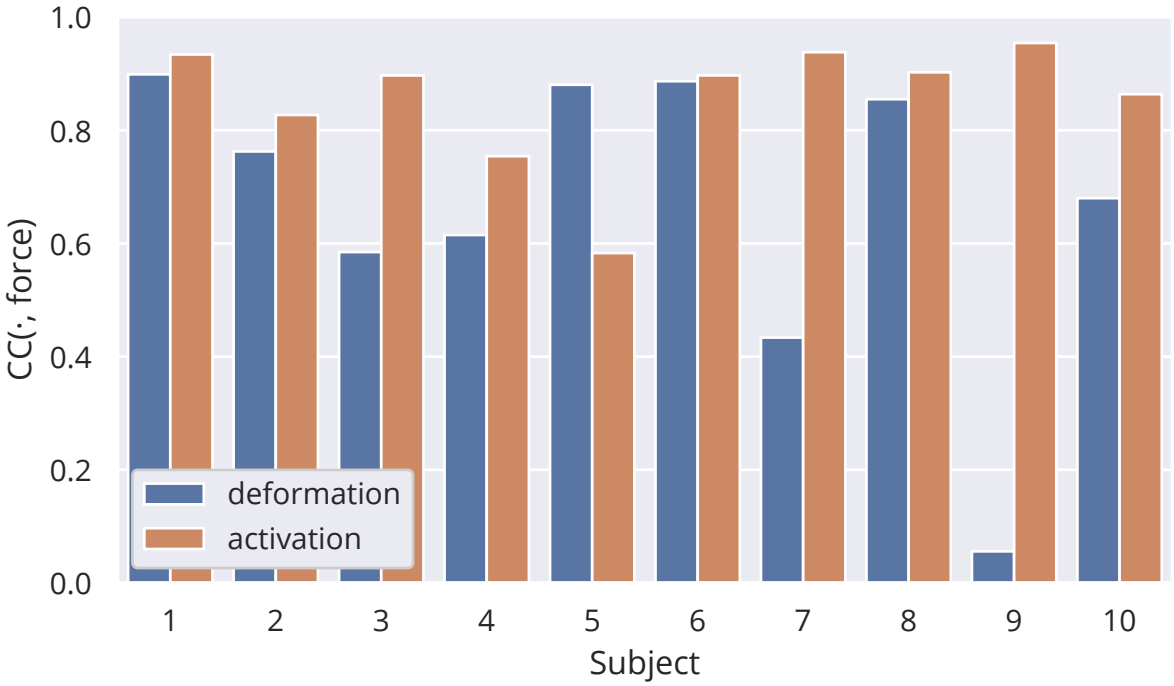

Figure 4.4: Correlation of muscle deformation (*blue*) and activation (*orange*) signals with elbow output force across all subjects. Despite substantial differences in morphology (illustrated below in Fig. 4.5), most subjects — aside from *Sub7* and *Sub9*, who displayed morphological quirks that resulted in poor signal quality, as discussed in section 4.4 — showed strong correlation between deformation and output force. ©2021 IEEE.

At the same time, two subjects — *Sub7* and *Sub9* — fail to show the same thickness–force correlation. While more principled analysis is needed to tease out the exact reasons for this lack of correlation, experimenters noted during collection that the observed deformation appeared qualitatively different, as illustrated in Fig. 4.5: unlike most subjects, for which vertical (deep–superficial) expansion was observed, widening the brachioradialis contour, these subjects showed substantial lateral motion, in which fibers appeared to slide side to side, but the deep and superficial fascial surfaces seemed to move very little. Given that this different motion paradigm appeared in multiple subjects, and could be the result of idiosyncratic morphology, a function of sensor placement, or differing elbow flexion strategy, we see other deformation measures (e.g., localized fiber motion in any direction) as worthy of future study.

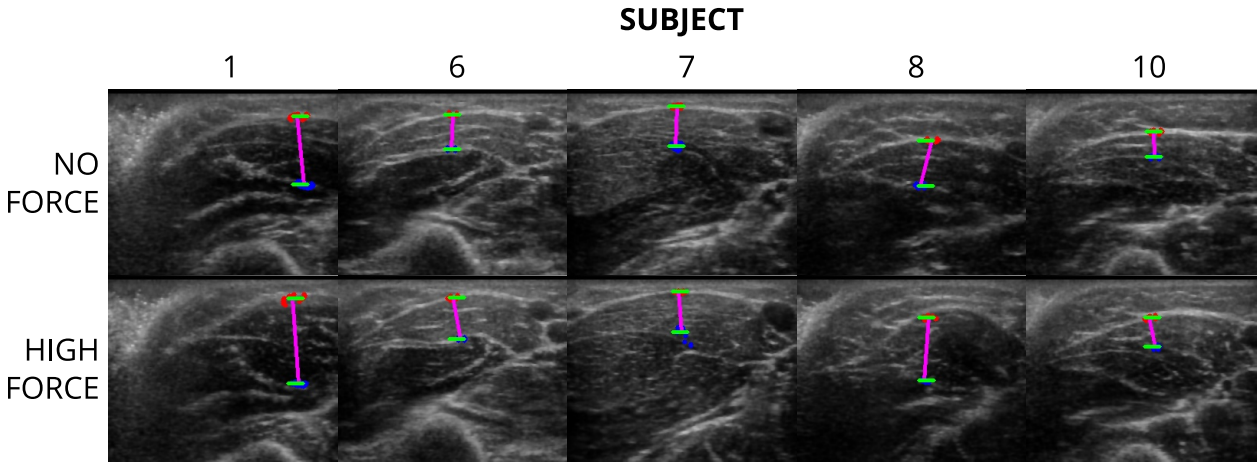

Figure 4.5: Example ultrasound frames from an illustrative subset of subjects, with tracked and annotated brachioradialis thickness, depicting no force (*top row*) and high output force (*bottom row*) for each subject. While subjects' morphology varies significantly, most subjects (like the pictured *Sub1*, *Sub6*, *Sub8*, and *Sub10*) display a reliable thickness increase with output force, while several (like *Sub7*) primarily display lateral motion, leaving thickness uncorrelated with force output. ©2021 IEEE.

Correlation by Trajectory Type

Generally, as shown in Fig. 4.6, both deformation and activation correlate well with force across all four trajectory types (*sustained*, *ramp*, *step*, and *sine*, as outlined above in section 4.3), with activation showing indistinguishable, and high, levels of correlation across all trajectory types. Force–deformation correlation was particularly high for *ramp* and *step* conditions, though lower and more variable for *sustained* and *sine* conditions. Based on qualitative examination of the data, we theorize that these lower values are the result of limitations in our tracking software, rather than an underlying physiological mechanism: specifically, we observe first, that under quick, dramatic force/thickness changes like those at the start and end of each trial, the optically tracked points fail to fully remain on the fascia, and thus do not return to baseline, and second, that these points generally drift over time, perhaps impacting the final (*sine*) section of the trajectory most dramatically.

Nevertheless, deformation — even as measured by this limited, drifting, proof-of-concept system — is consistently and highly correlated with output force, showing promise for use in device control or motion analysis under both fast and slow movement conditions.

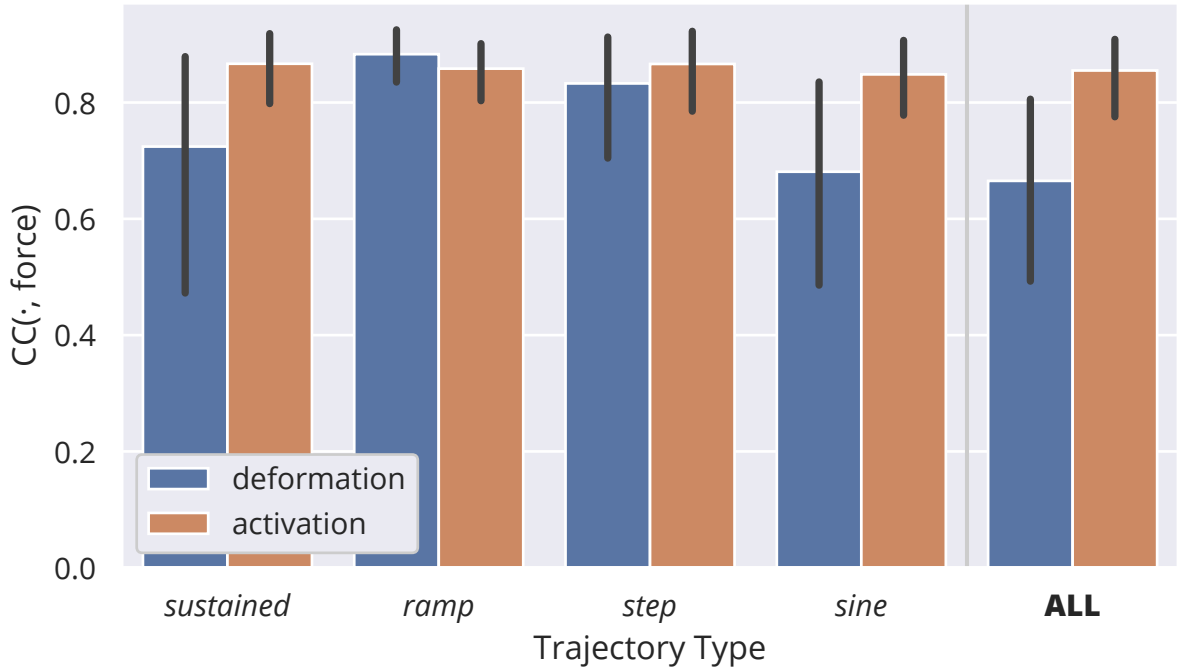

Figure 4.6: Correlation of muscle deformation (*blue*) and activation (*orange*) signals with elbow output force across various trajectory types and in aggregate, with noted standard deviation across subjects. Deformation remained strongly correlated with output force for all examined trajectory types, with slightly lower and more variable correlation during *sustained* and *sine* trajectories that is likely the result of limitations in the optical flow tracking system, as discussed in section 4.4. ©2021 IEEE.

4.5 Summary & Limitations

The results in this chapter validate and expand the preliminary correlation results in chapter 3 — i.e., they corroborate that brachioradialis muscle deformation (thickness change) correlates with output force at the elbow during varied isometric contraction and is consistent with simultaneous sEMG data under our refined experimental conditions.[6] These results — and the novel platform used to collect them — provide the basis for the proof-of-concept deformation-based control explored in chapter 5.

At the same time, these correlations were analyzed under limited conditions, and the cohort size remains small; in the future (though beyond the scope of this thesis), we aim

[6]In fact, as we expand our analyses beyond isometric conditions (e.g., to include joint motion), we expect the deformation–force correlation to surpass that of activation, as the signal measured by sEMG electrodes will likely be more heavily impacted by the relative motion of skin and muscle than that measured by ultrasound.

to further expand these analyses with additional subjects (of varying age and ability), additional force conditions (e.g., non-isometric, natural/unconstrained motion), and an improved (e.g., multi-channel) sEMG system for more equitable ultrasound comparison (especially important as we begin to examine dynamic motions, during which electrodes placed on the skin may slide relative to underlying structures). Such enhancements will allow for expanded understanding of the results above (including the impact of age on the deformation signal) and of phenomena not yet explored (e.g., temporal and spatial sEMG–deformation relationships, the impact of fatigue).

A further limitation of this analysis is our assumption that force, activation, and deformation signals occurred simultaneously: we might have achieved better correlation by accounting for the multi-millisecond electromechanical delay expected between sEMG-measured activation and output force [39]. We aim to incorporate this delay into future correlation and modeling analyses (and even leverage our data to study this delay, which is variable and remains poorly characterized).

Finally, this study also illustrates the challenge of extracting generalizable force–deformation relationships: as shown visually in Fig. 4.5's illustrative frames, a single muscle cross section, without accompanying 3D shape data, is difficult or impossible to interpret, and simple deformation measures like those analyzed above barely scratch the surface of its architectural nuance. Future chapters of this thesis address this challenge from two perspectives, through development of tracking systems that support real-time extraction of new and higher-dimensional signals from a given muscle cross section (chapter 6) and the construction and analysis of comprehensive 3D deformation data sets (parts III and IV).

Part II

Tracking & Control

Chapter 5

Trajectory Following & Control

This chapter is adapted in part from [50].[1]

In part I, we showed that several measures of muscle deformation correlate well with ouput force, and that we can track one measure (thickness change) in real time; now, in part II, we explore the use of these signals for device control, leveraging our muscle thickness tracker in a prototype trajectory tracking task and examining possible expansions to the tracking system to enable extraction of richer control signals.

In this chapter, we present a proof-of-concept case study in which subjects completed a trajectory tracking task by modulating either ultrasound-measured deformation or surface electromyography (sEMG)-measured activation. In the sections below, we outline our study scope and data collection procedure, then examine subjects' tracking performance and reported preferences, and lastly, discuss study limitations and applications to physical device control. We find that ultrasound-measured thickness can indeed be used for real-time trajectory tracking, and is even, in some cases, preferred by the user over sEMG-based control.

All data and analysis code have been released to the OpenArm project in the OpenArm Multisensor 2.0 package (see section 11.2).

5.1 Study Scope

Drawing on our successful correlation analyses in part I, we use ultrasound-measured brachioradialis thickness change and sEMG-measured biceps brachii activation — measured using the collection system described in section 4.2 — as two candidate control signals with which a cohort of subjects performed a one-dimensional trajectory tracking task using isometric elbow flexion exertion.

[1] ©2021 IEEE. Reprinted, with permission, from Laura A. Hallock, Bhavna Sud, Chris Mitchell, Eric Hu, Fayyaz Ahamed, Akash Velu, Amanda Schwartz, and Ruzena Bajcsy. "Toward Real-Time Muscle Force Inference and Device Control via Optical-Flow-Tracked Muscle Deformation". In: *IEEE Transactions on Neural Systems and Rehabilitation Engineering (TNSRE)* (2021, under review).

5.2 Data Collection

The same cohort of 10 subjects for which refined correlation analysis was performed in chapter 4 (*Sub1*–*Sub10*, detailed in section 4.3) was asked to perform several trajectory tracking tasks using the chapter 4 force–deformation–activation tracking system (section 4.2, Fig. 4.1). Specifically, for each trial, the subject was instructed to perform the same trajectory tracking task described in section 4.3, with identical tracking initialization and calibration steps, except that instead of the ground truth force trajectory, the monitor displayed only the normalized thickness/deformation or normalized and mean-filtered activation trace, alongside the same goal trace used in chapter 4 correlation analysis. These two tracking schemes (identified to the subject as "mode 1" and "mode 2") were presented in a random order, and each subject was only informed that each controller would be using some combination of sEMG and ultrasound data, with no additional details.[2] The subject executed the full tracking task using each signal twice, and the second trial was used in tracking performance evaluation, except where noted in the corresponding data release (e.g., due to sensor malfunction).

Each subject was then asked to complete a questionnaire, with both Likert scale and free-form response elements, to evaluate their controller preferences.

5.3 Quantitative Tracking Performance

In the following analyses, we evaluate tracking performance of an exemplar subject, as well as across subjects and trajectory types.

Illustrative Trajectories

Fig. 5.1 shows the time series tracking performance of subject *Sub1* using each candidate signal. This subject and most others were able to complete the trajectory tracking task qualitatively well using both thickness/deformation and activation signal traces.

On the other hand, these data highlight the drift that occurs in our thickness tracking system, which we also observed in the correlation analyses in section 4.4. While most subjects could compensate for this drift somewhat when guided by visual feedback, they were often unable to return the signal to baseline at later stages of each trial (a challenge reflected in subjects' written feedback, as discussed below in section 5.4). Activation-based tracking displayed no such drift, but showed many undesirable spikes and oscillations about the goal trajectory. In future iterations of this work, we will seek to ameliorate this noise with more aggressive mean filtering, though this will come at a cost of responsivity.

[2] In fact, these trajectories were much simpler than this statement implies: as stated, the two modes were simply the normalized thickness/deformation and normalized and mean-filtered activation measures extracted from ultrasound and sEMG data, respectively.

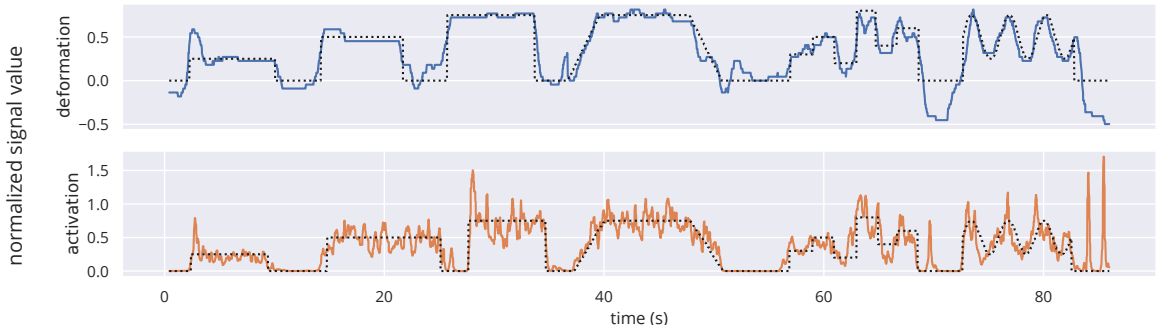

Figure 5.1: Example time series data collected from subject *Sub1* during separate deformation (*blue*, *top*) and activation (*orange*, *bottom*) tracking trials, alongside specified goal trajectories (*black*, *dotted*). Most subjects were able to complete the tracking task qualitatively well using both signal traces, despite drift in the deformation tracking system that sometimes made returning the system to baseline difficult (as is evident to the right of this trajectory), and uncontrolled, high-frequency oscillations in the activation signal. ©2021 IEEE.

Tracking Performance by Subject

As shown in Fig. 5.2, most subjects completed the tracking task with reasonably low RMS error using either sensor — generally performing best using thickness/deformation-based control, underscoring the value of this novel control paradigm. Perhaps unsurprisingly, *Sub7* (a subject discussed in section 4.4 as having irregular anatomy) was the exception to this success, and was largely unable to track the trajectory at all. *Sub9*, on the other hand, despite similarly irregular anatomy, was nevertheless able to adapt to the tracking system and achieve similarly low error to that of other subjects whose muscle deformation more reliably correlated with force. These outliers underscore the importance of continuing to examine additional possible deformation signals as we work to build generalizable control systems, but also the promising adaptability of human users to suboptimal control signals in this space.

Tracking Performance by Trajectory Type

As illustrated in Fig. 5.3, subjects were generally able to track all four trajectory types (first outlined in section 4.3). Relative to activation-based tracking — at which subjects performed comparably regardless of trajectory type, aside from slightly higher and more variable error during the *step* condition, perhaps due to the combined challenge of modulating the activation signal both quickly and to arbitrary levels — subjects consistently achieved lower error using the thickness/deformation signal. The one exception to this improved performance was during the *sine* portion of the trajectory; we theorize that this diminished

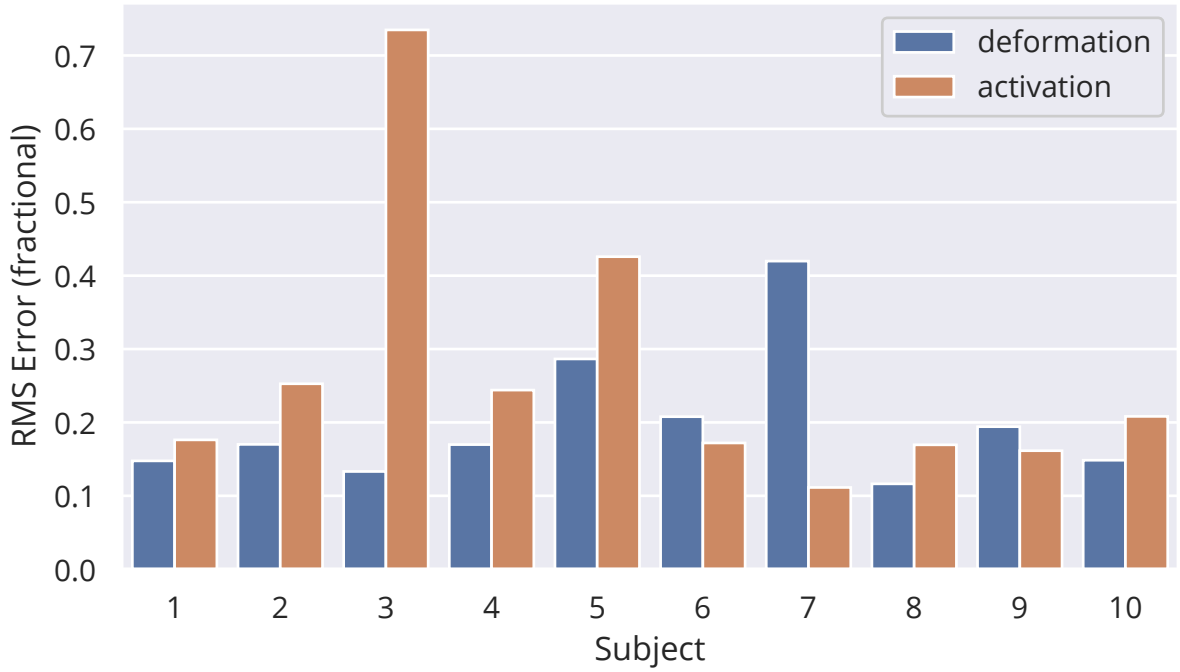

Figure 5.2: Tracking error during muscle deformation (*blue*) and activation (*orange*) trajectory tracking tasks across all subjects. With the exception of *Sub7* — whose irregular anatomy prevented almost any deformation-based control — subjects were largely able to achieve better, or at worst comparable, performance when performing deformation-based control, as compared with our baseline activation-based system. ©2021 IEEE.

performance is at least partially due to the drift observed in our tracking software (and discussed in section 4.4), though more exploration is needed to validate this claim (e.g., temporally rearranging trajectory elements).

5.4 Quantitative & Qualitative Tracking Preferences

Fig. 5.4 illustrates subjects' survey-reported tracking signal preferences, when evaluated both separately and comparatively. In both cases, subjects reported finding the thickness/deformation trajectory easier to control, and felt it better matched their perceived force trajectory, but that the activation trajectory was more responsive. Most subjects preferred the deformation-based controller overall — evidence that deformation-based tracking is not only feasible, but can be made intuitive for users.

In free-form responses, subjects noted many of the same qualitative characteristics of each controller that were observed by experimenters: many found that the deformation-

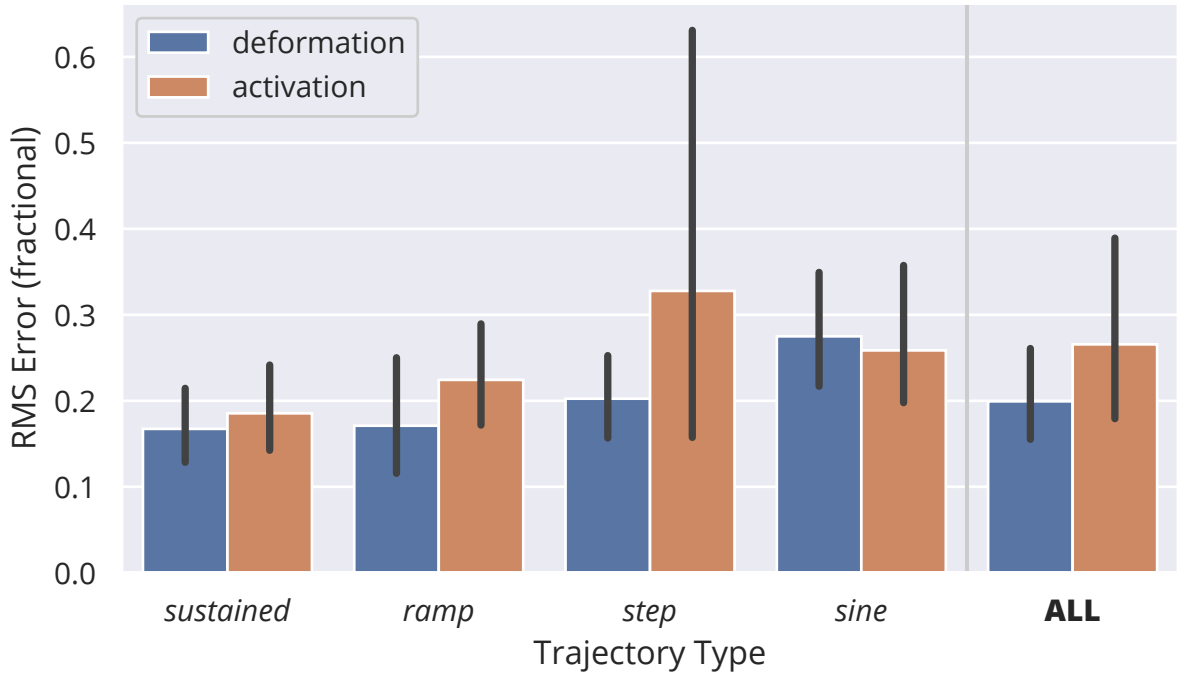

Figure 5.3: Tracking error during muscle deformation (*blue*) and activation (*orange*) trajectory tracking tasks across various trajectory types and in aggregate, with noted standard deviation across subjects. During all but *sine* trajectories — which we again theorize were impacted by the drift in our deformation tracking system — subjects were able to consistently achieve lower tracking error using the deformation signal. ©2021 IEEE.

based tracker was "better for maintaining a steady trajectory" but found it "hard to reduce the signal to 'at rest' level"; conversely, they noted the "quick response" of the activation-based tracker but found it "hard to maintain constant force" and "hard to not overshoot." One subject even noted that the two tracking schemes "essentially had opposite issues" and explicitly suggested sensor fusion. These comments — which were largely consistent with experimenters' own observations — will be used to inform future improvements to each tracking system in isolation (e.g., more aggressive low-pass filtering of the activation signal, improved tracking software to address drift and step artifacts) and in new, sensor-fused systems, as discussed below.

Subjects' full survey responses are included with the open-source data release.

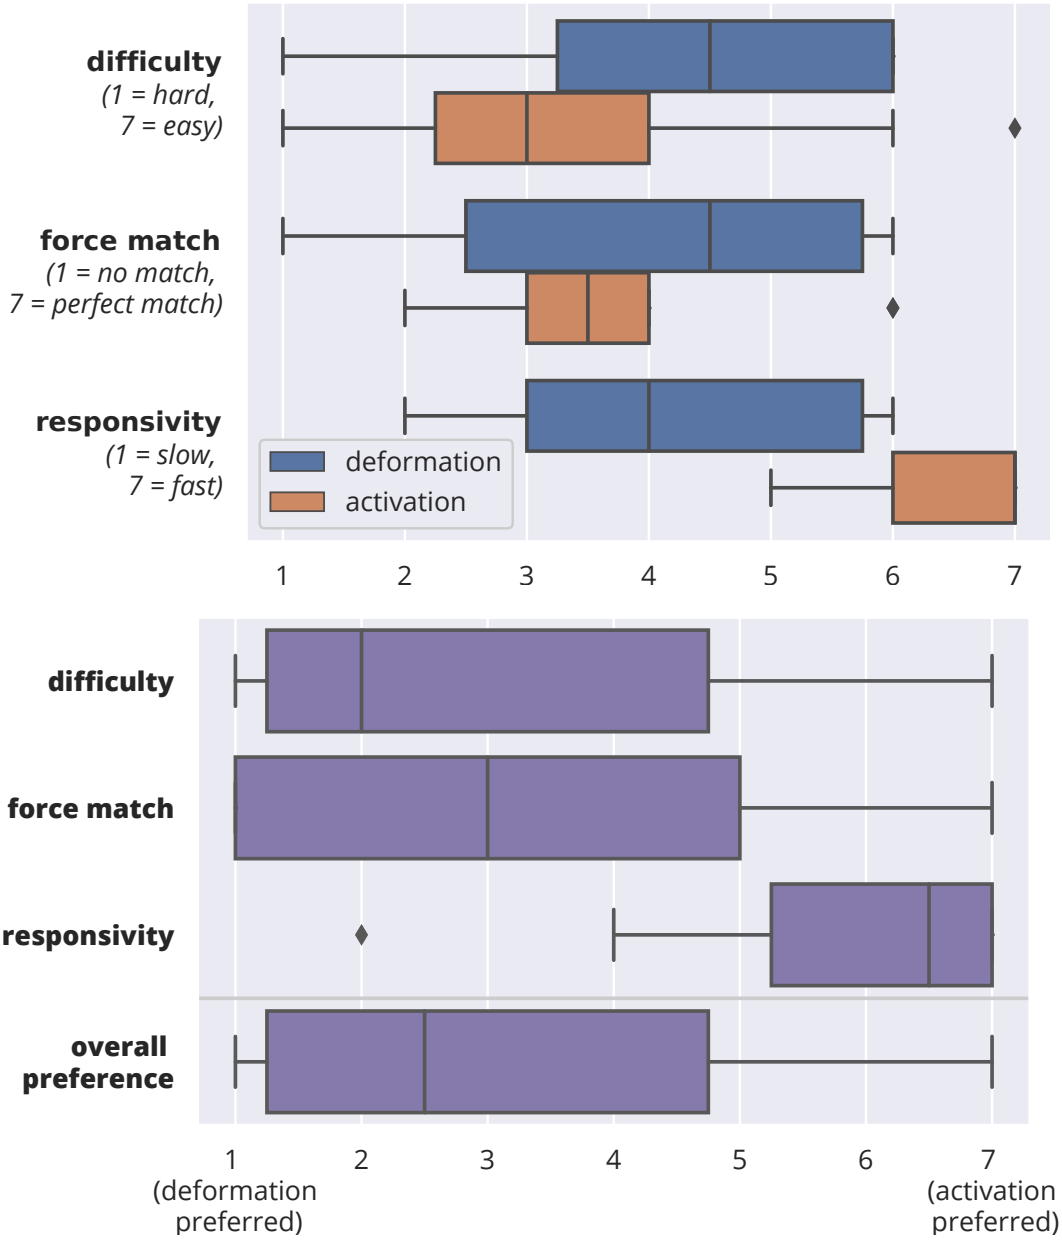

Figure 5.4: Subject preferences when asked to evaluate deformation- and activation-based tracking tasks separately (*top*) and in head-to-head comparisons (*bottom*). Subjects largely found the deformation-based tracking task easier and perceived the deformation signal to better match their output force, though they rated the activation-based tracker as more responsive, and most preferred the deformation-based tracker overall. Full survey questions and responses are included with the open-source data release. ©2021 IEEE.

5.5 Summary & Limitations

The promising results of this preliminary trajectory tracking study — in terms of both tracking accuracy and subject preferences — constitute strong evidence that it's possible, and even intuitive, to perform control with deformation-based signals. Beyond this dissertation, we are currently working to adapt this work to the control of physical devices, and in particular, to enable natural control of multiple degrees of freedom by extracting (highly localizable) deformation signals from multiple muscles simultaneously. Such expansions to real control applications will also require hardware enhancements, including the use of wearable ultrasound devices, to which this work should readily translate, many of which are under development [58, 106, 155].

The relative strengths and limitations of each sensing modality — namely, the slower response time and drift observed in ultrasound-measured deformation, and the oscillatory behavior of sEMG-measured activation — suggest compelling new control approaches leveraging these signals in complementary ways. In the future, we aim to explore a number of control schemes that combine these signals to exploit the strengths of each (e.g., using baseline sEMG signals as a trigger to reset tracked ultrasound points, or formulating a consensus approach to maintain responsiveness while avoiding erroneous motion).

Even without considering sensor fusion approaches, the prototype deformation tracking system used in this study is both simplistic (allowing tracking of only a single one-dimensional thickness measure) and unoptimized (using a basic implementation of Lucas–Kanade optical flow, as outlined in section 4.2). In the next chapter (6), we examine several optical-flow-based methods to enable tracking of the full muscle contour, and thus a richer class of possible control signals.

Chapter 6

Expanded Deformation Tracking

This chapter is adapted in part from [51].[1]

In the previous chapter, we showed that deformation-based control is both feasible and intuitive (for most subjects in a preliminary cohort) using a simple thickness-based measure under constrained experimental conditions; however, this control method was more effective for some subjects' morphology than others', and exploratory analyses in part I confirm the complexity of the deformation signal. In this chapter, we present a first effort at extraction of a richer class of possible control signals using optical-flow-based methods.

In the sections below, we present several algorithms to track both the candidate deformation signals discussed in chapter 3 (CSA, T, and AR, as defined in section 3.3) and the overall contour of the brachioradialis muscle. We first precisely define this tracking problem and its associated performance metrics, then discuss each algorithm's performance on the multi-angle, multi-subject (OpenArm Multisensor 1.0) data set generated and analyzed in chapter 3.

All data and analysis code have been released to the OpenArm project in the OpenArm Multisensor 1.0 package (see section 11.1).

6.1 Study Scope

This study constitutes a preliminary analysis of several optical-flow-based tracking algorithms on the OpenArm Multisensor 1.0 data set — of 5 subjects, denoted *Sub1–Sub5* — generated and analyzed in chapter 3. Although any number of tracking schemes, both sparse and dense, could be used to extract varied measures of deformation, we focus on the problem of muscle contour tracking — more precisely, tracking the edges of the brachioradialis muscle

[1] ©2020 IEEE. Reprinted, with permission, from Laura A. Hallock, Akash Velu, Amanda Schwartz, and Ruzena Bajcsy. "Muscle deformation correlates with output force during isometric contraction". In: *IEEE RAS/EMBS International Conference on Biomedical Robotics and Biomechatronics (BioRob)*. IEEE. 2020, pp. 1188–1195. DOI: 10.1109/BioRob49111.2020.9224391.

as it shifts over time. We choose this formulation for a number of reasons. First, the *CSA*, *T*, and *AR* deformation signals we showed in chapter 3 to be well-correlated with output force can be readily extracted from the tracked contour. Second, the manually-generated mask data used in chapter 3's analysis provides a ready ground truth signal for both overall tracking quality (via Jaccard distance computation, for example, as detailed in section 6.3 below) and for our previously examined deformation signals of interest. Third, the tracked contour admits future extraction of more complex deformation signals (e.g., statistical shape models or masked dense optical flow measures) while isolating a given muscle of interest from the surrounding tissue.

Lastly, tracking this cross section is feasible, if challenging: the brachioradialis, like many muscles of interest, is surrounded by fascia that appear brighter than surrounding tissue on an ultrasound scan. While the fascial structure is sometimes narrow and difficult to track — a fact that informs our more sophisticated tracking schemes — the preliminary results below indicate that even standard, untuned tracking methods perform relatively well in many cases.

Note that this study includes only offline data and analysis, which are used to determine promising directions for future tracking system development. Use of these signals for real-time tracking and control is left for future investigation.

6.2 Candidate Tracking Algorithms

Tracking biological structures within ultrasound data is an active area of research, with a number of general [77] and structure-specific [131, 162, 107, 4, 81, 30] approaches proposed. As a proof of concept, we examine the following candidate algorithms, each of which is built on the standard iterative Lucas–Kanade method of optical flow estimation [96] as implemented in the OpenCV Python library [14].

For the latter two algorithms, we present tracking error results for two sets of parameter values: "general" (tuned via expert-informed grid search on the *Sub1* data series collected at 69° elbow flexion) and "subject-specific" (tuned using each subject's ∼69° data series). The full details of each implementation, including algorithm-specific parameter values and tuning methods, have been released with the analysis codebase, and example frames from each tracker are shown in Fig. 6.1.

Naive Lucas–Kanade (LK). As a tracking baseline, we perform unmodified iterative Lucas–Kanade optical flow on every point along the brachioradialis contour, regardless of feature quality, initialized via the manually segmented contour mask in the first frame.

Feature-Refined Lucas–Kanade (FRLK). As a refinement of the LK procedure above, we track only the top 70% of feature points along the manually segmented contour as measured by their Shi-Tomasi corner score [134], a good indicator of the trackability of a point. In general, this refinement results in less drift-associated error from poor feature points but often fails to track narrower sections of the fascia altogether.

Bilaterally-Filtered Lucas–Kanade (BFLK). To further improve the tracking of individual contour points through noise removal without compromising edge integrity, we generate two

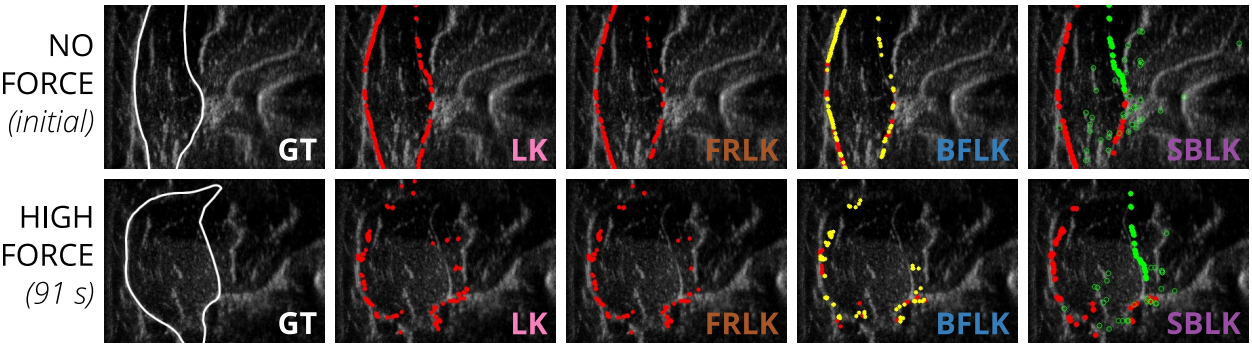

Figure 6.1: Example ultrasound frames from *Sub1* 69° data series and their respective tracked contour points using each candidate algorithm, at (relaxed) initialization (*top*) and approximately 91 s into tracking during force exertion (*bottom*). *Center left*: Complete set of contour points (*red*) tracked via naive Lucas–Kanade (*LK*) generally describes ground truth muscle shape (*GT*, *left*), but fails to track more extreme deviations and exhibits significant drift error. *Center*: Refined set of contour points based on feature quality (*red*) and tracked via Lucas–Kanade (*FRLK*) describe muscle shape only mildly more accurately than naive *LK* approach. *Center right*: Refined contour points tracked after aggressive (*red*) and less aggressive (*yellow*) bilateral filtering (*BFLK*) experience slightly less drift error, but still fail to capture deviation in the top right quadrant of the contour where fascia is narrow. *Right*: Combining Lucas–Kanade-tracked refined feature points (*red*) with contour points (*green, closed*) predicted based on supporter point locations (*green, open*) (*SBLK*) results in reasonable tracking of the full muscle contour as long as supporter point motion is well-correlated with contour motion. ©2020 IEEE.

denoised images at each time point via two separate bilateral filters [141]. The first of these filters is tuned to aggressively suppress speckle noise, but also removes narrower sections of fascia; the second generates a noisier image but retains even narrow fascia sections. Points along the initial manually segmented contour are again culled based on feature quality, this time in each of the filtered images, to generate two sets of high-quality feature points. (Points considered high quality on both filtered images are tracked on the more noise-suppressed frame.) These point sets are then tracked via standard Lucas–Kanade on their respective filtered images, and the contour at each time point is computed from the union of these points. Parameters tuned for subject-specific algorithm implementations include both the noise reduction characteristics of each bilateral filter and the fraction of highest-quality points used for tracking.

Supporter-Based Lucas–Kanade (SBLK). We further refine the *BFLK* method above in a manner similar to that described in [116]: high quality contour points (after aggressive bilateral filtering) are tracked via Lucas–Kanade, and the remainder of the contour is filled in based on the relationship of contour points to "supporter points" of high feature quality

extracted from throughout the image. To avoid irregularities caused by alternating Lucas–Kanade-tracked and supporter-tracked points — and observing that the narrowest, most featureless portion of the contour occurs in the top right quadrant of the contour — we enforce that all supporter-tracked points belong to this quadrant, and all Lucas–Kanade-tracked contour points remain outside. We then again apply two bilateral filters to each frame; using the first and most aggressive, we again obtain the top fraction of contour points by feature quality and track those that lie within the non-supporter quadrants via Lucas–Kanade. Next, using the second and less-aggressive filter, we select a set of "supporter points" of high feature quality, which we assume maintain a consistent distance to each contour point in the supporter-tracked quadrant. We then track these supporter points via Lucas–Kanade and use each supporter point to predict the new location of each contour point based on this constant-displacement assumption, in a manner similar to the "one-shot learning" method described in [41]. We calculate each contour point's final predicted location as the dynamically weighted mean over all supporter points' predictions, where we rely more heavily on supporter points that themselves deviate significantly over time from their initial position (to avoid relying too heavily on unmoving background feature points). Parameters tuned for subject-specific algorithm implementations include those tuned in the *BFLK* algorithm, as well as the number of supporter points, the fraction of Lucas–Kanade-tracked points, and properties of the (affine) prediction weighting function.

6.3 Tracking Error Metrics

In evaluating the performance of the algorithms above, we wish to characterize how well each algorithm tracks our *CSA*, *T*, and *AR* measures of interest, as well as how precisely it tracks the muscle contour overall. We thus present explicit *CSA*, *T*, and *AR* fractional error values along with contour segmentation error, computed as Jaccard distance (i.e., one minus intersection-over-union, or IoU). More precisely, Jaccard distance error is computed as

$$1 - \frac{|A_{\text{manual}} \cap A_{\text{tracked}}|}{|A_{\text{manual}} \cup A_{\text{tracked}}|}$$

at each frame, where A_{manual} is the manually segmented area of the brachioradialis contour and A_{tracked} is the area of the contour as tracked by the candidate algorithm.

6.4 Tracking Performance

In this section, we evaluate our candidate tracking algorithms on the chapter 3 OpenArm Multisensor 1.0 data set, remarking on each algorithm's general performance and its performance variation across subjects.

Table 6.1: Example Tracking Error

TRACKING ALGORITHM	Symbol	Tracking Error (*M*±*SD*) Jaccard Distance (1-IoU)	Fractional Error *CSA*	*T*	*AR*
Naive Lucas–Kanade	LK	0.17±0.05	0.37±0.22	0.11±0.08	0.29±0.13
Feature-Refined Lucas–Kanade	FRLK	0.19±0.05	0.31±0.18	0.09±0.06	0.28±0.15
Bilaterally-Filtered Lucas–Kanade, *general*	BFLK-G	0.18±0.06	0.38±0.20	0.11±0.09	0.30±0.12
Bilaterally-Filtered Lucas–Kanade, *tuned*	BFLK-T	0.16±0.04	0.24±0.13	0.10±0.07	0.31±0.12
Supporter-Based Lucas–Kanade, *general*	SBLK-G	0.20±0.06	0.30±0.22	0.09±0.05	0.29±0.09
Supporter-Based Lucas–Kanade, *tuned*	SBLK-T	0.18±0.05	0.19±0.11	0.08±0.05	0.36±0.13

Example tracking error of *LK*, *FRLK*, *BFLK*, and *SBLK* algorithms, both tuned and untuned, on the *Sub3* data series, computed as both Jaccard distance (one minus intersection-over-union, or IoU) and fractional error on predicted *CSA*, T, and AR deformation signals. In general, tuned algorithms outperform untuned algorithms, more structured algorithms (*BFLK* and *SBLK*) often outperform less structured ones, and T is tracked with lower mean error than *CSA* or AR. Values are presented as mean ± standard deviation. ©2020 IEEE.

An Illustrative Time Series

Table 6.1 shows aggregate tracking error values over the *Sub3* data series examined in chapter 3. While tracking performance varies by subject and kinematic configuration, these values illustrate several observed trends. First, adding structure to specifically track narrow sections of fascia via the *BFLK* and *SBLK* algorithms often improves performance over naive Lucas–Kanade baselines (though which of these two algorithms gives this performance boost varies by subject, as reflected in Fig. 6.1). Second, per-subject tuning generally improves tracking performance, as discussed in detail below. Third, thickness T is the easiest of our deformation measures to track, and AR the most challenging, suggesting that although AR correlates slightly more reliably with muscle output force, T — which also correlates reasonably well — may be a better choice of control signal; this is discussed more comprehensively below.

Multi-Subject Tracking

Fig. 6.2 shows tracking error for each proposed algorithm on the five subjects examined in chapter 3. Overall tracking quality (both mean and standard deviation of the error) varies substantially by subject, likely due to both the morphology differences illustrated in Fig. 3.4 (*Sub2*, for example, had extremely narrow fascia) and differences in motion strategy (*Sub4*, for instance, moved very abruptly). The importance of subject-specific tuning varies as well; *Sub5*, whose morphology varies substantially from that of *Sub1*, showed a large decrease in mean tracking error with tuning, while *Sub3*, whose morphology is more similar, showed a much more modest improvement.

Tracking Performance Across Deformation Measures

Fig. 6.3 illustrates the relative CSA, T, and AR tracking performance of our best-performing algorithm, *SBLK-T*, alongside Jaccard distance error, when computed in aggregate across all tracked trials for all subjects and elbow angles. In general, thickness T, the simplest signal, is tracked most reliably. The example (*Sub3*) time series in Fig. 6.4 corroborates this result: while tracked T values remain close to ground truth, CSA and AR values quickly begin to drift, a problem we also observed with our simpler thickness tracking system used in chapters 4 and 5. Generally, we observed this drift to be our most dominant source of error, suggesting that future tracking algorithms should take advantage of shape priors to periodically recalibrate to reasonable cross section shapes.

6.5 Summary & Limitations

In this chapter, we showed that optical-flow-based tracking can be used to track the brachioradialis muscle contour over time, and that imposing additional structure (e.g., separate,

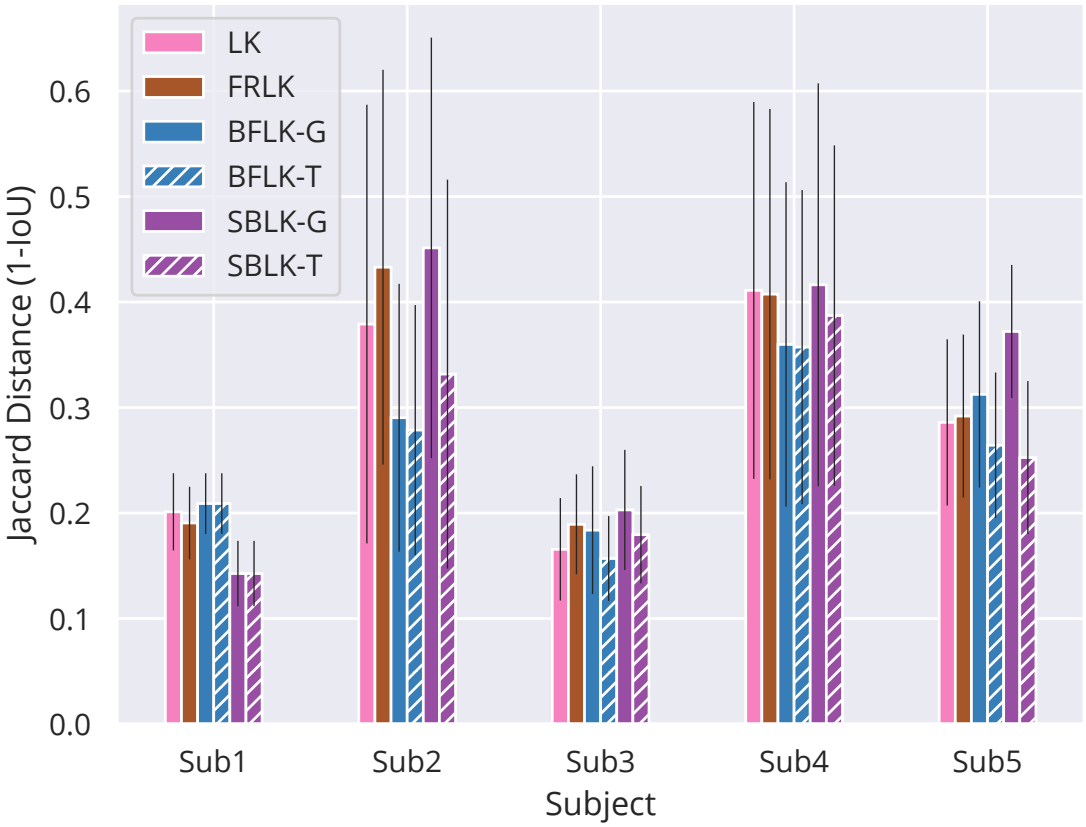

Figure 6.2: Mean tracking error and standard deviation of *LK* (*pink*), *FRLK* (*brown*), *BFLK* (*blue*), and *SBLK* (*purple*) algorithms, both tuned (*hashed*) and untuned (*solid*), computed as Jaccard distance, across various subjects, collected at ∼69° elbow flexion. Both the top performing algorithm and the level of tracking accuracy achieved, as well as the extent to which algorithm tuning mitigated errors, varied substantially by subject, likely due to variations in both morphology and motion qualia. Note that for *Sub1*, whose tuned values formed the generic baseline for all tracking, tuned and untuned tracking algorithms are the same. ©2020 IEEE.

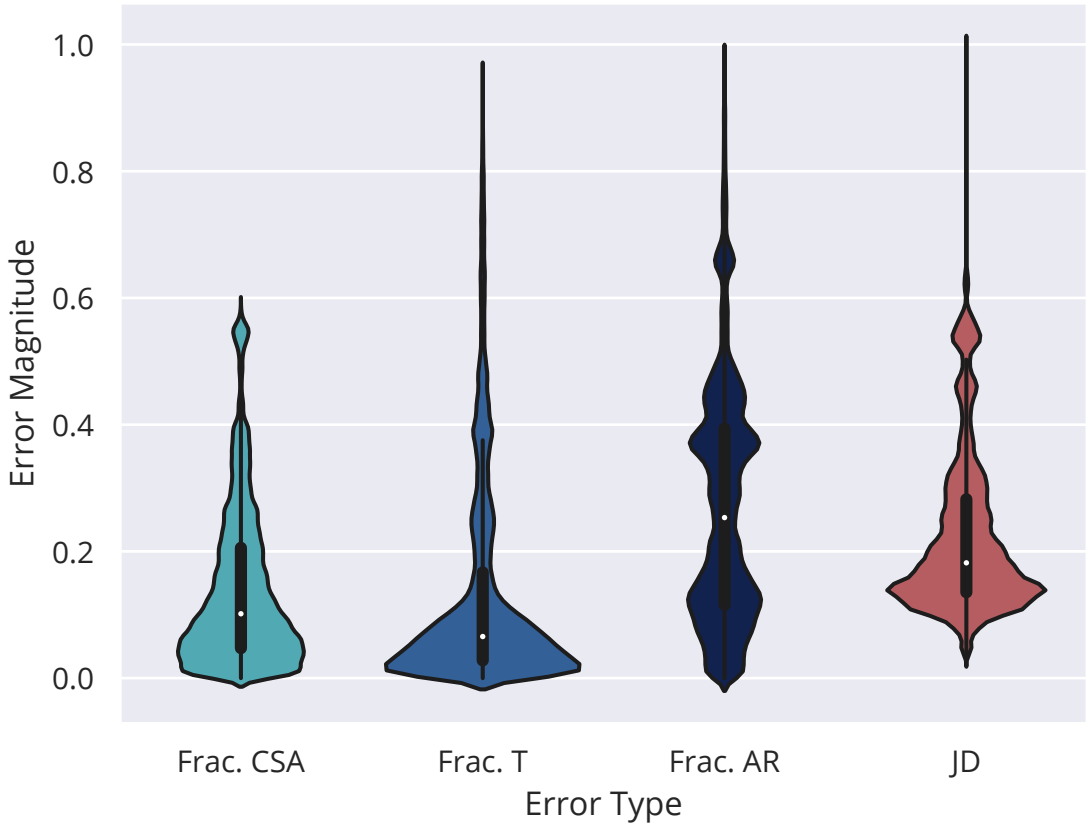

Figure 6.3: Overall tracking efficacy of *SBLK-T* algorithm for all aggregated time series, expressed as fractional *CSA* (*cyan*), *T* (*blue*), and *AR* (*dark blue*) error and Jaccard distance (*JD*, *red*). Histogram weights indicate that thickness *T* can be tracked most reliably, as compared with *CSA* and *AR*, making it the most promising control signal using our examined tracking algorithms.

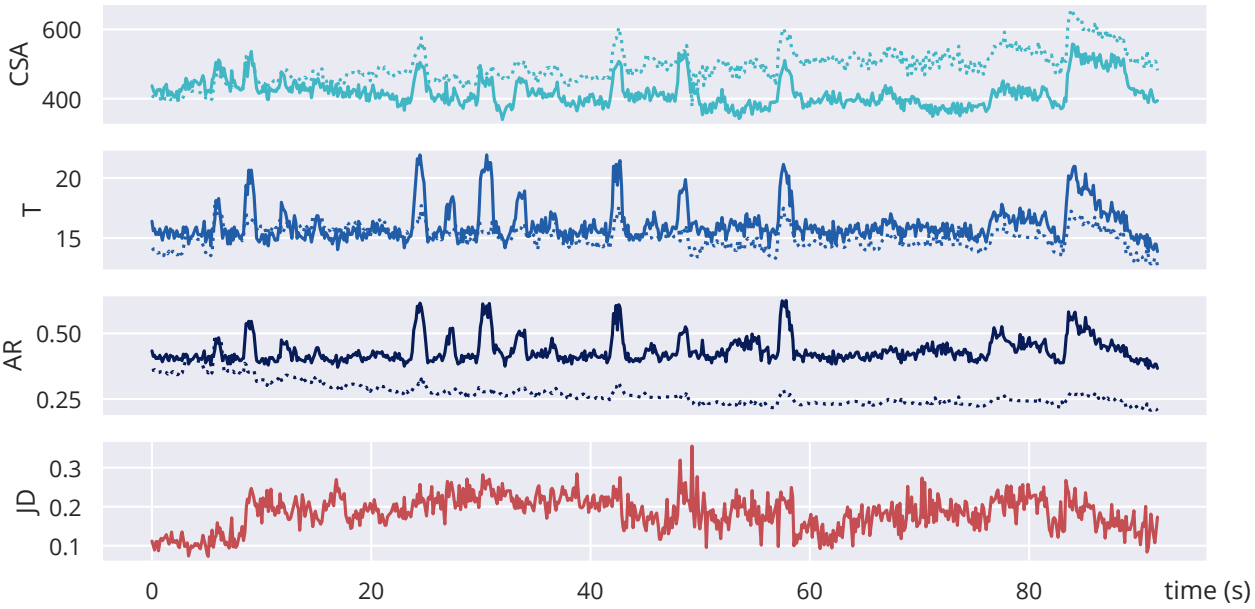

Figure 6.4: Example *CSA* (*cyan*), *T* (*blue*), and *AR* (*dark blue*) trajectories, both ground truth (*solid*) and tracked via *SBLK-T* algorithm (*dashed*), alongside time series Jaccard distance error (*JD*, *red*), for *Sub3* data series. For most data series, thickness *T* is tracked most reliably, while aspect ratio *AR* (and, to a lesser extent, cross-sectional area *CSA*) drift substantially.

supporter-based tracking of thick and thin fascial locations) can improve tracking performance. However, both the algorithms examined and our preliminary analyses of their quality are only the first steps toward real-time muscle cross section tracking, and the algorithms we selected were tailored to our particular data set of interest only in the most general sense. True assessment of their quality will require a much more rigorous examination of time series error characteristics, as well as collection of data sets more expressly designed to elucidate tracking error. (The data sets here, by contrast, were primarily collected to assess correlation, and thus contain little variation in speed of force application.)

In addition, although contour tracking enables new classes of possible control signals (e.g., mean overall magnitude of pixel motion, shape change from the relaxed condition), and there exist a multitude of more sophisticated tracking approaches we could explore (e.g., dense pixel tracking, edge tracking, applying shape priors), which signal types are most reliably related to force output remains unclear. Establishing these relevant signal types through the creation and systematic analysis of comprehensive 3D deformation data sets is the focus of parts III and IV.

Part III

Comprehensive Data Set

Chapter 7

3D Data Set Collection

This chapter is adapted in part from [49] and [113].[1]

In parts I and II, we showed that simple measures of muscle deformation correlate with output force during isometric contraction at the elbow, and that they can be used for control; at the same time, these preliminary studies made evident how much these signals vary across individuals and joint angles, such that substantially more data is needed to determine which deformation signals are most correlated, where they should be collected spatially on the arm, and how this relationship is modified by changes in kinematic configuration. Now, in part III, we build an extensive 3D data set to enable the comprehensive study of these relationships (a preliminary version of which appears in part IV).

In this chapter, we present two iterations of the core OpenArm data set, each of which comprises a collection of full-factorial 3D volumetric images of the arm under multiple kinematic configurations (i.e., elbow angles) and force conditions for multiple subjects. In the sections below, we first outline our data set scope and core imaging system design choices, then detail the subject cohort, as well as data collection and tissue annotation procedures, for each data iteration. Lastly, we discuss the limitations of these data and their potential application to device control, animation, and biomechanical study.

All data have been released to the OpenArm project in the OpenArm 1.0 and 2.0 packages as noted (see sections 11.3 and 11.4, respectively).

[1] ©2018–2019 IEEE. Reprinted, with permission, from Laura A. Hallock, Akira Kato, and Ruzena Bajcsy. "Empirical Quantification and Modeling of Muscle Deformation: Toward Ultrasound-Driven Assistive Device Control". In: *IEEE International Conference on Robotics and Automation (ICRA)*. IEEE. 2018, pp. 1825–1832. DOI: 10.1109/ICRA.2018.8462887 and Yonatan Nozik*, Laura A. Hallock*, Daniel Ho, Sai Mandava, Chris Mitchell, Thomas Hui Li, and Ruzena Bajcsy. "OpenArm 2.0: Automated Segmentation of 3D Tissue Structures for Multi-Subject Study of Muscle Deformation Dynamics". In: *International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*. IEEE. 2019, pp. 982–988. DOI: 10.1109/EMBC.2019.8857669. *Equal contribution.

7.1 Data Scope

As described in section 2.5, the relationship between muscle deformation and output force is straightforward: as (isovolumetric) muscles shorten along their line of action to pull the connected tendon and impart force, they necessarily expand and otherwise deform in other dimensions. At the same time, as noted in our preliminary correlation studies in chapter 3, deformation also occurs under changes in kinematic configuration, as muscles are passively stretched to accommodate new distances between their attachment points and deformed by interactions with the surrounding tissues. Thus, in order to build a framework that can model the force–deformation relationship during natural motion and determine what deformation signals are most correlated with force and where to collect them — the ultimate goal motivating this thesis work — we must account for and model these configuration-associated deformation changes.

The two data sets below (OpenArm 1.0 and 2.0) were collected to enable this kind of comprehensive, systematic analysis of both force- and configuration-associated deformation at the elbow. Each data set comprises full 3D volumetric scans of the anterior side of a given subject's arm under factorial isometric loading conditions and kinematic configurations (i.e., elbow angles) with select annotated tissue structures (primarily, the full volume of the biceps brachii and the anterior surface of the humerus, as noted below). OpenArm 1.0 comprises full-factorial data sets of 3 subjects, with manually annotated bone and muscle structures for a subset of 1 subject; OpenArm 2.0 comprises similar data, collected using a refined experimental procedure to enable more consistent force output, for 11 subjects, with manual- and neural-network-annotated[2] biceps brachii and humerus volumes for a select subset and all subjects, respectively.

7.2 3D Imaging Modality: Localized 2D Ultrasound

To generate the 3D tissue scans we require, MRI is perhaps the most obvious imaging modality and has been used in prior studies of muscle motion [161]. However, limited bore size makes scanning under natural arm configurations difficult or impossible, and long scan times make collecting data under multiple conditions prohibitively expensive and time-consuming. This is especially true when scanning under loaded configurations, as muscle fatigue may significantly influence the observed deformation signal.

Instead, we collect 3D scans via sweeps of a 2D brightness mode (B-mode) ultrasound probe — a technology that is cheaper, better studied, and more readily available than 3D ultrasound — whose position is spatially registered via motion capture. This technique has been widely used in both human- and robot-guided ultrasound imaging for surgical [146] and exploratory [69, 9] applications, and comprehensive tools for scan generation are widely available [88]. While the system is vulnerable to significant ultrasound-specific artifacts, including acoustic shadowing opposite ultrasound-impermeable bone and imaging of the

[2]This automated annotation process is the focus of chapter 8.

surface gel, structures of interest (bone surfaces, muscle fascia, etc.) are readily visible. Moreover, the comparatively short scan time and unrestricted workspace permit numerous scans of a single subject.

Specifically, to generate 3D images under each examined condition, we collect a dense data set of B-mode ultrasound scans along the anterior surface of the arm, each of which is spatially localized via calibration of the marker-tagged ultrasound probe with an active motion capture system. These scans are then used to build a 3D intensity map within which muscle fascia can be readily observed, thus allowing for segmentation of tissue structures of interest. The full details of this process and its usage in the generation of each data set are presented below.

7.3 Prototype Data Set (OpenArm 1.0)

In this section, we describe the generation of an initial volumetric data set, hereafter denoted OpenArm 1.0. We first outline our subject cohort and data scope, then detail the experimental procedures used to collect the raw ultrasound data, reconstruct the spatial location of each ultrasound scan, extract the relevant muscle volumes and bone surfaces, and align all scans for analysis in future chapters.

Subject Biometric Data & Consent

Data were collected from the right arms of three total subjects, one primary subject for which tissue annotation was performed for select scans (male, age 25, mass 55 kg, height 1.6 m) and two additional subjects for future analysis (female, age 24, mass 70 kg, height 1.8 m and male, age 21, mass 66 kg, height 1.6 m). All subjects were healthy and right-handed.

The study protocol was approved by the University of California Institutional Review Board for human protection and privacy under Protocol ID 2016-01-8261 (first approved 4 April 2016) and written informed consent was obtained from each subject.

Data Collection

During data collection, the test subject lay supine and relaxed, with legs comfortably extended and right arm extended laterally from the body at a 90° shoulder abduction angle. The forearm was fully supinated, with the upper arm supported at the distal end of the humerus, as shown in Fig. 7.1. Scans were then collected with the subject's elbow held statically at each of four angle values (0°, 30°, 60°, and 90° as measured from from full extension) under five separate loading conditions (fully supported by a jig at the wrist, unsupported but unloaded, and unsupported while lifting wrist weights of comparatively low, intermediate, and high mass), for a total of 20 trials. Angle conditions were selected to allow for both observable kinematic-associated muscle deformation and free manipulation of the ultrasound

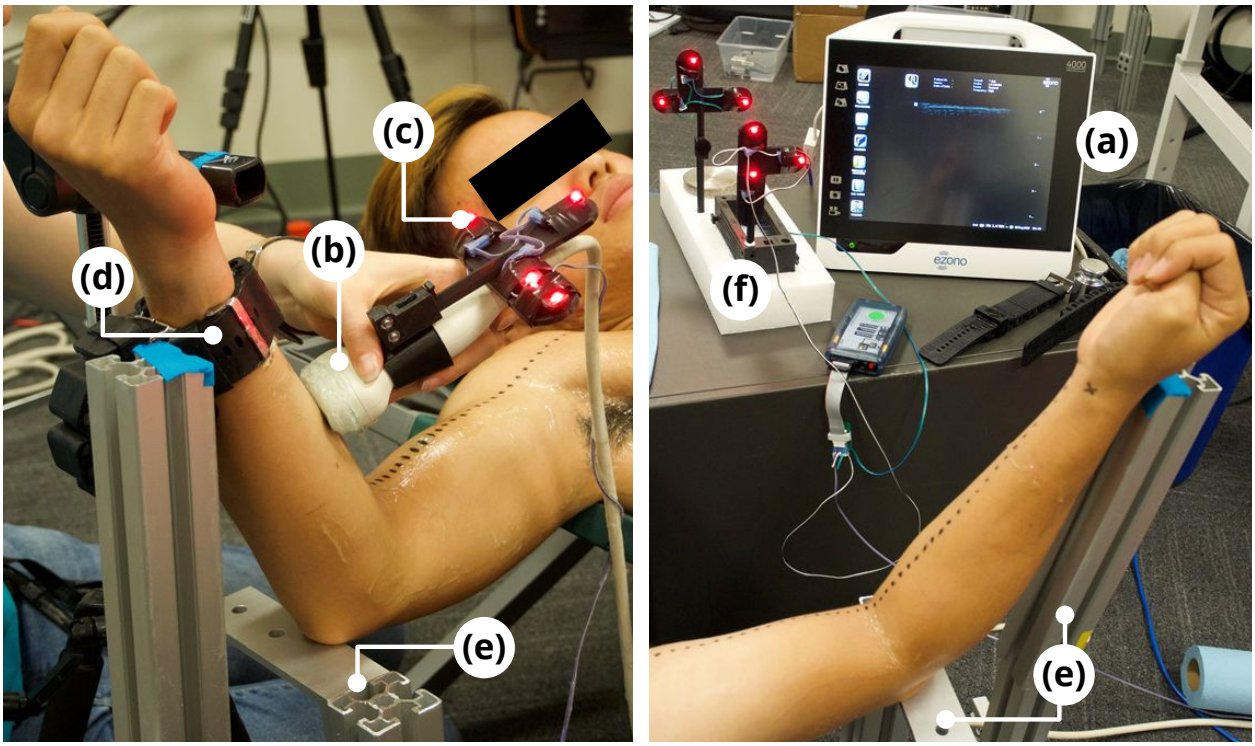

Figure 7.1: Experimental setup for the collection of OpenArm 1.0 full-arm upper-limb morphology data under multiple elbow angles and loading conditions, enabling a factorial study of the sources of muscle deformation (shown here at a 60° angle of elbow flexion under HF (*left*) and FS (*right*) loading conditions). Setup includes ultrasound scanner (a) and probe (b) (with attached active motion capture markers (c) used for spatial tracking); weight bands (d) used to load the elbow flexors (during LF, MF, and HF trials); mechanical jig (e) used to support the elbow (during all trials, *left*) and the forearm (during FS trials, *right*); and the phantom devices (f) required to calibrate the coordinate transformation between the motion capture world frame and the measured ultrasound scans. ©2018 IEEE.

Table 7.1: Volumetric Data Collection Conditions, OpenArm 1.0

Manipulated Factors		**Levels***
θ	elbow flexion angle	*0°*, 30°, 60°, 90°
LC	elbow load condition	fully supported (FS), *gravity compensation (GC)*, under load of 227g (LF), *under load of 727g (MF)*, under load of 954g (HF)

*Italicized values denote conditions for which raw ultrasound data were collected but for which tissue annotation was not completed (in deference to the time-intensive nature of manual annotation, addressed via automated segmentation in the data set's second iteration and detailed in chapter 8). Note that "fully supported" (FS) trials occurred while the arm was fully supported by the experimental jig (i.e., to measure "pure" kinematic deformation), while "gravity compensation" (GC) trials occurred while the arm was unsupported but unloaded (i.e., the elbow flexors performed gravity compensation for the arm's mass, but nothing more). Elbow flexion angles are measured from full extension. ©2018 IEEE.

probe around the elbow joint (the latter of which precluded scans at larger elbow flexion angles). Similarly, loading masses were selected to allow for observable force-associated muscle deformation while not being so heavy that subjects were unable to remain still during the (several minute) duration of the scan. During unsupported weight-bearing trials, subjects were asked to maintain contact between a designated point on the anterior surface of the wrist and a guide bar in order to maintain constant elbow flexion angles both within and between trials. Loading weights were attached at the wrist to avoid the confounding activation of wrist and finger flexors that would occur if weights were held in the hand. These conditions are summarized in Table 7.1.

During each trial, ultrasound images were collected using a portable commercial ultrasound scanner (eZono 4000, eZono AG, Jena, Germany) equipped with a 3–12 MHz linear transducer (L3–12 NGS, eZono AG, Jena, Germany). The machine was configured to collect B-mode data at a depth of 4 cm, with a 3.8 cm transducer footprint.

To collect full volumetric data of the anterior surface of the arm, the ultrasound probe was held perpendicular to the subject's skin and swept by a practiced operator along the arm's surface, using the operator-estimated minimum pressure necessary to maintain probe contact in order to deform the tissue as little as possible.[3] To scan the full anterior surface

[3]While the operator took care to maintain this pressure as consistently as possible, the deformation imposed at the probe contact point is likely a source of error in our collected images. There exist a number

of the arm (from medial to lateral edges) required multiple parallel sweeps of the probe; to ensure that scans were sufficiently evenly distributed in space, the subject's arm was marked in 1 cm increments using a non-toxic marker, and the operator maintained a metronome-guided constant rate of approximately 80 cm/min when sweeping proximally-to-distally along the arm, deviating as necessary near the shoulder and elbow joints to acquire sufficient numbers of scans. (Note that the volumetric reconstruction of these scans used here does not impose specific requirements on the spatial distribution of data, so this process need not be completely rigorous and simply improves signal quality.)

To permit volumetric reconstruction of the ultrasound data, the spatial location of the probe was tracked from a set of four optical markers using a PhaseSpace active motion capture system (PhaseSpace Inc., San Leandro, CA, USA). Prior to data collection, the transformation between the ultrasound probe location and the measured image was calculated (both spatially and temporally) using the open-source PLUS calibration toolkit [88], with a reported probe calibration error of 1.6 mm. During data collection, data were streamed to an external computer at a rate of 30 Hz through an OpenIGTLink server [140] and later reconstructed using the volume reconstruction application provided by the PLUS toolkit.

The full experimental setup is shown in Fig. 7.1, and representative volumetric data can be seen in Fig. 7.2 as the spatial intensity map from which volumes were manually segmented.

Tissue Segmentation

To provide deformation data for a factorial range of force and kinematic conditions, while acknowledging the time intensive nature of 3D image annotation, we selected nine scans from our primary subject data set for tissue segmentation (at 30°, 60°, and 90° elbow flexion and fully supported, lightly loaded, and heavily loaded force conditions, as noted in Table 7.1).

A complete study of muscle deformation about the elbow would require characterization of all muscles that actuate the joint, including flexors and extensors, as well as the muscles of the surrounding joints that contact and collide with those above. This study includes preliminary segmentation results of all elbow flexors[4] (biceps brachii, brachialis, and brachioradialis), but we report primarily on the observed deformation of the biceps brachii, which can be most cleanly observed in all nine examined scans and — as a surface muscle — is a natural target for eventual use in assistive device control signals.[5]

of strategies for mitigating this distortion via robot-assisted scanning [128, 10, 101, 28] or correcting for it in post-processing based on collected force information [148], but a number of technical challenges (robot arms occluding motion capture markers, the necessity of irregular and subject-specific scanning paths about the shoulder and elbow, the number of scans requiring consistent processing, etc.) prevented their usage in the collection of this preliminary data set.

[4]Elbow extensor muscles — primarily the triceps brachii — could not be segmented due to the placement of jig supports.

[5]While the brachialis muscle is also known to exert significant force during elbow flexion, the fact that its belly is largely centered over the elbow joint itself makes deformation difficult to observe and characterize (unlike the biceps brachii, whose mass is concentrated along the upper arm). The brachioradialis, despite its usage in our chapter 3 and 4 analyses, was not prioritized for scanning, as it provides a smaller fraction of

For initial segmentation, we selected the scan collected at a 30° elbow flexion angle while fully supported, as this condition represents a natural baseline for all subsequent conditions in our targeted data subset that increase angle and loading. Muscle volumes for the biceps brachii, brachialis, and brachioradialis were observed and manually annotated in the axial, coronal, and sagittal planes of the scan using ITK-SNAP [156], as were the anterior surfaces of the humerus, ulna, and radius (in order to establish the locations of both muscle attachment points and the elbow joint itself).

The full manual segmentation of this initial scan represents a bottleneck in the analysis process: it took two operators tens of hours to complete the full segmentation at the desired level of precision. To segment additional scans, a manual rigid transformation was performed in 3D Slicer [32] to align the segmentation of the initial scan with each of the remaining eight volumetric scans, such that the humerus of each scan perfectly overlapped. These segmentations were then manually modified in ITK-SNAP to include the new muscle deformation observed and any other necessary clean-up (a process that took several hours per scan rather than several days).

In addition to decreasing the required segmentation time, this process of modifying an existing segmentation instead of segmenting from scratch mitigated the chances of the human segmenter encountering areas of scan ambiguity (e.g., wide muscle fascia, acoustic shadowing, poor signal quality around the elbow joint) and making different choices across scans, generating a deformation signal that could be erroneously attributed to configuration- or force-associated deformation. To further reduce the prevalence of these types of errors, all nine scans were manually aligned (again using humerus alignment as ground truth) and were simultaneously examined by the same segmenter, slice by slice, to ensure that segmentation ambiguities were resolved consistently across scans.

The final result of this segmentation process is a set of nine aligned upper-arm scans at the factorially-varying elbow flexion and loading conditions described above, as shown in Fig. 7.2. These scans constitute the core OpenArm 1.0 data set and are used in our preliminary 3D deformation analyses in part IV.

7.4 Refined, Autosegmented, Multi-Subject Data Set (OpenArm 2.0)

While the OpenArm 1.0 release provides a proof-of-concept data set for 3D deformation analysis — one that we leverage in our geometric analyses in chapter 9 — we highlight two substantial limitations to our experimental procedure above and the resultant data. First, the method used to generate varied force conditions — attaching weights at the elbow — generated force conditions that were not quantitatively comparable across force conditions: holding up a weight applied to the arm at 30° elbow flexion requires substantially more

output joint force (especially when the wrist is supinated), and is not always completely visible within the frame of our scans.

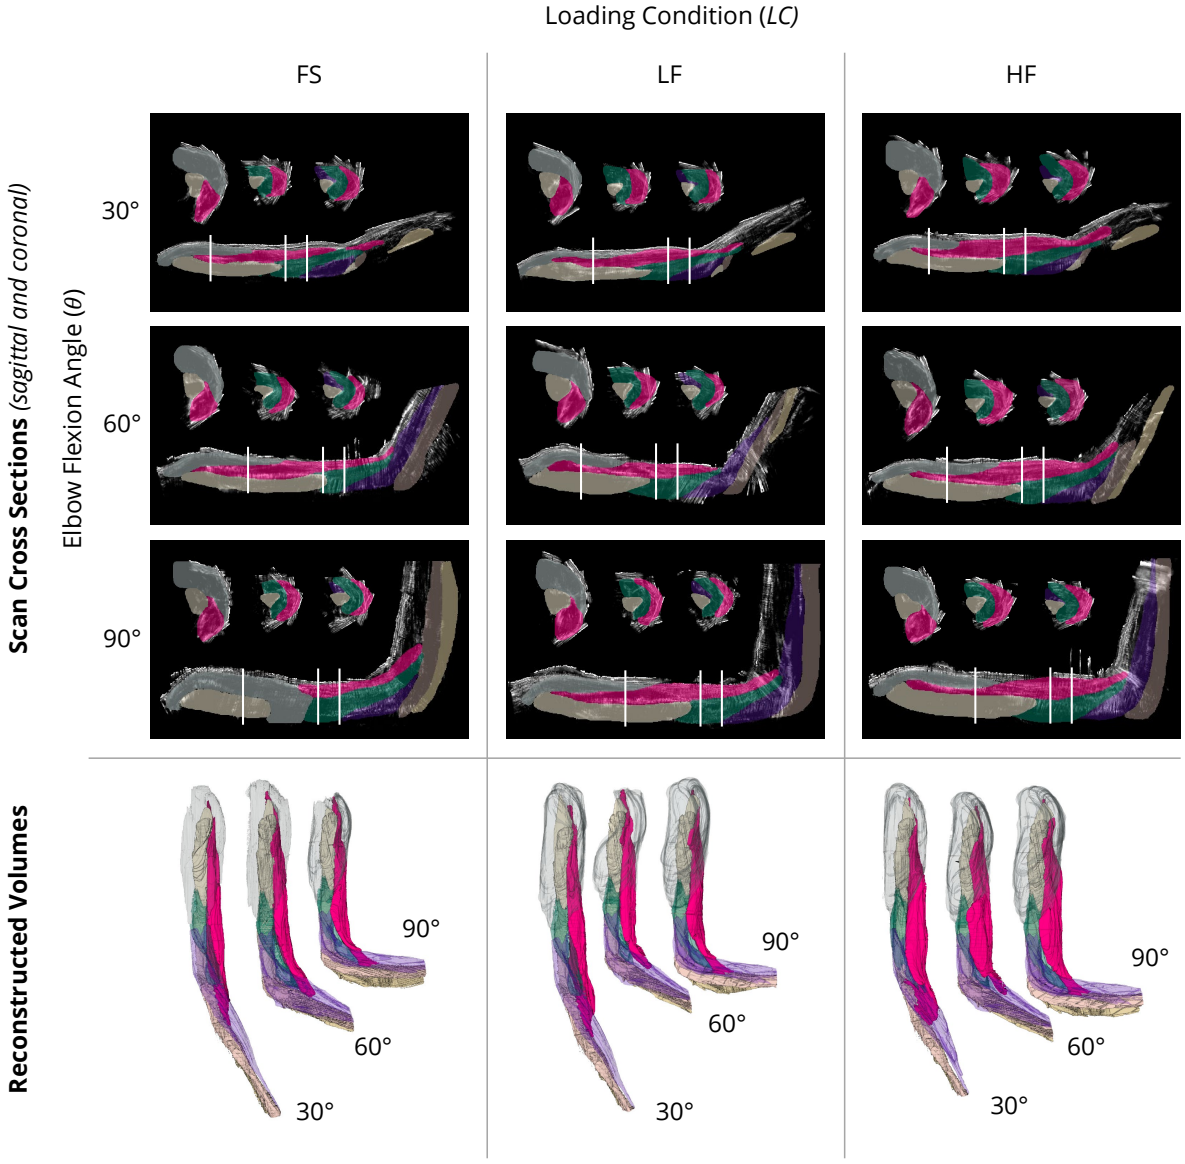

Figure 7.2: Tissue structures from the OpenArm 1.0 data set release, including the (*magenta*) biceps brachii and surrounding muscles (*turquoise* brachialis, *purple* brachioradialis, and *gray* deltoid) under multiple loading conditions and kinematic configurations, as segmented from volumetric reconstruction of ultrasound data. Locations of the coronal cross-sectional scans shown are noted by lines transecting the associated sagittal scan in each inset. Volumetric changes across both kinematic configurations and loading conditions are readily observable, confirming the necessity of modeling both signal sources when employing muscle deformation as a device control signal and informing our analyses in part IV. ©2018 IEEE.

exertion than at 90° flexion, for example, at which point it becomes chiefly a balancing task. Second, the time intensive nature of the manual segmentation process prevented the creation of a fully annotated multi-subject data set.

In this section, we describe the generation of a refined volumetric data set, hereafter denoted OpenArm 2.0, that addresses both limitations above, leveraging a new experimental setup to enable subjects to maintain specified output forces using visual feedback and a novel neural-network-based pipeline for automated image segmentation (described in detail in chapter 8). In the subsections below, we first outline our new subject cohort and data scope, then detail our refined procedures for data collection and tissue annotation.

Subject Biometric Data & Consent

Data were collected from the right arm of 10 subjects (6 male, 4 female, all right-handed, age 21.4 ± 2.46, mass 66.9 ± 10.1 kg, height 1.72 ± 0.0876 m, body mass index 22.5 ± 2.63), hereafter denoted *Sub1*–*Sub10*.[6] All subjects were healthy, with a wide variety of exercise regimes and body types. An incomplete set of scans, as noted in Table 7.2, was collected from an additional elderly subject (female, right-handed, age 85, mass 61.2 kg, height 1.52 m, body mass index 26.4), hereafter denoted *Sub11*, for preliminary evaluation of generalizability across age groups.

The study protocol was approved by the University of California Institutional Review Board for human protection and privacy under Protocol ID 2016-01-8261 (first approved 4 April 2016) and written informed consent was obtained from each subject.

Data Collection

Data collection methods were largely identical to those used in the generation of the OpenArm 1.0 data set described in section 7.3, aside from updates to subject demographics, posture, and loading conditions, as noted above and below.

During data collection, the test subject sat erect in a low stadium chair, with legs comfortably extended and right arm extended laterally from the body at a 90° shoulder abduction angle. The forearm was fully supinated, with the upper arm supported at the distal end of the humerus, as shown in Fig. 7.3. Scans were then collected with the subject's elbow held statically at the same four angle values examined in OpenArm 1.0 data collection under five refined loading conditions (fully supported by a jig at the wrist and unsupported while pressing upward on a force-torque sensorized handle with four prescribed levels of force), for a total of 20 trials. These conditions are summarized in Table 7.2. Subjects wore a brace to limit wrist flexion force and more completely isolate the elbow.

Loading conditions were selected for each participant based on the subject's maximum voluntary contraction (MVC) value. To measure this MVC value, subjects were asked to

[6] Statistics are reported as mean ± standard deviation. For additional demographic data, broken down by subject, see the full open-source data release. Note that the subjects examined here and their identifying numbers bear no relation to those in the OpenArm 1.0 data set or examined in previous part I and II studies.

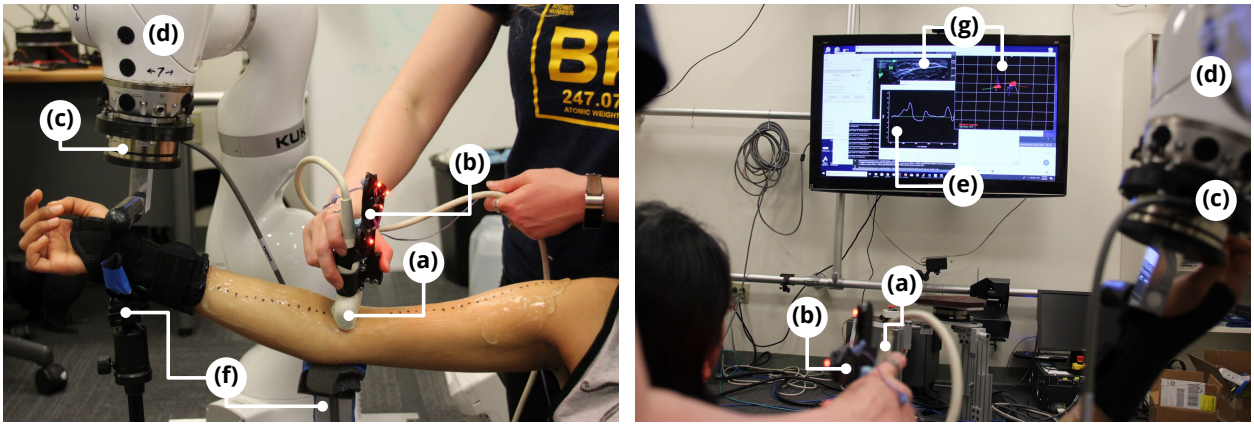

Figure 7.3: Experimental setup for the collection of OpenArm 2.0 full-arm upper-limb morphology data under multiple elbow angles and loading conditions (shown here at a 30° angle of elbow flexion under FS (*left*) and P1 (*right*) loading conditions). Setup includes ultrasound probe (a) (with attached active motion capture markers (b) used for spatial tracking); force-torque sensor (c), held statically in place by KUKA LBR iiwa 14 R820 robot (d) and used by the subject via real-time visual feedback (e) to maintain constant force output during loaded trials; mechanical jig (f) used to support the elbow (during all trials) and the forearm (during FS trials, *left*); and real-time ultrasound and motion capture data (g) for continuous system status monitoring. ©2019 IEEE.

press upward on the handle with maximum possible force, then hold for several seconds, then release. Subjects performed this sequence twice, both at full elbow extension, and the maximum overall force value was recorded as MVC. (Empirically, this value varied substantially across subjects, with mean and standard deviation 66.1 ± 28.8 N for the 10 primary subjects tested.)[7] To allow subjects to maintain a given force for the several minutes required to generate a full 3D scan, force conditions were chosen as 0, 10, 30, and 50 percent of the MVC value. Note that the same MVC value, and thus the same force conditions, were used for all angle conditions for the same subject, to allow for development of models explicitly relating muscle deformation to force. Collecting this value at full extension — the angle at which muscles are weakest [37] — ensured that subjects could maintain the forces required under all tested conditions. Subjects maintained the prescribed force during each trial by matching visual feedback from the attached force-torque sensor (ATI Axia80, ATI Industrial Automation, Apex, NC, USA) to a marked goal value on a real-time series plot, as shown in Fig. 7.3. All subjects were able to consistently maintain this value within several newtons.

During each trial, ultrasound images were collected by an experienced operator using the same motion-capture-localized ultrasound system and scanning protocol used to generate

[7]This sequence was performed at 30° of extension for the single elderly subject to avoid hyperextension injury, at which the measured MVC was 42.5 N.

Table 7.2: Volumetric Data Collection Conditions, OpenArm 2.0

Manipulated Factors		**Levels**‡
θ	elbow flexion angle	0°, 30°, 60°, 90°
LC	elbow load condition	fully supported (FS), 0% MVC (P0), 10% MVC (P1), 30% MVC (P3), 50% MVC (P5)

‡Data were collected from 10 subjects at each of the 4 flexion angles and 5 loading conditions listed above in a factorial manner, for a total of 20 scans per subject. Additional data were collected from a single elderly subject at 30° and 60° flexion angles (all loading conditions) and 0° and 90° (FS only), for a total of 12 scans. Note that fully supported (FS) trials occurred while the arm was held in position by the experimental jig (i.e., to measure "pure" kinematic deformation), while the latter 4 loading conditions denote percentages of maximum voluntary contraction (MVC) level. Elbow flexion angles are measured from full extension. ©2019 IEEE.

the OpenArm 1.0 data set and detailed in section 7.3 above, with all the same configuration parameters and calibration procedures.

The full experimental setup is shown in Fig. 7.3, and representative volumetric data can be seen in Fig. 8.3 as the spatial intensity map from which volumes are segmented.

Tissue Segmentation

To enable annotation of the full OpenArm 2.0 data set (over 200 scans) while minimizing the need for manual annotation, we restricted our segmentation to the biceps brachii (for deformation analysis) and the humerus (for alignment across scans) and developed an automated neural-network-based segmentation system. This system, including all examined architectures and modifications, is detailed in chapter 8.

7.5 Summary & Limitations

In this chapter, we described the generation of the OpenArm 1.0 and 2.0 data sets, which comprise the first sets of 3D scans of the arm under factorial kinematic configurations and loading conditions and enable comprehensive study of both muscle deformation sources across multiple individuals. In part IV, we leverage this data to begin this type of analysis, and in

the future, the full volumetric nature of the scans make the data sets an excellent platform with which to verify the wide variety of existing muscle deformation models and quickly test the feasibility of new ones. For this reason, we consider these data sets a core thesis contribution in their own right, and they have been open-source released for future study by the biomechanics, robotics, and animation communities.

At the same time, these data sets remain limited, and are only a starting point for study of the complex force–deformation relationship. Subject cohort sizes remain small, and the substantial morphological variation observed (here and in part I studies) makes it likely that a robust understanding of muscle deformation signals will require many more subjects of much more varied demographics. In addition, the length of time required to collect a single scan means that all scans were collected under static, though loaded conditions; how well models based on these data will translate to natural, unconstrained motion remains an open question. Lastly, as discussed above, while these data primarily support analysis of a single elbow flexor (the biceps brachii) and surrounding tissue structures, a more complete understanding of the elbow joint and its actuators would require analysis of all flexors and extensors — and even those of surrounding joints.

Despite these limitations, the OpenArm data sets are the first of their kind and enable unprecedented exploration of the force–deformation relationship. We aim to use the insights gained in analyzing these data to build more targeted data sets in the future, enabling expansion to larger subject cohorts, dynamic conditions, and additional muscles and tissue structures.

Chapter 8

Automated Tissue Segmentation

This chapter is adapted in part from [113].[1]

In the previous chapter, we described the generation of the OpenArm 1.0 and 2.0 data sets — 3D volumetric scans of the arm under factorial force conditions and kinematic configurations to enable comprehensive study of the force–deformation relationship. This chapter describes the development of a novel neural-network-based pipeline to enable automated segmentation of tissue structures of interest from these scans, for which manual annotation is prohibitively time intensive.

In the sections below, we first define our annotation system requirements (i.e., target tissue structures) and core (neural-network-based) approach, then present several candidate system architectures and evaluate their segmentation performance on the OpenArm 2.0 data set as compared with a registration-based baseline approach. Lastly, we comment on the system's limitations and speculate on its application to future annotation tasks.

All code and trained neural network models have been released to the OpenArm project in the OpenArm 2.0 and Annotation Source Code packages (see sections 11.4 and 11.5, respectively).

8.1 Segmentation Targets & Scope

While the segmentation pipeline we present could easily be re-trained for various tissue segmentation tasks, our primary objective is the segmentation of elbow flexors and other relevant tissue structures in the OpenArm 2.0 data set detailed in chapter 7. In particular, in a concession to the time-intensive nature of manual ground truth training data segmentation,

[1] ©2019 IEEE. Reprinted, with permission, from Yonatan Nozik*, Laura A. Hallock*, Daniel Ho, Sai Mandava, Chris Mitchell, Thomas Hui Li, and Ruzena Bajcsy. "OpenArm 2.0: Automated Segmentation of 3D Tissue Structures for Multi-Subject Study of Muscle Deformation Dynamics". In: *International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*. IEEE. 2019, pp. 982–988. DOI: 10.1109/EMBC.2019.8857669. *Equal contribution.

we limit our segmentation to the two tissue structures most immediately relevant to the deformation analyses undertaken in part IV: the biceps brachii (for direct analysis) and the humerus (for alignment across scans).[2]

8.2 Segmentation Approach: Convolutional Neural Networks

Convolutional neural networks (CNNs) are a promising tool in many medical imaging domains [92], including echocardiogram annotation [158], brain lesion segmentation [73], and localization of organ structures [110].

In this work, we adapt existing CNN-based methods to segment our tissue structures of interest, the humerus and the biceps brachii. Specifically, building from the U-Net [124] neural network framework, we examine various data augmentation techniques and training data sets to both optimize the network's performance on our data set and hypothesize strategies to better select training data, minimizing manual annotation time while maximizing performance. We then employ this pipeline to generate segmentations of biceps brachii and humerus tissue volumes for the OpenArm 2.0 data set detailed in chapter 7.

On a conceptual level, this problem is well suited to a deep learning approach: rules for segmentation are difficult to enumerate a priori (due to both the modality-specific artifacts mentioned above and the varied geometric behaviors of different tissue structures, which make classical registration approaches difficult); at the same time, the domain is well-specified and highly constrained, such that generating domain-spanning training data is a tractable problem.

As discussed below in section 8.5, we find that CNN-based segmentation significantly outperforms classical registration-based approaches in both quantitative segmentation accuracy and qualitative manual cleanup time.

8.3 Candidate Segmentation Architectures & Modifications

To segment the humerus and biceps brachii structures from the generated OpenArm 2.0 scans, we explored the following neural net architectures and data augmentation techniques. All networks were trained in Tensorflow [100] using the Adam optimizer [79] and a cross-entropy loss function on a custom-built desktop machine with an INTEL Core i7-5820K six-core CPU and an NVIDIA Titan Xp GPU. A full list of tested architectures and baselines

[2] An additional benefit is that these two structures are representative of the wider tissue segmentation challenge in that they have significantly different appearance and properties (e.g., different elasticity properties, tissue-specific artifacts like bone shadow); thus, successful segmentation of these structures constitutes strong evidence that our methods could be extended to segmentation of other bones and muscles in the future.

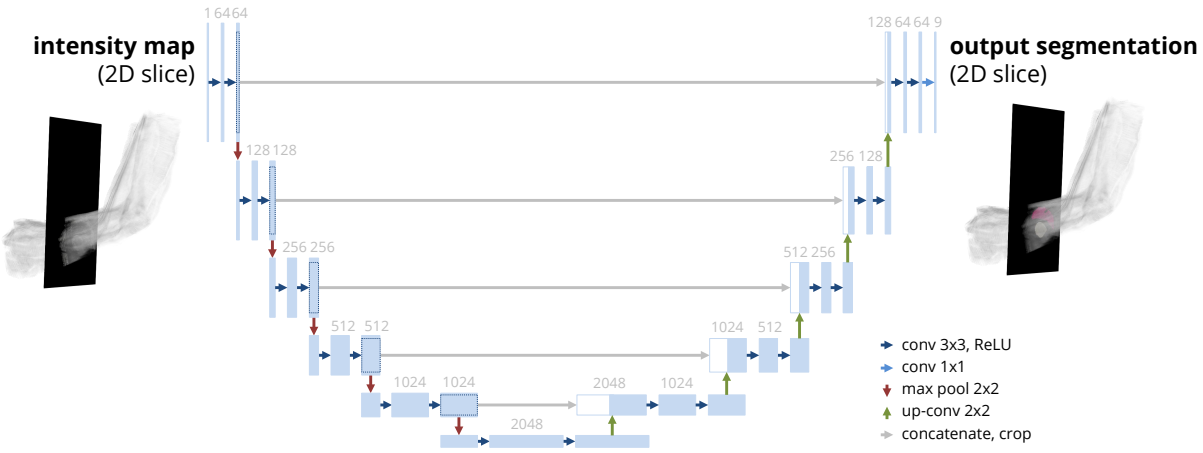

Figure 8.1: Baseline U-Net architecture for segmenting 2D slices of 3D volumetric scans, which are then compiled to generate full 3D segmentation. Architecture is almost exactly that of the original U-Net [124], with only a single additional concatenation block (corresponding to four additional 3x3 convolutions, one additional 2x2 max pooling operation, and one additional "up-convolution").

is shown in Table 8.1, and all hyperparameter values and architecture details are available with code release.

Baseline Architecture: 2D U-Net

Although our objective was to generate 3D scans, we chose the 2D U-Net [124] as our baseline segmentation architecture, which we applied to axial slices of the full upper-arm volume; each slice of each test scan was then predicted individually to generate full volumetric predictions. The U-Net was designed to perform well when trained on relatively small data sets and has been widely applied to various biomedical image segmentation tasks [158, 3], and it is more computationally efficient than many of its 3D counterparts. All examined networks built upon the original U-Net architecture, with one additional concatenation block, as shown in Fig. 8.1.

Data Augmentation

To generate additional training data without prohibitively time-consuming manual annotation, we artificially increased the size of our training data set via both rotational and elastic deformation, a useful and common practice in neural network training when comprehensive data is not available [136]. Specifically, we generated additional scans through arbitrary

rotation and random elastic deformation [145] and trained networks both with and without this augmented data, as noted in section 8.5 and Table 8.1.

Baselines/Controls for Comparison

We evaluate these neural network approaches against a classical registration-based approach in which tissue structures from one scan are mapped to another by finding the optimal transformation between the two spatial intensity maps. Specifically, we compare the methods above with both pure rigid registration (as a simple baseline) and a set of sequentially higher-degree-of-freedom transformations (rigid, affine, and B-spline-parameterized) in which each transformation is used to initialize the next registration (as a more complex baseline that better represents the upper limits of classical registration's efficacy). Optimal transformations were calculated using the SimpleElastix image registration library [98]; transformation quality was evaluated at each optimization step using the mutual information criterion. Manual translation was performed prior to automated registration to better align the scans, and sensitive hyperparameters were tuned via grid search to further optimize the results. Baseline registration code and associated hyperparameter values, like those used in neural-network-based methods, are available with code release.

8.4 Ground-Truth Data Set Generation

To train and evaluate the segmentation pipeline across kinematic configurations, loading conditions, and subjects, we manually segmented full-scan volumetric data for the following subjects:

- *Sub1* – all angles and loading conditions (20 scans)
- *Sub2* – all loading conditions at 30° (5 scans)
- *Sub3* & *Sub4* – single scans at 30°, FS (1 scan each)

Scans were selected to allow for comparison across variables of interest while remaining tractable. Each scan required approximately 10–12 hours of expert annotation time.

8.5 Segmentation Performance

The objectives of our investigations into neural network architectures were twofold: first, we sought a network that could reliably generate annotations across a variety of subjects and configurations; second, we aimed to derive principles to guide future architecture design in similar domains.

Two major factors influenced neural network performance: use of data augmentation and selection of training data. While a full factorial analysis under all possible augmentations

Table 8.1: Segmentation Accuracy

Architecture / Strategy	Accuracy (IoU, *Pixel Accuracy*)			
	same angle, same force, same subject[1]	**new** angle, same force, same subject[2]	same angle, **new** force, same subject[3]	same angle, same force, **new** subject[4]
RR (*rigid registration*)	0.431, *0.980*	0.320, *0.970*	0.537, *0.985*	0.244, *0.953*
RANR (*rigid-affine-nonlinear hierarchical registration*)	0.722, *0.991*	0.545, *0.980*	0.450, *0.980*	0.386, *0.960*
U-NET (*unmodified U-Net [124]*)	0.875, *0.996*	0.593, *0.982*	0.604, *0.988*	0.422, *0.961*
U-NET+RA (*U-Net + rotational augmentation*)	0.929, *0.998*	0.560, *0.984*	0.464, *0.986*	0.393, *0.971*
U-NET+EA (*U-Net + elastic deformation augmentation*)	0.950, *0.999*	0.677, *0.988*	0.573, *0.989*	0.533, *0.978*
U-NET+RA+EA	0.936, *0.998*	0.577, *0.984*	0.544, *0.988*	0.499, *0.978*
Multi-Angle U-NET	0.886, *0.997*	0.691, *0.989*	0.614, *0.989*	0.470, *0.972*
Multi-Angle U-NET+EA	0.906, *0.997*	0.717, *0.989*	0.651, *0.990*	0.523, *0.975*
Multi-Force U-NET	0.885, *0.997*	0.617, *0.985*	0.770, *0.994*	0.452, *0.972*
Multi-Force U-NET+EA	0.902, *0.997*	0.682, *0.988*	0.793, *0.994*	0.519, *0.977*
Multi-Subject U-NET	0.884, *0.997*	0.657, *0.987*	0.536, *0.988*	0.885, *0.995*
Multi-Subject U-NET+EA	0.908, *0.998*	0.687, *0.989*	0.565, *0.989*	0.909, *0.996*

RR, RANR, U-NET, U-NET+RA, U-NET+EA, and U-NET+EA were all trained on (or mapped from) the single *Sub1* scan at 30° and FS conditions, with and without augmented data from the same scan(s) as noted. Multi-Angle U-NETs were trained on *Sub1* scans at all angle conditions and FS loading; similarly, Multi-Force U-NETs were trained on *Sub1* scans at 30° and all force conditions, and Multi-Subject U-NETs were trained on all subjects at 30° and FS conditions. Note that grayed values constitute network "memorization" — i.e., predictions are calculated over data included in network training. These values are presented as baselines for the maximum performance we expect to achieve from a given strategy. ©2019 IEEE.

[1] Accuracy on *Sub1* scan at 30° angle and FS loading conditions.

[2] Mean accuracy on *Sub1* scans at 0°, 60°, and 90° angle and FS loading conditions.

[3] Mean accuracy on *Sub1* scans at 30° angle and P0, P1, P3, and P5 loading conditions.

[4] Mean accuracy on *Sub2*, *Sub3*, and *Sub4* scans at 30° angle and FS loading conditions.

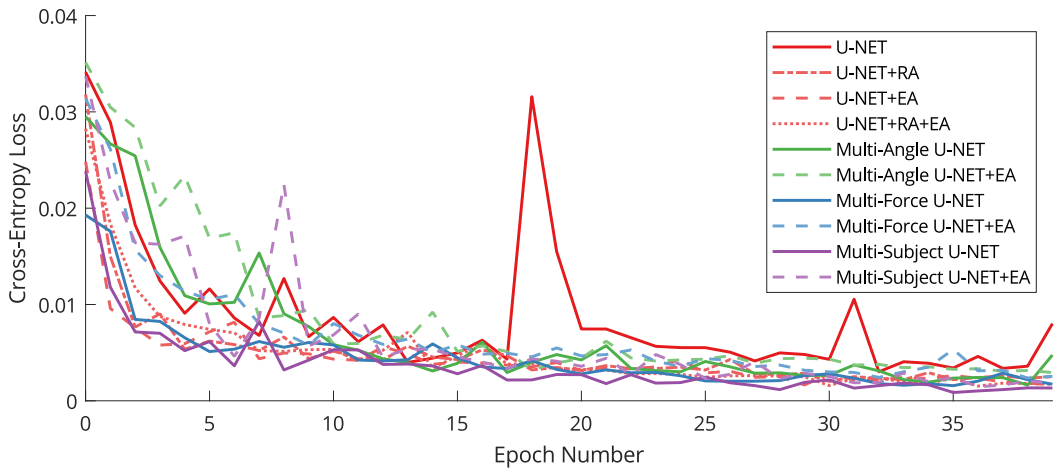

Figure 8.2: Training error for all reported network architectures over 40 epochs, over which networks reliably — if sometimes messily — converged. Accuracy values reported in Table 8.1 and predictions shown in Fig. 8.3 are computed from the minimum-loss epoch network for each respective architecture. ©2019 IEEE.

and training data sets was intractable, we tested both the augmentation types above and various promising sets of training data (incorporating single or multiple angle conditions, loading conditions, and subjects) in a principled manner, in both isolation and combination, as outlined in Table 8.1. Networks were restricted to using at most 500 slices of manually segmented data to allow fair comparison across training strategies. For networks in which multiple scans of training data were used, these 500 slices were distributed uniformly at random across scans. Up to 1000 slices of total augmented data were used, again distributed uniformly at random across augmented scans, to maintain tractable training time. Of these slices, 5% were reserved for validation at each epoch and 15% for final testing.[3] Networks were trained for 40 epochs — after which, empirically, they had reliably converged, as shown in Fig. 8.2 — and the epoch with best performance was selected for comparison.

The performance of these models under various conditions is shown in Table 8.1, and example qualitative segmentation data from select architectures are shown in Fig. 8.3.

Registration vs. CNN-Based Methods

Quantitatively, as shown in Table 8.1, neural networks perform almost uniformly better than registration-based methods in terms of both intersection-over-union (IoU) values and overall

[3]Segmentation accuracy on this test data was used internally to confirm that networks were not overfitting; here, we instead report accuracies across new 3D scans not used in training, as generalization across conditions, rather than across slices of the same 3D image, is of primary interest.

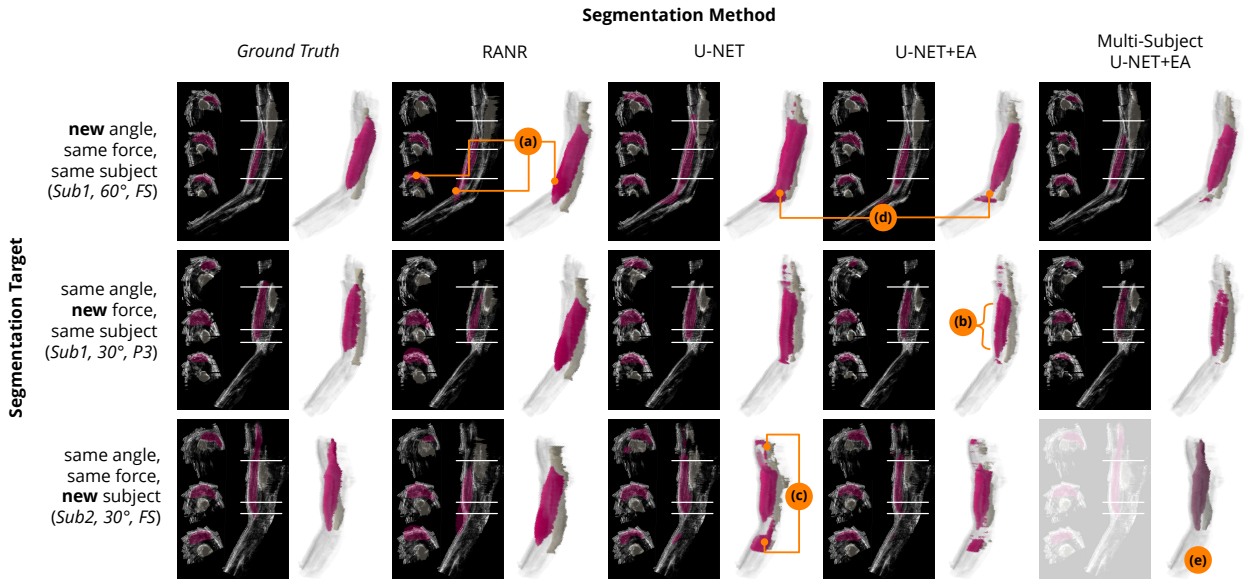

Figure 8.3: Exemplar volumetric data, as segmented manually (*Ground Truth*), via optimized classical image registration (RANR), and via neural network, unmodified (U-NET), using elastic deformation data augmentation (U-NET+EA), and using an augmented multi-subject data set (Multi-Subject U-NET+EA). Data used for RANR were ground truth values of (*Sub1*, 30°, FS); data for training the optimized neural networks are those described in Section 8.5 and Table 8.1. Data were trained, tested, and predicted only on the upper part of the arm, above the elbow; raw lower-arm intensity maps are provided for context. Although superficially smooth and well-formed, RANR segmentation poorly localizes biceps and humerus (a), resulting in poor segmentation accuracy; in contrast, neural network methods perform reliably along the middle section of the biceps (b) but segment more poorly near the ends of the structure (c). Adding elastic deformation augmentation data generally helps smooth the data and improve accuracy (d), though many artifacts remain. Note that *Sub2* Multi-Subject U-NET prediction (e) shows a scan used in network training; its high level of accuracy thus represents network memorization, and the scan is presented for completeness only. ©2019 IEEE.

pixel accuracy (calculated as the total fraction of correctly classified voxels). Registration is particularly inadequate when segmenting new subjects. In addition, as shown in Fig. 8.3, the sources of this error are fundamentally different. Registration-based segmentation generates well-formed tissue structures, but these structures are significantly misaligned, so errors are distributed along the length of the scan; in contrast, neural network methods perform well in the central belly of the muscle, but struggle with endpoints, such that errors are concentrated at the top and bottom of the muscle. Manual error correction is thus much more time-efficient on neural-network-segmented scans, as a much smaller subset of slices must be corrected. Alternatively, analysis can be restricted to the center of the muscle with a reasonable expectation of data reliability. This difference in error type also suggests future work in ensemble approaches that incorporate the successful aspects of each strategy.

Elastic Deformation Augmentation

The augmentation of the training data set with elastically deformed scans substantially improved performance in most tested cases, suggesting that this method of data augmentation is useful in tissue annotation. In fact, as shown in Table 8.1, a U-Net trained on a single scan along with the elastically-augmented version of that scan (U-NET+EA) performed almost as well on scans at new angles as a network trained on multiple angles (Multi-Angle U-NET). Gains across force conditions and subjects were more modest, but elastic augmentation still showed improved performance as compared with unaugmented counterparts. Qualitatively, elastic augmentation seems to sharpen scans and and remove some artifacts, as shown in Fig. 8.3.

The addition of rotationally augmented data did not improve performance; in fact, performance was diminished across all tested categories, most likely because both test and training scans were collected at similar poses.

A Note on Multi-Subject Training and Performance

To effectively scale up muscle deformation analysis requires generalization to new subjects. While neural network segmentation performed reasonably well in this respect, considering the small number of scans seen — even an unmodified U-Net trained on a single scan successfully segmented the middle portion of most subjects' biceps and humerus — performance remains significantly lower than when generalizing across angle or force conditions for the same subject. More explicitly, a network trained on multiple subjects performs better on new angles and force conditions than a network trained on multiple angles or force conditions performs on new subjects; thus, our objective of multi-subject data sets will be best served in the future by focusing on generation of training data across multiple subjects rather than multiple angle or force conditions.

8.6 Summary & Limitations

In this chapter, we presented a novel CNN-based pipeline for generating 3D annotated scans of bone and muscle structures and used it to annotate the biceps brachii and humerus of the OpenArm 2.0 data set developed in chapter 7. We found that this pipeline generally compared favorably with classical registration-based segmentation approaches in terms of segmentation accuracy and distribution of errors, and it saved hundreds (possibly thousands) of hours of estimated expert annotation time on this target data set. In the future, we hope to leverage this segmentation process to further expand the data we examine, which will include multiple muscles under both static and dynamic conditions, as well as greater variety in subjects' age and pathology. Beyond muscle segmentation, this work could easily be extended to measurement and modeling of other deformable tissues throughout the body.

At the same time, while we found 2D CNNs to be sufficiently expressive to generate our data set with the desired level of accuracy for preliminary modeling, segmentations remain messy in some cases, and a number of improvements to the pipeline could be made to achieve even greater accuracy (though possibly at the cost of computational efficiency), including use of a 3D U-Net [17], consensus methods across image slices collected from different dimensions, or, treating one dimension as time, techniques used in video segmentation [38].

Lastly, as with most deep-learning-based approaches and despite the use of data augmentation, our segmentation pipeline assumes a substantial body of expert-annotated training data that remains time-intensive to collect, and cannot immediately generalize to untrained additional tissue structures. Future research will explore both additional augmentation approaches and the feasibility of training generalized networks to recognize arbitrary numbers and configurations of muscles based on the appearance of fascia and the edges of bones, to allow for further expansion of the force–deformation modeling work begun in this thesis.

Part IV

Systematic Modeling

Chapter 9

Exploratory Geometric Analysis

This chapter is adapted in part from [49].[1]

In part III, we described the construction of two 3D volumetric data sets of the elbow to enable comprehensive study of muscle deformation under changes in both kinematic configuration and loading conditions. Now, in part IV, we present both preliminary geometric analyses of these data and a proposed framework for expanded modeling in the future.

In this chapter, we broaden our exploratory part I analyses of simple deformation signals to three dimensions, examining kinematic- and force-associated deformation of the biceps brachii along the length of the arm using the OpenArm 1.0 data set described in chapter 7. These analyses inform the systematic modeling framework we propose in chapter 10.

In the sections below, we first define the scope of our analyses and examined deformation signals, then quantify these signals' spatial variation on the OpenArm 1.0 data set. Lastly, we summarize the limitations of these analyses and their implications in subsequent systematic modeling and control signal extraction, including insights on the best spatial locations from which to reliably measure force-associated deformation signals.

9.1 Study Scope

While the OpenArm data sets provide volumetric data that supports explicit 3D modeling (e.g., finite element analysis), we focus our analysis on signals relevant for use in real-time control and biomechanical modeling — i.e., those that are correlated with force output, can be measured from a single 2D B-mode ultrasound transducer (or even 1D A-mode transducer), can be tracked in real time without extensive computation, and are robust to slight variation in sensor location. Thus, we start our examination with simple cross-sectional

[1] ©2018 IEEE. Reprinted, with permission, from Laura A. Hallock, Akira Kato, and Ruzena Bajcsy. "Empirical Quantification and Modeling of Muscle Deformation: Toward Ultrasound-Driven Assistive Device Control". In: *IEEE International Conference on Robotics and Automation (ICRA)*. IEEE. 2018, pp. 1825–1832. DOI: 10.1109/ICRA.2018.8462887.

signals (area, etc., as defined below) similar to those explored in chapter 3 and examine their behavior along the proximal-to-distal length of the biceps brachii (the elbow flexor for which the cleanest and most complete segmentation was obtained). We first examine these signals' behavior across pure, unloaded changes in kinematic configuration (i.e., elbow angle), then examine, at each angle, changes under loading, remarking on these measures' consistency with the simple ellipsoidal deformation expected based on our understanding of the deformation mechanism (as described in section 2.5).

9.2 Definition of Deformation Signals

The following three deformation measures (noted spatially on a representative scan slice in Fig. 9.1) were extracted from the segmented biceps brachii at each cross section of each of the nine segmented OpenArm 1.0 scans and analyzed along the axis of the humerus (hereafter denoted x).

Cross-sectional area (CSA). Similarly to the measure in chapter 3, we define *CSA* as the area of the muscle cross section sliced perpendicular to the length of the humerus.

Thickness (T). While thickness was shown in chapter 3 to be correlated with output force when measured at a single cross section of the brachioradialis, it is not obvious how best to define this value spatially across muscle cross sections, since defining it relative to any single axis along the humerus cannot accommodate the nonlinear deviation of the muscle along its length. For these preliminary analyses, we define thickness T as the mean thickness of the muscle measured from the anterior arm surface down toward the humerus within a 1 cm window surrounding the centroid of the measured muscle cross section, a region selected to average out irregularities resulting from the manual annotation process.

Eccentricity (E). We define the eccentricity E of the muscle cross section as the condition number of the covariance matrix of the cross section when treated as spatially-varying point data. Intuitively, the condition number represents the ratio of the major to minor axes of the best-fit ellipse to the cross-sectional area and is thus a measure of muscle eccentricity along the length of the arm. This measure is analogous to an inverted version of the aspect ratio measurement explored in chapter 3, but incorporates point data from the entire cross section rather than only the points of widest extent in each dimension.

9.3 Spatial Deformation Analysis

Based on the simple force generation mechanism outlined in section 2.5 — specifically, that because muscles are isovolumetric, force-related length changes occur alongside corresponding cross-sectional area changes — we expect that, when measured in an appropriate location along the muscle, both passive elbow flexion and force exertion that result in muscle shortening will result in corresponding increases in cross section size in a roughly ellipsoidal manner.

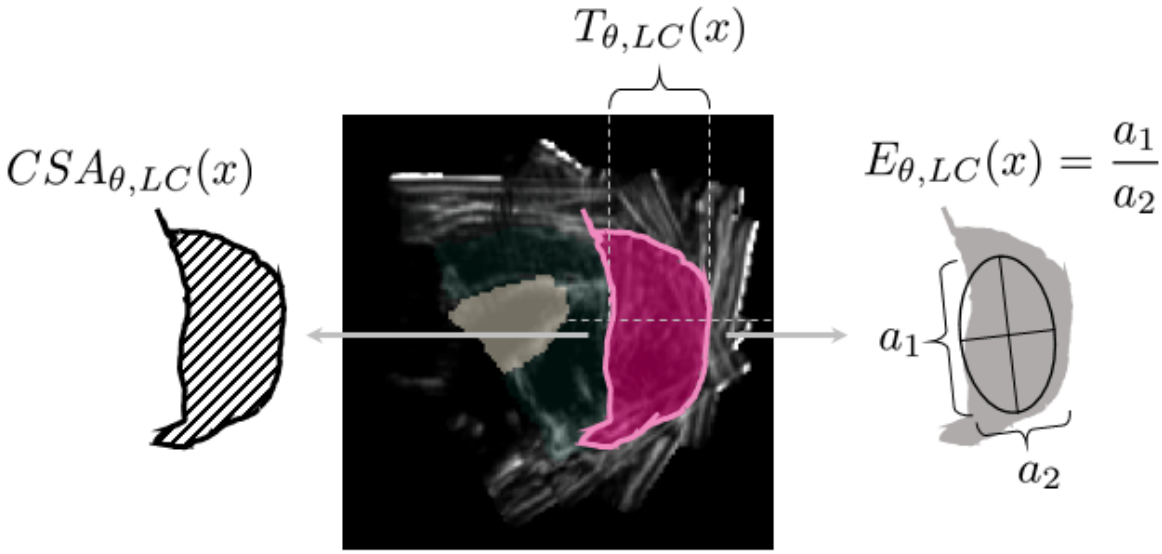

Figure 9.1: Visualization of deformation measures calculated for spatial geometric analysis at each elbow angle θ, loading condition LC, and location x along the length of the biceps brachii on a single representative coronal cross section. Cross-sectional area $CSA_{\theta,LC}(x)$ was directly extracted from the segmented volumetric biceps volume, and thickness $T_{\theta,LC}(x)$ was computed as the mean of the measure shown about a 1 cm region surrounding the centroid of the computed area. Eccentricity $E_{\theta,LC}(x)$ was computed as the major-to-minor-axis ratio of the best fit ellipse to the cross-sectional spatial data values in the least-squares sense. The data collected suggest that *CSA* changes are consistently associated with muscle output force, and that thickness measurements T, when combined with a model of eccentricity E, could be used to estimate the *CSA* signal using cheaper 1D A-mode ultrasound sensors. ©2018 IEEE.

In this section, we investigate the extent to which this model is evident in our prototype OpenArm 1.0 data set via spatial analysis of the deformation metrics defined above.

Kinematic-Related Deformation

We first analyze purely kinematic-associated muscle deformation, examining the spatial behavior of *CSA*, T, and E measures under fully supported loading conditions. Based on the simple ellipsoidal model discussed above, as well as qualitative observation of the multi-angle data set explored in chapter 3, we expect maximum *CSA* and T measures to increase with increased flexion angle. (The behavior of the E signal in this early data is more difficult to parse, so we do not offer any particular hypothesis for its behavior.)

The top and middle plots of Fig. 9.2 show raw *CSA* and T data, respectively, for the biceps brachii under the fully supported condition at all three angles examined, as well as the best-fit quadratic regressor for the data (in the least-squares sense), from the proximal extremum of the observable biceps volume to the location of the elbow (i.e., where the humerus meets the radius and the ulna).

Examining both raw data values and these quadratic regressors, we indeed observe the anticipated increase in maximum *CSA* and T values with increased elbow flexion. In addition, despite messily segmented muscle endpoints, the regressors reflect the muscle's isovolumetricity: the model for the 30° condition is not only shortest in the *CSA* dimension, but longest in the x, and the 90° model is both tallest in the *CSA* dimension and shortest in the x. The fitted quadratic functions consistently peak at a distance of 7.4 cm (46%) from the proximal end of the muscle across all angle conditions (all within a range of 0.3 cm), and no obvious deviation from this mean is observable under load; this is interesting given that when observing natural flexion motion, the biceps muscle belly appears to slide from distal to proximal, and we might expect maximum *CSA* and T measures to do the same. We do not observe these trends in the measured T data, aside from an expected lower thickness value at the most extended elbow 30° near the muscle belly at around 6 cm; we conjecture that this is because our thickness measure does not sufficiently incorporate the 3D nature of the muscle's line of action, and can be substantially impacted by the muscle's deviation from the humerus-aligned scans we consider. It's also possible that, unlike the trends we observed in the brachioradialis muscle in chapter 3, thickness alone is not a reliable indicator of biceps length changes.

Nevertheless, thickness T remains a desirable signal for deformation measurement in that it can be measured with cheap, simple sensors like 1D A-mode ultrasound transducers, and we hypothesize that they could be used as a proxy for more representative *CSA* values when combined with an offline model of eccentricity values E (as *CSA* can be roughly calculated when T and E are known). The final plot of Fig. 9.2 characterizes this E value for our prototype data set across all examined joint angles and loads, showing it to be relatively consistent across experimental conditions, with narrow, eccentric muscle endpoints and a wider, rounder belly, indicating that construction of a low-dimensional eccentricity model may be possible. Once appropriately fitted to the data, a system leveraging such an

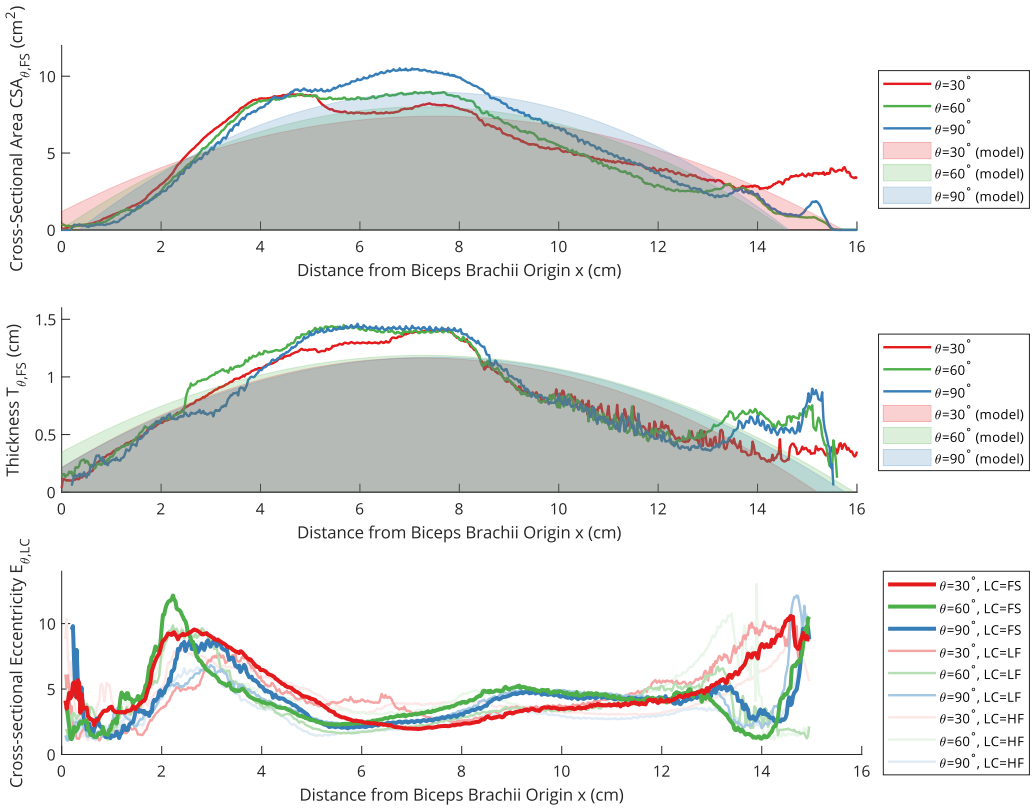

Figure 9.2: Variation in cross-sectional area *CSA* (*top*), thickness T (*middle*), and eccentricity E (*bottom*) along the length of the biceps brachii from shoulder to elbow of the prototype OpenArm 1.0 subject, under multiple elbow flexion angles with the forearm fully supported (*CSA* and T) and at all tested conditions (E). The location of maximum CSA/T (as measured from the corresponding quadratic regression models shown in overlay) was not shown to vary with angle, but the changes in width of the fitted *CSA* quadratics reflect the compression of the muscle as elbow flexion increases, a preliminary and intuitive insight that suggests that building low-dimensional predictive models of *CSA* change may be possible. The steep increase in eccentricity E near the muscle's ends is reflected in the cross-sectional images shown in Fig. 7.2, and the consistent shape of the eccentricity map across all tested conditions indicates that a spatial eccentricity map could be of use in developing a predictive model of muscle *CSA* from 1D thickness T data. ©2018 IEEE.

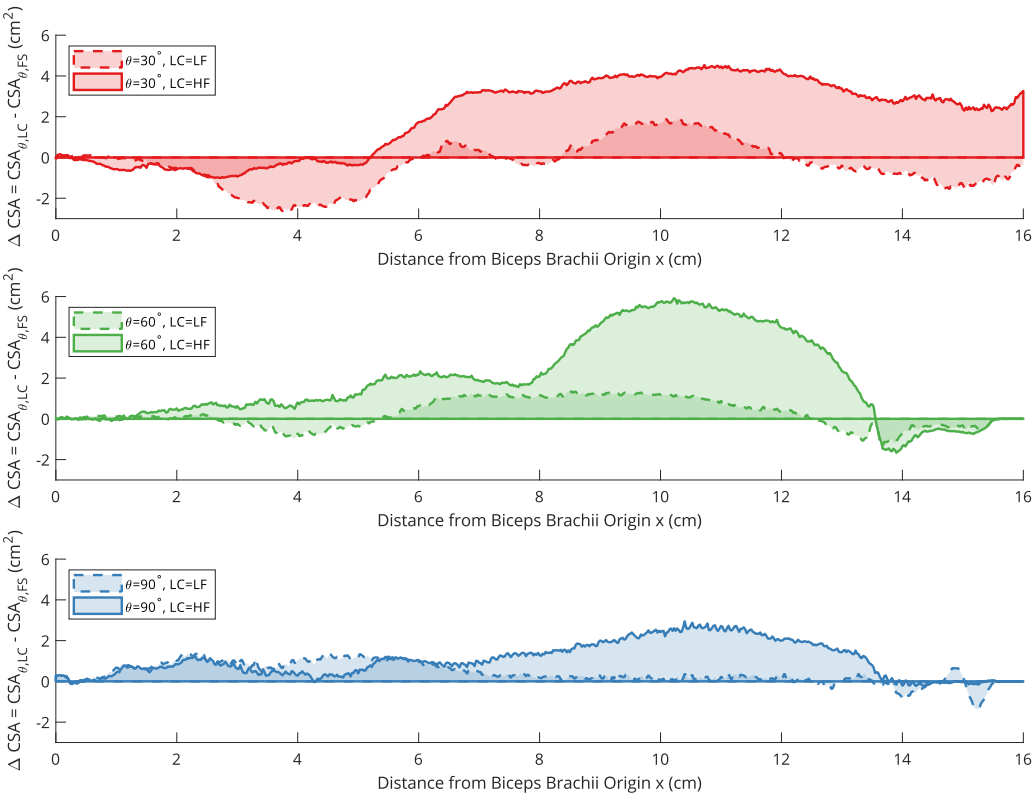

Figure 9.3: Spatial variation of change in biceps brachii cross-sectional area (*CSA*) from that of the fully-supported (FS) volume under low (LF, *dashed*) and high (HF, *solid*) loading conditions. Significant variation is consistently observed in a range centered approximately 2.6 cm distal from the location of maximum absolute *CSA* and is larger under higher loading at each configuration. These observed ΔCSA values thus indicate a candidate location from which to extract a spatially robust assistive device or biomechanical modeling control signal. ©2018 IEEE.

eccentricity model and real-time thickness data could form the bases of a *CSA*-estimating control scheme.

Force-Related Deformation

We next examine force-associated deformation — the key relationship we seek to model for enhanced inference of human dynamics. For initial analysis, we consider this force–deformation relationship separately at each kinematic configuration.[2]

Fig. 9.3 shows the spatially-varying change in muscle *CSA* under low- and high-force conditions (LF and HF, respectively) detailed in section 7.3 and Table 7.1, defined as the *CSA* under loading less the *CSA* while fully supported (FS) at each spatial location. The observed changes peak consistently across configurations and loading conditions at a distance of approximately 10 cm from the proximal end of the muscle, roughly 63% of the way down the muscle and 2.6 cm (16%) distal from the peak of the absolute maximum *CSA* shown in Fig. 9.2. Overall, substantial deformation — up to 5.9 cm^2, or 54% of the corresponding unloaded *CSA* value — is observed.

While the magnitudes of this *CSA* change are not comparable across configurations, as the same wrist weights induced different moments at each angle, greater changes in *CSA* are observed for larger loads than those of smaller loads at all angle conditions within a considerable region (8.8 cm, or 55% of the total examined biceps length, roughly centered about the location of maximum *CSA* change, with much larger regions in the 30° and 60° cases). This suggests that an assistive device control signal could be gathered from this location and remain reasonably robust under moderate levels of sensor movement. Interestingly, this location of maximum *CSA change* is not centered about the location of maximal *CSA overall*, which is consistently approximately 2.6 cm more proximal, as illustrated in Fig. 9.2, illustrating a breakdown of the simple ellipsoidal model. We hypothesize that this inconsistency is related to contact forces with nearby tissue structures — in particular, with the brachialis, which also provides substantial flexion force and presses upward from below — but more analysis on a wider subject cohort will be necessary to confirm this conjecture.

Although changes in muscle thickness across loading conditions were similarly defined and examined, the data proved extremely noisy, and observed signal changes were small (on the order of 0.4 cm) and difficult to characterize, with no discernible peak or trend across loading conditions. As with our kinematic-related deformation analysis above, this suggests that the thickness signal as defined neglects substantial muscle deformation in the unexamined dimension, and may therefore be most useful in concert with a model of muscle

[2]Note that the nature of our imposed force conditions — weights applied at the wrist — does not allow us to disambiguate between muscle exertions that increase output force and those that increase stiffness or impedance. (Indeed, at the 90° condition, effectively all exertions relate to stiffness, as the loading vector points along the length of the forearm, as discussed in section 7.4.) Here, we describe both of these exertions as being "force-related," but target this disambiguation for future study (e.g., using the OpenArm 2.0 data set, which imposes force conditions in a more precisely specified manner, and by building new data sets that examine both agonist and antagonist muscles together).

eccentricity, as described above. We therefore plan to undertake a more thorough study of these models, and also of the most effective ways to gather a 1D muscle thickness signal (e.g., by considering different potential positions for an ultrasound sensor about the coronal plane of the arm) in future investigations.

9.4 Summary & Limitations

In this chapter, we undertook a preliminary geometric analysis of the kinematic- and force-associated variation of several simple deformation signals on the biceps brachii of a prototype individual from the OpenArm 1.0 data set. We found these data to be largely consistent with a simple isovolumetric, ellipsoidal model of deformation — despite several inconsistencies likely caused by tissue contact dynamics and other geometric complexities — and located a promising spatial location at which force-related deformation changes reliably appear, suggesting an appropriate region from which to extract a control or modeling signal.

Due to the complex sliding and contact dynamics involved, expanded analysis of configuration-related deformation in particular is likely to require an especially complex model; at the same time, it may prove less necessary to use this model explicitly during real-time device control, as a number of sensors already exist (including motion capture, electrogoniometers, and inertial measurement units) that can measure real-time kinematics. Our ultimate objective is the extraction of force-related signals — which, given an offline model of configuration-associated deformation, may admit a clean and low-dimensional formulation, allowing for real-time, high-dimensional measurement of human dynamics.

These exploratory single-subject analyses inform our systematic modeling work in chapter 10, in which we present an example framework for incorporating both types of deformation in a principled manner while also considering more complex known biological phenomena impacting the signal.

Chapter 10

Systematic Force–Deformation Modeling

In the previous chapter, we presented a simple geometric analysis of kinematic- and force-associated muscle deformation and found it to be largely consistent with ellipsoidal muscle deformation; at the same time, our results showed substantial spatial irregularities in the deformation signal, making it difficult to determine precisely what model formulations would best describe the force–deformation relationship. In this chapter, leveraging these basic insights, we propose a novel theoretical framework to systematically examine a wide variety of possible force–deformation models while considering both empirical data and literature-based assumptions to varying degrees.

In the sections below, we first motivate the creation of this modeling framework and define its scope, then present an example suite of models — verifiable by the chapter 7 OpenArm 2.0 data set — relating biceps deformation and elbow angle to output torque at the elbow. Lastly, we comment on the model selection process and on the applications of both the models themselves and our wider strategy in conceptualizing the black box/white box trade-off, as well as limitations of this approach.

10.1 Model Scope

As with all system identification problems, the "best" model is defined by its final application.[1] In this thesis, we are primarily concerned with models suitable for integration in real-time device control tasks and biomechanical modeling, and thus scope our proposed models accordingly, avoiding computationally intensive techniques like finite element analysis and models that require 3D inputs that cannot be collected in a single time instant.

Even limiting our target applications in this manner, there exist a multitude of possible force–deformation model types, ranging from purely data-driven and unstructured "black

[1] "*All models are approximations. Essentially, all models are wrong, but some are useful. However, the approximate nature of the model must always be borne in mind.*" —George E. P. Box [13]

box" models to "white box" models that impose known physiological relationships with minimal free parameters. In the field of human modeling, there exist centuries of study from which to draw these physiological modeling assumptions; at the same time, given the geometric, mechanical, and neurological complexity of the neuromusculoskeletal system discussed at length in section 2.3, as well as the substantial morphological differences across individuals, it is not at all evident which assumptions hold, and under what conditions. This presents a modeling trade-off: while fully data-driven models offer little generalizability or physiological insights, introduction of erroneous values and assumptions from literature may adversely affect model quality in unpredictable ways.

Thus, rather than advocating for a single model relating muscle force, deformation, and configuration, we present a suite of possible models along this black box–white box continuum that make this empirical/theoretical trade-off in a principled manner, allowing for equitable comparison across candidate models. This model series constitutes both a readily implementable approach for immediate control signal extraction (compatible with, for example, the trajectory tracking control task explored in chapter 5) and an example of our more abstract modeling contribution: a systematic, quantifiable way to approach the black box/white box trade-off.

10.2 Example Modeling Framework

In this section, we present an example modeling framework mapping muscle deformation and kinematic configuration to output force at the elbow, designed for calibration and validation with the OpenArm 2.0 data set described in chapter 7. Specifically, in the subsections below, we detail a suite of models with common inputs biceps deformation D_{biceps} and elbow angle θ and common output elbow joint torque τ_{ext}, as shown in Fig. 10.1–10.4, which make varying trade-offs between relying exclusively on empirical data and building in assumptions from biomechanics literature. Note that while we envision D_{biceps} as being extracted from a single ultrasound cross section (to generate models readily applicable to real-time control scenarios), many of the model classes we describe support arbitrary parameterizations of this deformation in arbitrary dimensions, and the D_{biceps} signal — alongside more precisely defined θ and τ_{ext} — can be extracted from the OpenArm 2.0 data set in numerous ways.

While we formulate this framework primarily as a mapping from D_{biceps} and θ to τ_{ext} as noted, later, more structured models also enable inference of individual muscle forces — a key future objective of this line of modeling inquiry — as discussed below (though model calibration and validation of such extracted force values will likely require additional sensor data, like that proposed in section 13.2).

"Model-Free" Baseline

We begin with a "black box" baseline model, in which joint angle θ and deformation D_{biceps} are mapped directly to output torque τ_{ext} via arbitrary function $f_0(\cdot, \cdot)$ using only empirical

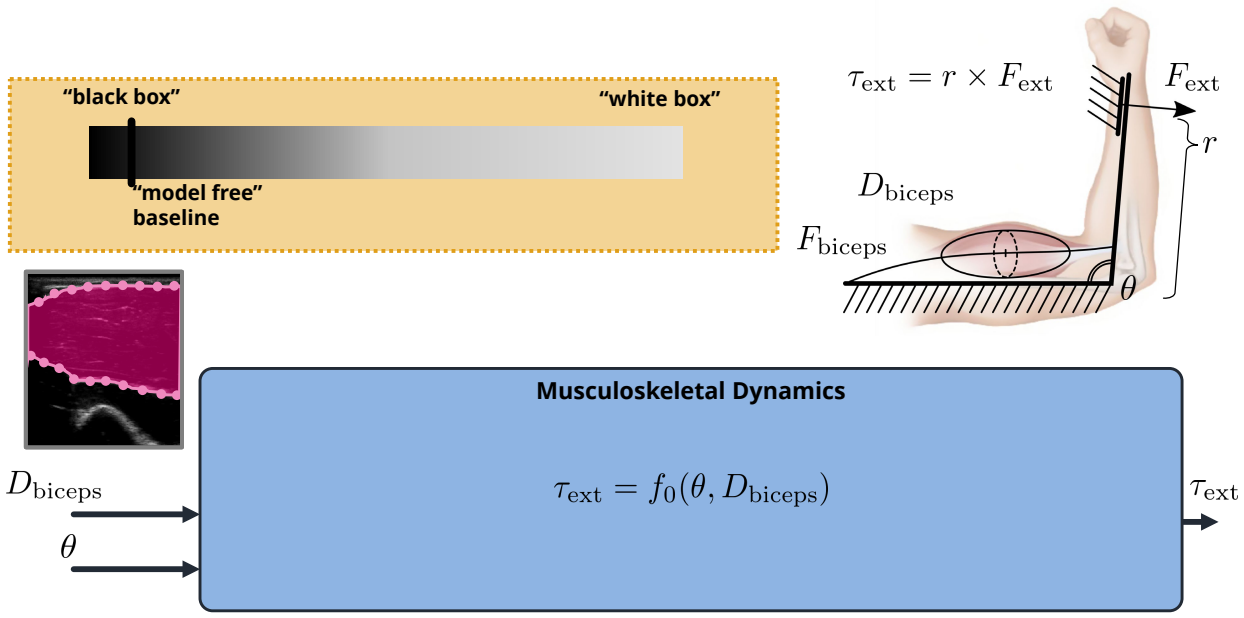

Figure 10.1: Fully empirical "black box" baseline model mapping elbow angle θ and arbitrarily-parameterized muscle deformation D_{biceps} to elbow torque τ_{ext} through arbitrary function $f_0(\cdot, \cdot)$. This model enables data-driven prediction of output torque but lends little to no insight into underlying physiology. Image adapted from [67].

data, as shown in Fig. 10.1. This model class supports a wide array of possible deformation and function parameterizations — D_{biceps}, for example, could be any of the deformation signals discussed in previous chapters, the location of contour points, or even raw pixel values, and $f_0(\cdot, \cdot)$ could be anything from a linear model to a neural network.

Adding Multi-Muscle Dynamics

We can impose structure on the above black box baseline model by incorporating the structure of multi-muscle dynamics, as shown in Fig. 10.2. In this model class, joint angle θ and deformation D_{biceps} are again mapped through arbitrary function $f_1(\cdot, \cdot)$ in an equally data driven manner (supporting all of the above possible parameterizations), but are instead mapped to biceps force F_{biceps}, which is now modeled as a fraction of the total output torque measured at the elbow. Note that this model requires assuming several values and properties from literature — either function $c(\cdot)$, a configuration-dependent multiplier for the percentage of force attributable to the biceps, or the actual torque contributions of all other elbow flexors and extensors (τ_{brach}, τ_{brachrad}, τ_{triceps}) and that of the forelimb's mass (τ_{mg})

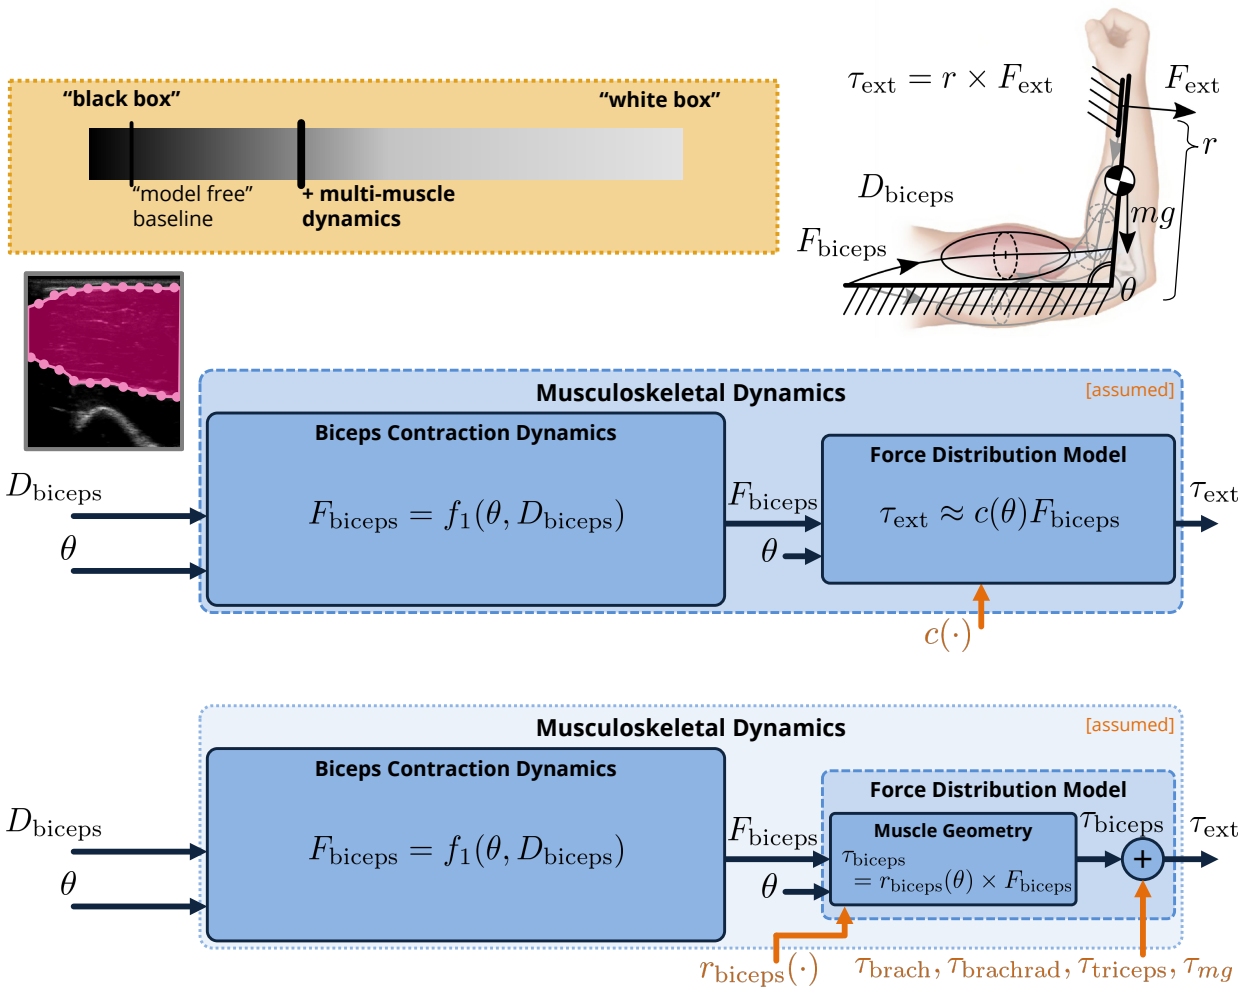

Figure 10.2: Imposing multi-muscle structure on the baseline model in Fig. 10.1 enables extraction of biceps force F_{biceps} and is more representative of real physiology, but requires knowledge of either the configuration-varying percentage of elbow force attributable to biceps exertion $c(\cdot)$ (*top*) or actual torque contributions of other actuating muscles τ_{brach}, τ_{brachrad}, and τ_{triceps} and mass contribution τ_{mg} (*bottom*). Using the currently available OpenArm 2.0 data set, these values must be assumed from literature, but as data sets expand, they too could be determined from corresponding muscle deformation signals. Image adapted from [67].

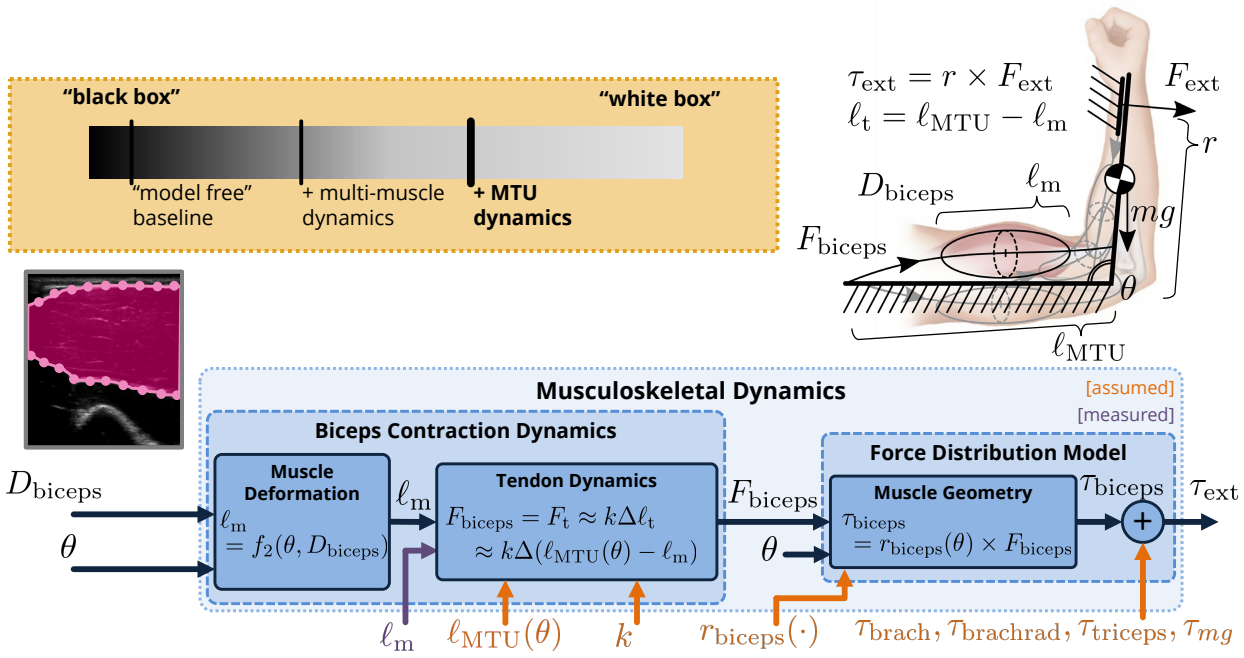

Figure 10.3: The Fig. 10.2 model above can be further refined by incorporating dynamics of the muscle–tendon unit (MTU), assuming springlike behavior of the tendon with spring constant k. This model class requires additional assumption of both k and kinematically varying MTU length $\ell_{\text{MTU}}(\cdot)$ but can be partially validated through measurement of muscle length ℓ_{m} on the OpenArm 2.0 data set. Image adapted from [67].

along with (possibly kinematically varying) biceps moment arm $r_{\text{biceps}}(\cdot)$. Alternatively, as our deformation data sets expand, these additional muscle torque values could be inferred from their corresponding deformation signals.

By imposing this structure, we can now use this model (and all models below) directly for individual muscle force inference if desired, treating F_{biceps} as our output.

Adding Muscle–Tendon Unit Dynamics

We can further impose structure on the above models by incorporating the dynamics of the muscle–tendon unit (MTU), as shown in Fig. 10.3, mapping θ and the (still arbitrarily parameterized) D_{biceps} to muscle length change ℓ_{m} via arbitrary function $f_2(\cdot, \cdot)$ and modeling the tendon length change as that of a stiff spring with constant k. This requires assuming several additional values from literature, including the total (kinematically varying) MTU length $\ell_{\text{MTU}}(\cdot)$ and tendon spring constant k. Note that models in this class can be partially validated on (or fitted to) the OpenArm 2.0 data set by extracting muscle length ℓ_{m} directly

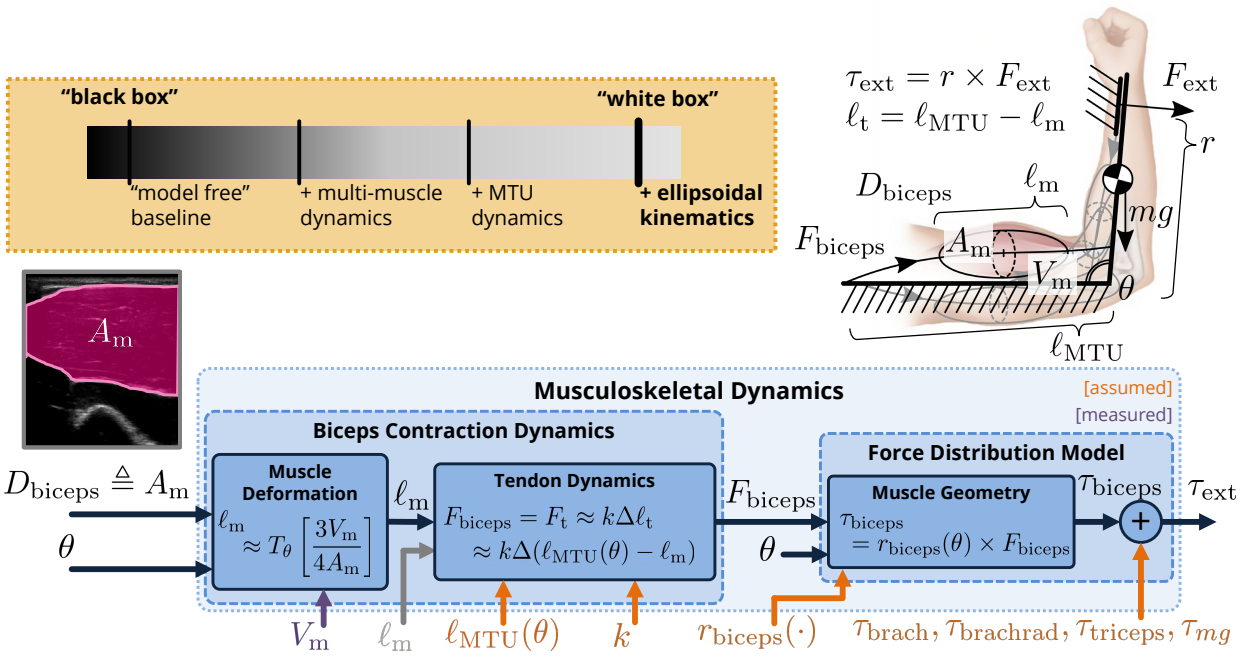

Figure 10.4: Isometric ellipsoidal muscle kinematics can be imposed on the Fig. 10.3 model above to generate a fully "white box" model requiring little to no parameter fitting. This additional structure enforces that deformation is specifically parameterized as cross-sectional area A_{m} and requires knowledge of (assumed constant) muscle volume V_{m}, which can be extracted from the OpenArm 2.0 data set. Image adapted from [67].

from the full volumetric scans.

Adding Ellipsoidal Deformation Kinematics

Lastly, we can impose the isometric ellipsoidal deformation kinematics implied by the mechanism described in section 2.5 to generate a fully "white box" model, as shown in Fig. 10.4, defining D_{biceps} specifically as muscle cross-sectional area A_{m} and assuming purely ellipsoidal deformation, such that muscle length ℓ_{m} is always equal to $3V_{\text{m}}/(4A_{\text{m}})$ (for constant muscle volume V_{m}) when viewed in appropriate reference frame T_θ. This requires knowledge of V_{m}, which can be extracted from the OpenArm 2.0 data set volumes or assumed from literature.

10.3 Thoughts on Model Selection

As discussed above in section 10.1, the suitability of a particular model on the continuum discussed above is completely application dependent, and its quality can be evaluated on a number of disparate factors (time series error, generalizability across individuals or conditions, etc.). For device control applications, for example, data-driven models may be sufficient; for biomechanical study, however, and for applications in which muscle exertions are evaluated to prevent injury or maintain comfort (e.g., exoskeletal control), physiology-based models are critical. While in the short term, lacking a comprehensive muscle force–deformation body of literature (and especially given individual morphological variation), we may be forced to rely primarily on black box techniques, we ultimately advocate for leveraging as much modeling structure as possible to enable safety, interpretability, and novel scientific understanding.

10.4 Summary & Limitations

In this chapter, we presented a theoretical framework to enable the consideration of a range of muscle force–deformation–configuration models incorporating both empirical data and existing phyiological modeling assumptions and evaluation of their relative performance in a quantifiable manner. While implementing these models in the context of a particular application is beyond the scope of this thesis, we are currently developing software infrastructure to allow for comparison of the models above on a number of performance metrics (e.g., mean prediction error, generalizability across subjects, sensitivity to probe location) on the OpenArm 2.0 data set developed in chapter 7. Ultimately, we envision the wide application of these and similar models — as well as our overall approach to quantifying the black box/white box trade-off — in the many domains discussed in our concluding chapter 13, including assistive device control and enhanced understanding of the human motor control system, both healthy and pathological.

Part V

Data, Codebases, Alternative Sensing, & Conclusions

Chapter 11

The OpenArm Codebase

This chapter delineates all current releases of the OpenArm data set, released open-source on the SimTK database at https://simtk.org/projects/openarm. Beyond the work in this dissertation, the OpenArm data sets are designed to enable study of force- and kinematic-induced muscle deformation for applications in biomechanics research, computer graphics, and assistive device development, and are presented as a core thesis contribution. We invite anyone in the research community to use the OpenArm and OpenArm Multisensor data sets to validate existing muscle deformation models or to devise new ones.

Each package, along with all available downloads at time of publication, is enumerated below.

11.1 OpenArm Multisensor 1.0

The OpenArm Multisensor 1.0 package [51] contains the first iteration of the OpenArm Multisensor data set and associated muscle contour tracking code. Included are *1)* a multi-subject data set of ultrasound-based time series deformation data of the brachioradialis muscle alongside surface electromyography (sEMG), acoustic myography (AMG), and force data, and *2)* all code used in analyzing deformation across subjects and optical flow tracking of muscle deformation over time.

Specifically, the data release includes:

- *Data set.* All OpenArm Multisensor data, including raw and segmented ultrasound frames used in brachioradialis tracking and processed time series data.
- *Source code.* GitHub repository of all optical flow tracking and analysis code.

11.2 OpenArm Multisensor 2.0

The OpenArm Multisensor 2.0 package [50] contains the second iteration of the OpenArm Multisensor data set and associated muscle contour tracking code. Included are *1)* a multi-

subject data set of ultrasound-based time series deformation data of the brachioradialis muscle alongside surface electromyography (sEMG), force, and goal trajectory data, and *2)* all code used in analyzing deformation across subjects, optical flow tracking of muscle deformation over time, and real-time visualization of time series data.

Specifically, the data release includes:

- *Data set.* All OpenArm Multisensor data, including raw and processed ultrasound frames used in brachioradialis tracking, processed time series data, demographic data, and survey results on user preferences during a trajectory tracking task.
- *Ultrasound streaming & tracking source code.* GitHub repository of all real-time streaming and optical-flow-based muscle thickness tracking code.
- *Visualization & data collection source code.* GitHub repository of all real-time signal visualization and data collection code.
- *Analysis source code.* GitHub repository of all analysis code.

11.3 OpenArm 1.0

The OpenArm 1.0 package [49] contains the first iteration of the OpenArm data set, a set of full volumetric scans of the human arm collected using ultrasound and motion capture. Data are factorial under multiple elbow angles and loading conditions, and thus allow for separable analysis of force- and configuration-associated muscle deformation.

Specifically, the data release includes:

- *Data set.* full volumetric data of the arm for three subjects (*Sub1*, *Sub2*, and *Sub3*) under 20 force and elbow angle conditions. A subset of nine *Sub1* scans are annotated with the full visible volumes of the humerus and biceps brachii, as well as sections of the ulna, radius, deltoid, brachialis, and brachioradialis.

11.4 OpenArm 2.0

The OpenArm 2.0 package [113] contains the second iteration of the OpenArm data set, a multi-subject set of full volumetric scans of the human arm collected using ultrasound and motion capture. Data are factorial under multiple elbow angles and loading conditions, and thus allow for separable analysis of force- and configuration-associated muscle deformation. Improvements from OpenArm 1.0 include improved data collection procedures allowing for more explicit comparison of force conditions across angles, more subjects, and partial annotations for all subjects enabled by neural-network-based segmentation.

Specifically, the data release includes:

- *Data set.* Full volumetric data of the arm for ten subjects under 20 force and elbow angle conditions, as well as 12 scans from one additional subject. Ground-truth segmentation data of the biceps brachii and ventral surface of the humerus are included for one full set of subject scans, as well as several scans of other subjects. Partial neural-network-generated segmentation data are included for all subjects and scans.

11.5 Annotation Source Code

The Annotation Source Code package [113] contains all code and neural network models used to generate tissue segmentations of the biceps and humerus for the OpenArm 2.0 data set, using both convolutional neural networks and classical image registration. Code and models used in network training, prediction, analysis, and data augmentation are included.

Specifically, the data release includes:

- *Neural network models.* ZIP archives of all neural network models used in segmenting the OpenArm 2.0 data set, including all models used in the associated publication [113] and the best performing model overall at time of release.
- *Segmentation source code.* GitHub repository of all segmentation code.

Chapter 12

Alternative Sensing: Force–Vibration Correlation via Acoustic Myography

This chapter is adapted in part from [48].[1]

In parts I–IV, we explored the use of ultrasound-extracted deformation signals for force inference and device control. In this chapter, we present preliminary work leveraging an alternative sensing technology — acoustic myography, or AMG — to measure muscle vibration during contraction to perform similar inference of musculoskeletal dynamics.

In the sections below, we first present the scope of this (highly preliminary, single-subject) study and provide background on existing models relating muscle force and vibration. We then present a proof-of-concept model of the elbow joint for examining force–vibration correlation under isometric elbow flexion and show early evidence that these values are indeed correlated. Lastly, we remark on the study's limitations and implications for future device control and biomechanical modeling.

12.1 Study Scope

Building on current scientific understanding of the force–vibration relationship, we present a simple, novel model relating output force at the elbow to AMG-measured vibration from both agonist (biceps) and antagonist (triceps) muscles, whose interaction is critical to modulating joint stiffness during everyday tasks and largely ignored in existing musculoskeletal modeling frameworks. We then probe this model's validity during isometric contraction via a preliminary single-subject study.

[1] ©2018 IEEE. Reprinted, with permission, from Laura A. Hallock and Ruzena Bajcsy. "A Preliminary Evaluation of Acoustic Myography for Real-Time Muscle Force Inference". In: *International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*. IEEE. 2018.

12.2 Existing Muscle Force–Vibration Measurement & Models

In this section, we describe existing sensing paradigms used in muscle vibration measurement, competing theories for the mechanistic source of this vibration signal, and current understanding of its relationship to muscle output force.

Muscle Vibration Measurement

While the fact that skeletal muscles produce vibrations during contraction has been known for centuries [43, 152], the vibration signal has only recently been used in scientific studies, enabled by the development of appropriate sensors (including sufficiently sensitive piezoelectric transducers and condenser microphones) to measure the low-amplitude, low-frequency signal [114]. The term acoustic myography (AMG) was first used to describe the measurement of these vibration-induced pressure waves in [7]; other terms, including vibromyography [75, 8] and mechanomyography [114] have been used to describe the acquisition of the same signal. These terms are frequently conflated, though some literature draws a distinction based on the specific modality used (microphones, accelerometers, etc.); all terms above refer to the measurement of the same underlying vibrational phenomenon, and we will here employ the term AMG to describe our own signal acquisition procedure.

Competing Theories of AMG Signal Source

Despite advances in sensing technology, there remain three major competing theories for precisely what phenomenon causes the vibratory signal. The *cross-bridge cycling* theory [115] states that the vibration is a direct product of the myosin-actin cross-bridge cycle that causes each sarcomere to contract. Proponents support this theory by relating the dominant frequency of muscle sounds during isometric contractions (both voluntary and elicited by electrical stimulation) with the rate of cross-bridge cycling; however, they provide no model-based evidence to support this conclusion.

The *vibrating string* theory [36] states that the muscle sound pressure can be explained by the model of a vibrating string. Proponents' studies [36] have shown that plucking of isolated muscle produced decaying pressure oscillations in which the first half cycle was largest, as would be expected from a string under tension.

The *unfused motor unit* theory [86] states that there exists a relationship between motor unit recruitment strategies and associated force transients and the AMG signal at different levels of voluntary effort. Citing earlier work in [71], proponents [86] showed that the vibratory signal of a single motor unit can be recorded at very weak contractions using an accelerometer, and that these values were consistent with their proposed model.

Each of the three theories above offers multiple concrete predictions about the nature of the AMG signal; based on these predictions, [147] performed thorough discriminative studies

on both human and animal models, both electrically stimulated and voluntary, including a fatigue study, and concluded that the unfused motor unit theory was best supported by the available evidence. Our preliminary results below are largely consistent with these findings.

Muscle Force and the AMG Signal

Muscle force is a function of both the number of fibers recruited and the firing rate of each motor unit; while these phenomena are difficult or impossible to disambiguate from traditional surface electromyography (sEMG)-based inference methods, they are readily observable via AMG in the signal's amplitude and frequency, respectively [53]. This knowledge forms the basis of the preliminary force–vibration model explored below.

12.3 Proof-of-Concept Model

In this section, we define a proof-of-concept model based on the above understanding of the force–vibration relationship relating AMG-measured vibrational signals to output torque at the elbow. We then pose several hypotheses implied by this model and use them to interrogate the model's validity on a prototype subject.

Model Definition

We can model the observed AMG amplitude A as approximately proportional to the number of activated fibers n, while the AMG frequency ν is approximately proportional to the mean fiber force $\bar{F}_f$ [53]. The output force F_m of a muscle can then be written as

$$F_m = n\bar{F}_f = \alpha A\nu$$

for some (positive) constant α.

Although this model is difficult to validate in vivo, as measurements of external forces admit an infinite range of possible muscle synergies (as discussed in section 2.3), we perform preliminary validation analysis using the simplified sagittal model of the elbow shown in Fig. 12.1. Assuming a static configuration of the elbow, the following two hypotheses should hold if the model is valid: (1) under relaxed conditions (i.e., minimal elbow stiffening), $A_1\nu_1$ of the elbow flexor(s) should correlate positively with output torque τ; and (2) for a given τ, under varying elbow stiffness, $A_1\nu_1$ of the elbow flexor(s) should correlate positively with $A_2\nu_2$ of the elbow extensor(s) to maintain a constant output torque. We evaluate these hypotheses on a preliminary data set from a prototype subject below.

Subject Biometric Data & Consent

Data were collected from the right (dominant) arm of a healthy female subject, age 24, of mass 70 kg and height 1.8 m. The study protocol was approved by the University of

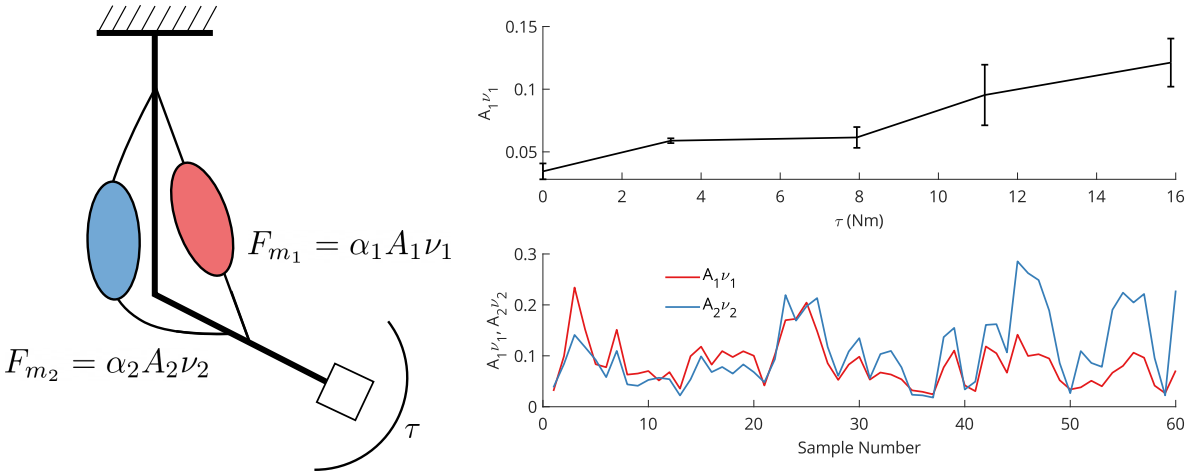

Figure 12.1: Preliminary acoustic myography (AMG) data of the biceps and triceps show substantial correlation with muscle output force. *Left*: Simplified sagittal model of the elbow used in data analysis. *Right, top*: $A_1\nu_1$ of the biceps is highly correlated with output output torque τ ($r = 0.9, p < 10^{-6}$). *Right, bottom*: Example $A_1\nu_1$ and $A_2\nu_2$ trajectories (of the biceps and triceps, respectively) during random elbow stiffness modulation, showing significant correlation between the two data series ($r = 0.6, p < 10^{-7}$), consistent with maintaining constant output torque. ©2018 IEEE.

California Institutional Review Board for human protection and privacy under Protocol ID 2016-01-8261 (first approved 4 April 2016) and written informed consent was obtained.

Data Collection & Preliminary Results

Muscle vibration data were collected from the biceps and triceps brachii of the single subject above using a Myodynamik CURO AMG collection system (MyoDynamik ApS, Bagsværd, Denmark), under multiple loads τ at the same elbow angle, while encouraging the subject to modulate their muscle stiffness at 3 s intervals over which $A_1\nu_1$ and $A_2\nu_2$ were calculated. The data were consistent with hypotheses (1) and (2) above, as shown in Fig. 12.1, suggesting that $A\nu$ is well-correlated with muscle force.

12.4 Summary & Limitations

The data above show extremely preliminary evidence that muscle vibration, as measured via AMG, is well correlated with muscle output force under isometric elbow flexion, suggesting that AMG is a feasible and useful modality for assistive device control and biomechanical study. Because the AMG signal, like the ultrasound-measured deformation that is the pri-

mary focus of this thesis, is mechanically, rather than neurologically coupled to force output, it presents a second, promising signal with which to measure musculoskeletal dynamics. The fact that different properties of the vibration signal are thought to correspond to different aspects of fiber recruitment could also enable more detailed study of motor unit recruitment strategies and how they vary across muscles and individuals.

Beyond the scope of this thesis, we are currently working to improve on this preliminary model — including separate correlation analysis of amplitude and frequency with force, fitting parameters α to enable actual force inference, and investigating both temporal and spatial resolution of the AMG signal (toward determining the shortest interval over which ν can be robustly calculated — which determines the "real-time" potential of this control signal formulation — and the number of independent signals that can be extracted from neighboring muscles, respectively) — and to test its validity on a wider cohort of subjects. In the future, as with our muscle deformation work, we aim to expand the resultant models to complex multi-muscle systems and employ them in device control and other real-time force inference applications.

Chapter 13

Conclusions & Future Directions

In this dissertation, we presented preliminary work prototyping the use of muscle deformation for force-related device control, as well several novel data sets and a comprehensive modeling framework to enable expanded study of the force–deformation relationship. In this final chapter, revisiting many of the application areas presented as motivation in chapter 2, we discuss both immediate and long-term plans for the application and expansion of this work in device control, human dynamics modeling, human–robot collaboration, and rehabilitation domains. We also present some final thoughts on the future importance of enhanced human modeling in general. On the whole, this thesis poses more questions than it answers, and the expansions below range from concrete to highly speculative.

13.1 Safe & Capable Device Control

As discussed in section 2.1, we envision two main respects in which our deformation-based muscle force modeling can enable enhanced device control: improving *safety* and improving *capability*, particularly for devices in which we wish to confer a high degree of control authority to the human user (e.g., prostheses, exoskeletons, teleoperation) and for automated systems that need to deeply understand human dynamics (e.g., to predict human motion with or without robotic intervention to avoid collision). In the subsections below, we provide several examples of specific control scenarios in which we envision these deformation-modeling-based safety and/or capability enhancements.

Simultaneous Position, Force, & Stiffness Control

One way in which deformation data could be utilized is as an additional control signal atop existing biosignal-based control schemes, enabling improved control in cases requiring precise impedance modulation (e.g., manipulating cane-like supports, using a screwdriver, performing catching maneuvers). Fig. 13.1 illustrates a simple teleoperation example of this usage, in which deformation, activation, and position signals are combined to enable

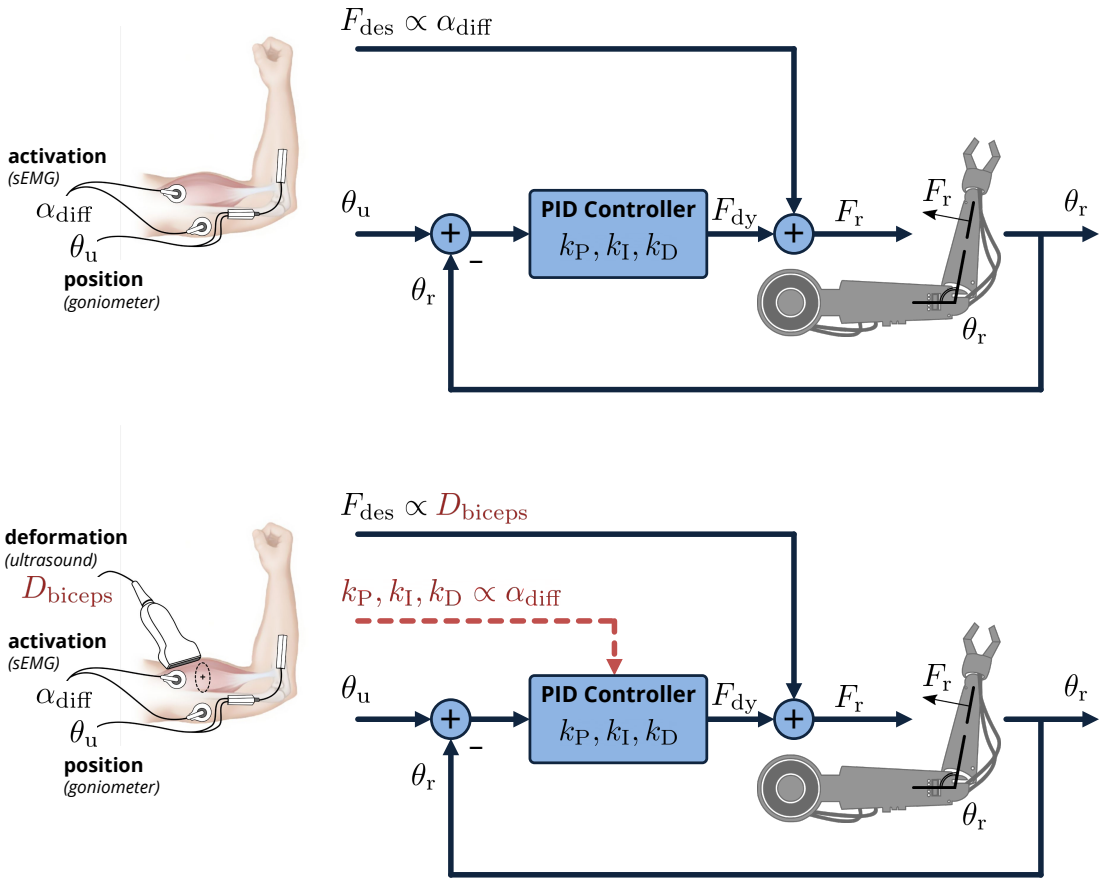

Figure 13.1: Idealized example of using deformation-based muscle force inference to enable simultaneous position, force, and stiffness control of a single joint in a teleoperation scenario. *Top*: Classical force-based control scenario in which human user's electrogoniometer-measured elbow angle θ_u is mapped to robot elbow angle θ_r through PID control loop with constants k_P, k_I, and k_D, with goal force trajectory F_{des} specified proportional to differential biceps–triceps surface electromyography (sEMG) signal α_{diff} and added to feedback-loop-generated position-correcting force F_{dy} to generate robot elbow force F_r. *Bottom*: By introducing ultrasound-measured muscle deformation D_{biceps} — a better analog for the user's absolute output force than the differential sEMG signal — we can configure the robot to track this output force more directly by specifying F_{des} as proportional to D_{biceps} and remap α_{diff} to generate controller stiffness constants consistent with the user's own elbow stiffness. This enables enhanced control in scenarios requiring precise independent modulation of position, force, and stiffness parameters (e.g., manipulation of cane-like weight-bearing tools, catching delicate objects). Note that this prototype example assumes ideal controllable-torque actuators; implementation would require careful actuator selection and likely innovation. Select images adapted from [67].

simultaneous and intuitive force, stiffness, and position modulation on a single robot joint, a paradigm shown to be feasible [1, 120] but that remains technically challenging.

High-Dimensional Device Control

A second potential deformation-based control application is our motivating example in chapter 1: intuitive control of devices with many degrees of freedom (DoF). While there exist a number of multi-jointed assistive robots and prostheses (e.g., anthropomorphic hands), with mechanical potential for substantial dexterity, most support only single-DoF control (e.g., opening and closing a grasp). A primary reason for these limited control schemes is humans' inability to consciously modulate the number of control signals required for greater dexterity: a fully biomimetic hand, for example, would require independent actuation of approximately 27 DoF, while humans often struggle to manipulate even a few disjoint control signals.[1]

We envision three possible strategies for addressing these challenges to enable high-DoF, dexterous control: substantial (likely prohibitive) levels of user training on classical control interfaces (e.g., joysticks), the introduction of considerable device autonomy, and the organic extraction of signals from the neuromuscular system humans are pre-wired to control. While we hypothesize that successful approaches will leverage all of these strategies, as well as additional augmentations like haptic feedback, we see the latter strategy as particularly powerful for increasing dexterity without diminishing human control authority.

The work in this thesis is directly applicable to precisely this multi-signal extraction problem, and is complementary to existing biosignal-based control schemes, which mostly employ surface electromyography (sEMG). Unlike sEMG, which provides an aggregate, noisy signal that varies greatly with even small shifts in sensor placement and largely cannot measure deep muscle activation, ultrasound-measured deformation is easily localized to a particular muscle and can be measured for deep and shallow muscles alike, allowing for straightforward extraction of multiple independent control signals to actuate multiple DoF. Ultimately, such signals could be used to control not only the biomimetic exoskeletons and prostheses discussed above, but other high-dimensional robotic systems that are currently difficult for a single operator to manage (e.g., fleets of quadrotors, supernumerary arms).

Human–Robot Collaboration

Beyond human-controlled assistive devices, deformation-based human dynamics models could improve physical human–robot interaction systems more broadly in two distinct ways: first, by enhanced modeling of human *capabilities*, and second, enhanced measurement of human *actions and intentions* in real time.

In the former case, such modeling systems could improve the safety and comfort of collaborative human–robot systems by monitoring human exertions and adjusting robot actions when these exertions surpass thresholds of comfort or safety (e.g., inducing a robot arm to

[1] This is something to which anyone who has played QWOP can attest.

take more weight when performing a collaborative lifting task). Scientifically, these monitoring systems could also be used to evaluate the quality of ergonomic cost functions currently used by these collaborative frameworks, evaluating whether the predicted ergonomicity of a system state in consistent with underlying musculoskeletal forces.

In the latter case, deformation-based measures of human dynamics could provide full dynamics trajectories of human actions, which could then be used directly in robot control schemes (e.g., inducing a robot to help lift an object simultaneously with a human collaborator) or to learn new robot control policies (e.g., monitoring human dynamics during such a lifting task to teach the robot via reinforcement learning which actions helpfully reduce required human exertions). At an even higher level of abstraction, these deformation-derived dynamics signals could also be used to evaluate the quality of the cognitive models often used to predict human intention in these domains by providing a concrete measure of human actions.

13.2 Neuromuscular Dynamics Modeling

In addition to device control, we highlight three key areas in the human neuromusculoskeletal modeling domain in which we see deformation-based muscle force modeling as particularly and immediately impactful.

Individual Muscle Force Inference

While our analyses in this dissertation are restricted to relating deformation to net joint output force, the true power in the deformation signal is its potential use in measuring individual muscle forces. As discussed throughout this thesis (and particularly in sections 2.3 and 2.4), there exists no framework for noninvasive, in vivo muscle force measurement of both deep and shallow muscles, a core limitation of biomechanics literature, and this thesis provides substantial evidence that deformation-based force inference could aid in filling this modeling gap.

To truly validate a force–deformation model in absolute terms would require an invasive (likely animal) study that enables muscle–tendon unit isolation and direct transducer-based strain measurement [6, 122]. In the future, we aim to investigate both these invasive validation strategies and others that don't require surgery (e.g., the "tendon tapping" method [99] discussed in section 2.4, a new technology shown to enable force measurement of muscle–tendon units for which the tendon is exposed).

Multi-Muscle Modeling & Existing Framework Validation

Beyond single-muscle force inference, (localizable) deformation provides a path toward improved modeling of synergistic and antagonistic muscle behaviors. In particular, we see deformation modeling as directly applicable in two ways, as shown in Fig. 13.2: first, by

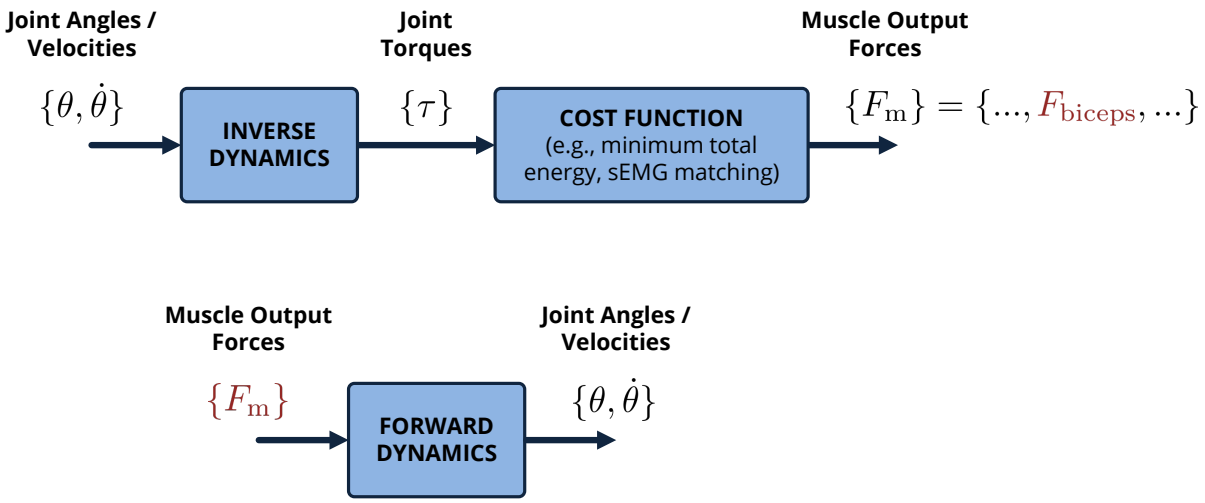

Figure 13.2: Example usage of deformation-based force inference in inverse and forward musculoskeletal dynamics modeling. *Top*: Current musculoskeletal simulation frameworks (e.g., OpenSim [22], AnyBody [19]) convert joint angles and velocities $\{\theta, \dot{\theta}\}$ to joint torques $\{\tau\}$ via inverse dynamics, then convert these torques to muscle forces $\{F_{\mathrm{m}}\}$ by optimizing over cost functions (e.g., minimum metabolic energy expenditure, gait cycle or sEMG matching) that may or may not represent true human dynamics. By measuring even a subset of muscle forces directly (e.g. F_{biceps} at the biceps brachii), we can begin to evaluate the validity of these cost functions, determine under what conditions they hold, and develop new ones. *Bottom*: Direct, simultaneous, deformation-based measurement of multiple muscle forces $\{F_{\mathrm{m}}\}$ would enable the creation of forward dynamics models of the human musculoskeletal system, a class of models that is not yet widely used in the biomechanics community, as these individual muscle forces are not often otherwise available even during invasive study.

providing force information on individual muscles that can be used to validate predictions made by existing (inverse) musculoskeletal modeling frameworks [22, 19], and second, by enabling the creation of novel forward models that predict human dynamics directly from muscle forces.

Force–Activation Modeling

While we have primarily focused in this dissertation on the (purely mechanical) force–deformation relationship, abstracting away the neurological control loop, the data and methodologies we have developed provide a compelling new approach to study of the nervous system. As illustrated in Fig. 13.3, a mechanical measure of output force constitutes a "loop closure" of sorts in our ability to measure the resulting physical output of neural control, not only the neurological input (already accessible peripherally via surface electromyography, or

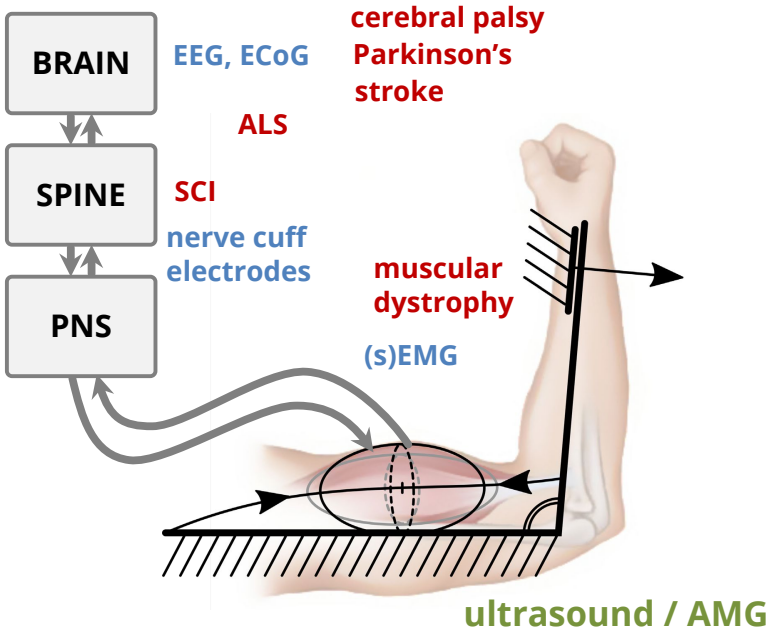

Figure 13.3: By providing mechanical, output-force-related deformation and vibration signals, both ultrasound and acoustic myography (AMG), respectively, "close" the neuromuscular control loop in terms of available sensor measurements, enabling enhanced understanding of neurological signals measured (via nerve cuff, EEG, ECoG, etc., *blue*) throughout the brain, spine, and peripheral nervous system (PNS), as well as a more holistic view of the way in which pathologies (*red*) throughout the system impact physical capabilities and behaviors. Image adapted from [67].

sEMG, or via nerve cuff electrodes, electroencephalography, or electrocorticography more centrally). Most immediately, as touched on in section 2.1, we are working to use these deformation signals to make sEMG more interpretable, by studying delays between activation and actual force onset and by inferring from which muscles (or even motor units) sEMG signals spatially originate.

13.3 Quantifying, Tracking, & Rehabilitating Pathology

The "loop closure" discussed above enables enhanced study of not only healthy individuals, but also those exhibiting musculoskeletal pathology. As illustrated in Fig. 13.3, pathologies influencing motion capabilities can occur at many different points along the neuromuscular control loop, and while the cause of these injuries is often fundamentally neurological

(e.g., spinal cord injury, abnormal brain development causing cerebral palsy), they generally result in physical morphological changes (e.g., muscle atrophy, permanent shortening of the muscle–tendon unit). Deformation-based force inference would allow for a more holistic understanding of these pathologies and would aid in quantifying the relationships between the causes of pathologies and their downstream effects, including variation over time (e.g., during the degeneration associated with muscular dystrophy to track disease progression, or during stroke recovery to track treatment efficacy).

Beyond providing these measures of efficacy, our proposed real-time force inference methods have direct applications in rehabilitation robotics. One immediate use for measures is to evaluate whether a given robotic device is actually inducing the desired rehabilitative exertions (e.g., whether a robotic exoskeleton causes target muscles to be activated or is simply performing the motion for the patient) — or, conversely, whether the device is inducing undesirable or pathological exertions that could be injurious. Another, longer-term application would be integrating these deformation-based force measures directly into the robotic control loop to enable the system to adapt in a bespoke manner to an individual's capabilities and rehabilitation goals.

13.4 Additional Fields of Interest

While we have focused our attention in this thesis on sensing and modeling schemes conducive to real-time control and modeling, the frameworks we developed — and in particular, the open-source data we have generated and released, as detailed in chapter 11 — could be used and adapted for many other applications and modeling paradigms. Our 3D OpenArm 1.0 and 2.0 data sets, for example, could be used to inspire and validate deformation models at many levels of detail (from finite-element-based to linear best-fit approximations), enabling advances in fields as diverse as muscle histology (through improved understanding of viscoelastic tissue mechanics) and animation (enabling enhanced graphical rendering of human motion). Likewise, the tracking and annotation software we developed could be applied to a wide variety of additional biomechanical analyses (e.g., by leveraging our optical flow tracking system to measure real-time motion of other tissues of interest, or training our neural-network-based autosegmentation framework to annotate additional structures).

13.5 Final Thoughts

Understanding detailed human dynamics will only become more critical as humans and robots increasingly occupy the same physical space and engage in progressively more complex physical interactions. We hope the work in this thesis will form the blueprint for novel deformation-based musculoskeletal modeling paradigms, enabling not only safe and capable human–robot systems, but enhanced understanding of neurological motor control and diagnosis and treatment of pathology.

Bibliography

[1] Cary J. Abul-Haj and Neville Hogan. "Functional Assessment of Control Systems for Cybernetic Elbow Prostheses — Part I: Description of the Technique". In: *IEEE Transactions on Biomedical Engineering* 37.11 (1990), pp. 1025–1036. DOI: 10.1109/TBME.1990.1438510.

[2] Nima Akhlaghi, Clayton A. Baker, Mohamed Lahlou, Hozaifah Zafar, Karthik G. Murthy, Huzefa S. Rangwala, Jana Kosecka, Wilsaan M. Joiner, Joseph J. Pancrazio, and Siddhartha Sikdar. "Real-time Classification of Hand Motions Using Ultrasound Imaging of Forearm Muscles". In: *IEEE Transactions on Biomedical Engineering* 63.8 (2016), pp. 1687–1698. DOI: 10.1109/TBME.2015.2498124.

[3] Md Zahangir Alom, Mahmudul Hasan, Chris Yakopcic, Tarek M. Taha, and Vijayan K. Asari. *Recurrent Residual Convolutional Neural Network based on U-Net (R2U-Net) for Medical Image Segmentation*. 2018. arXiv: 1802.06955 [cs.CV].

[4] Elsa D. Angelini and Olivier Gerard. "Review of Myocardial Motion Estimation Methods from Optical Flow Tracking on Ultrasound Data". In: *International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*. 2006, pp. 1537–1540. DOI: 10.1109/IEMBS.2006.259640.

[5] Lee Barber, Rod Barrett, and Glen Lichtwark. "Validity and reliability of a simple ultrasound approach to measure medial gastrocnemius muscle length". In: *Journal of Anatomy* 218.6 (June 2011), pp. 637–642. ISSN: 00218782. DOI: 10.1111/j.1469-7580.2011.01365.x.

[6] G. R. G. Barnes and D. N. Pinder. "In vivo tendon tension and bone strain measurement and correlation". In: *Journal of Biomechanics* 7.1 (1974), pp. 35–42. DOI: 10.1016/0021-9290(74)90068-2.

[7] Daniel T. Barry, Steven R. Geiringer, and Richard D. Ball. "Acoustic myography: a noninvasive monitor of motor unit fatigue". In: *Muscle & Nerve* 8.3 (1985), pp. 189–194. DOI: 10.1002/mus.880080303.

[8] Daniel T. Barry, Timothy Hill, and Dukjin Im. "Muscle fatigue measured with evoked muscle vibrations". In: *Muscle & Nerve* 15.3 (1992), pp. 303–309. DOI: 10.1002/mus.880150308.

[9] Zachary Baum, Tamas Ungi, Andras Lasso, Ben Church, Christopher Schlenger, and Gabor Fichtinger. "Visual aid for identifying vertebral landmarks in ultrasound". In: *Medical Imaging 2018: Image-Guided Procedures, Robotic Interventions, and Modeling*. Vol. 10576. International Society for Optics and Photonics. 2018. DOI: 10.1117/12.2292820.

[10] Muyinatu A. Lediju Bell, H. Tutkun Sen, Iulian Iordachita, and Peter Kazanzides. "Force-controlled ultrasound robot for consistent tissue pre-loading: Implications for acoustic radiation force elasticity imaging". In: *IEEE RAS/EMBS International Conference on Biomedical Robotics and Biomechatronics (BioRob)*. IEEE. 2014, pp. 259–264. DOI: 10.1109/BIOROB.2014.6913786.

[11] Aaron Bestick, Ravi Pandya, Ruzena Bajcsy, and Anca D. Dragan. "Learning Human Ergonomic Preferences for Handovers". In: *IEEE International Conference on Robotics and Automation (ICRA)*. IEEE. 2018, pp. 3257–3264. DOI: 10.1109/ICRA.2018.8461216.

[12] Mustapha Bouhrara, Benoist Lehallier, Sylvie Clerjon, Jean-Louis Damez, and Jean-Marie Bonny. "Mapping of muscle deformation during heating: in situ dynamic MRI and nonlinear registration". In: *Magnetic Resonance Imaging* 30.3 (2012), pp. 422–430. DOI: 10.1016/j.mri.2011.10.002.

[13] George E. P. Box. "Robustness in the Strategy of Scientific Model Building". In: *Robustness in Statistics*. Elsevier, 1979, pp. 201–236. DOI: 10.1016/B978-0-12-438150-6.50018-2.

[14] Gary Bradski. "The OpenCV Library". In: *Dr. Dobb's Journal of Software Tools* 25 (2000), pp. 120–125.

[15] F. Buchthal and H. Schmalbruch. "Motor Unit of Mammalian Muscle". In: *Physiological reviews* 60.1 (1980), pp. 90–142. DOI: 10.1152/physrev.1980.60.1.90.

[16] Xin Chen, Yong Ping Zheng, Jing Yi Guo, and Jun Shi. "Sonomyography (SMG) Control for Powered Prosthetic Hand: A Study with Normal Subjects". In: *Ultrasound in Medicine and Biology* 36.7 (2010), pp. 1076–1088. DOI: 10.1016/j.ultrasmedbio.2010.04.015.

[17] Özgün Çiçek, Ahmed Abdulkadir, Soeren S. Lienkamp, Thomas Brox, and Olaf Ronneberger. "3D U-Net: Learning Dense Volumetric Segmentation from Sparse Annotation". In: *International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI)*. Springer. 2016, pp. 424–432. DOI: 10.1007/978-3-319-46723-8_49.

[18] Mathilde Connan, Eduardo Ruiz Ramírez, Bernhard Vodermayer, and Claudio Castellini. "Assessment of a Wearable Force- and Electromyography Device and Comparison of the Related Signals for Myocontrol". In: *Frontiers in Neurorobotics* 10 (2016), p. 17. DOI: 10.3389/fnbot.2016.00017.

[19] 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 (2006), pp. 1100–1111. DOI: 10.1016/j.simpat.2006.09.001.

[20] Mohammad Darainy and David J. Ostry. "Muscle cocontraction following dynamics learning". In: *Experimental Brain Research* 190.2 (2008), pp. 153–163. DOI: 10.1007/s00221-008-1457-y.

[21] Carlo J. De Luca. "The Use of Surface Electromyography in Biomechanics". In: *Journal of Applied Biomechanics* 13.2 (1997), pp. 135–163. DOI: 10.1123/jab.13.2.135.

[22] Scott L. Delp, Frank C. Anderson, Allison S. Arnold, Peter Loan, Ayman Habib, Chand T. John, Eran Guendelman, and Darryl G. Thelen. "OpenSim: Open-Source Software to Create and Analyze Dynamic Simulations of Movement". In: *IEEE Transactions on Biomedical Engineering* 54.11 (2007), pp. 1940–1950. DOI: 10.1109/TBME.2007.901024.

[23] A. S. Deshmukh, D. E. Steenberg, M. Hostrup, J. B. Birk, J. K. Larsen, A. Santos, R. Kjøbsted, J. R. Hingst, C. C. Schéele, M. Murgia, B. Kiens, E. A. Richter, M. Mann, and J. F. P. Wojtaszewski. "Deep muscle-proteomic analysis of freeze-dried human muscle biopsies reveals fiber type-specific adaptations to exercise training". In: *Nature Communications* 12.1 (2021), pp. 1–15. DOI: 10.1038/s41467-020-20556-8.

[24] Ananya S. Dhawan, Biswarup Mukherjee, Shriniwas Patwardhan, Nima Akhlaghi, Guoqing Diao, Gyorgy Levay, Rahsaan Holley, Wilsaan M. Joiner, Michelle Harris-Love, and Siddhartha Sikdar. "Proprioceptive Sonomyographic Control: A novel method for intuitive and proportional control of multiple degrees-of-freedom for individuals with upper extremity limb loss". In: *Scientific Reports* 9.1 (2019), pp. 1–15. DOI: 10.1038/s41598-019-45459-7.

[25] James Dolbow and Zachary Throckmorton. "Neuroanatomy, Spinal Cord Myotatic Reflex". In: *StatPearls* (2020). URL: https://www.ncbi.nlm.nih.gov/books/NBK551629/.

[26] Anahid Ebrahimi, Robyn L. Kuchler, Robin L. Pomeroy, Isaac F. Loegering, Jack A. Martin, and Darryl G. Thelen. "Normative Achilles and patellar tendon shear wave speeds and loading patterns during walking in typically developing children". In: *Gait & Posture* (2021). DOI: 10.1016/j.gaitpost.2021.05.023.

[27] G. J. C. Ettema, G. Styles, and V. Kippers. "The moment arms of 23 muscle segments of the upper limb with varying elbow and forearm positions: implications for motor control". In: *Human Movement Science* 17.2 (1998), pp. 201–220.

[28] Ting-Yun Fang, Haichong K. Zhang, Rodolfo Finocchi, Russell H. Taylor, and Emad M. Boctor. "Force-assisted ultrasound imaging system through dual force sensing and admittance robot control". In: *International Journal of Computer Assisted Radiology and Surgery* 12.6 (2017), pp. 983–991.

[29] Dominic James Farris, Benjamin D. Robertson, and Gregory S. Sawicki. "Elastic ankle exoskeletons reduce soleus muscle force but not work in human hopping". In: *Journal of Applied Physiology* 115.5 (2013), pp. 579–585.

[30] Dominic James Farris and Gregory S. Sawicki. "Human medial gastrocnemius force-velocity behavior shifts with locomotion speed and gait". In: *Proceedings of the National Academy of Sciences (PNAS)* 109.3 (Jan. 2012), pp. 977–982. ISSN: 00278424. DOI: 10.1073/pnas.1107972109.

[31] John A. Faulkner. "Terminology for contractions of muscles during shortening, while isometric, and during lengthening". In: *Journal of Applied Physiology* 95.2 (2003), pp. 455–459.

[32] Andriy Fedorov, Reinhard Beichel, Jayashree Kalpathy-Cramer, Julien Finet, Jean-Christophe Fillion-Robin, Sonia Pujol, Christian Bauer, Dominique Jennings, Fiona Fennessy, Milan Sonka, John Buatti, Stephen Aylward, James V. Miller, Steve Pieper, and Ron Kikinis. "3D Slicer as an Image Computing platform for the Quantitative Imaging Network". In: *Magnetic Resonance Imaging* 30 (9 Nov. 2012), pp. 1323–1341. DOI: 10.1016/j.mri.2012.05.001.

[33] Angelo Emanuele Fiorilla, Francesco Nori, Lorenzo Masia, and Giulio Sandini. "Finger impedance evaluation by means of hand exoskeleton". In: *Annals of Biomedical Engineering* 39.12 (2011), p. 2945.

[34] Tamar Flash and Neville Hogan. "The coordination of arm movements: an experimentally confirmed mathematical model". In: *Journal of Neuroscience* 5.7 (1985), pp. 1688–1703. DOI: 10.1523/JNEUROSCI.05-07-01688.1985.

[35] Julaine M. Florence, Shree Pandya, Wendy M. King, Jenny Diveley Robison, Linda C. Signore, Mindy Wentzell, and Michael A. Province. "Clinical Trials in Duchenne Dystrophy: Standardization and Reliability of Evaluation Procedures". In: *Physical Therapy* 64.1 (Jan. 1984), pp. 41–45. DOI: 10.1093/ptj/64.1.41.

[36] J.V. Frangioni, T.S. Kwan-Gett, L.E. Dobrunz, and T.A. McMahon. "The mechanism of low-frequency sound production in muscle". In: *Biophysical Journal* 51.5 (1987), pp. 775–783. DOI: 10.1016/S0006-3495(87)83404-5.

[37] Laura A. Frey-Law, Andrea Laake, Keith G. Avin, Jesse Heitsman, Tim Marler, and Karim Abdel-Malek. "Knee and elbow 3D strength surfaces: peak torque-angle-velocity relationships". In: *Journal of Applied Biomechanics* 28.6 (2012), pp. 726–737.

[38] Nir Friedman and Stuart Russell. "Image segmentation in video sequences: A probabilistic approach". In: *Proceedings of the Thirteenth Conference on Uncertainty in Artificial Intelligence*. Morgan Kaufmann Publishers Inc. 1997, pp. 175–181.

[39] Shanette A. Go, William J. Litchy, Loribeth Q. Evertz, and Kenton R. Kaufman. "Evaluating skeletal muscle electromechanical delay with intramuscular pressure". In: *Journal of Biomechanics* 76 (2018), pp. 181–188. DOI: 10.1016/j.jbiomech.2018.05.029.

[40] A. M. Gordon, Andrew F. Huxley, and F. J. Julian. "The variation in isometric tension with sarcomere length in vertebrate muscle fibres". In: *Journal of Physiology* 184.1 (1966), pp. 170–192.

[41] Helmut Grabner, Jiri Matas, Luc Van Gool, and Philippe Cattin. "Tracking the invisible: Learning where the object might be". In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. 2010, pp. 1285–1292. DOI: 10.1109/CVPR.2010.5539819.

[42] Henry Gray. *Anatomy of the Human Body*. Philadelphia: Lea & Febiger, 1918. ISBN: 0-691-08322-3. URL: https://www.bartleby.com/107/.

[43] F. M. Grimaldi. *Physico-mathesis de lumine, coloribus et iride aliisque adnexis*. 1665.

[44] J. Gulati and A. Babu. "Tonicity effects on intact single muscle fibers: relation between force and cell volume". In: *Science* 215.4536 (1982), pp. 1109–1112.

[45] Jing Yi Guo, Yong Ping Zheng, Hong Bo Xie, and Xin Chen. "Continuous monitoring of electromyography (EMG), mechanomyography (MMG), sonomyography (SMG) and torque output during ramp and step isometric contractions". In: *Medical Engineering and Physics* 32.9 (Nov. 2010), pp. 1032–1042. DOI: 10.1016/j.medengphy.2010.07.004.

[46] Mark Halaki, Nicholas O'Dwyer, and Ian Cathers. "Systematic nonlinear relations between displacement amplitude and joint mechanics at the human wrist". In: *Journal of Biomechanics* 39.12 (2006), pp. 2171–2182.

[47] Laura A Hallock, Robert Peter Matthew, Sarah Seko, and Ruzena Bajcsy. *Sensor-Driven Musculoskeletal Dynamic Modeling*. Tech. rep. University of California, Berkeley, 2016, UCB/EECS–2016–66. URL: http://www.eecs.berkeley.edu/Pubs/TechRpts/2016/EECS-2016-66.html.

[48] Laura A. Hallock and Ruzena Bajcsy. "A Preliminary Evaluation of Acoustic Myography for Real-Time Muscle Force Inference". In: *International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*. IEEE. 2018.

[49] Laura A. Hallock, Akira Kato, and Ruzena Bajcsy. "Empirical Quantification and Modeling of Muscle Deformation: Toward Ultrasound-Driven Assistive Device Control". In: *IEEE International Conference on Robotics and Automation (ICRA)*. IEEE. 2018, pp. 1825–1832. DOI: 10.1109/ICRA.2018.8462887.

[50] Laura A. Hallock, Bhavna Sud, Chris Mitchell, Eric Hu, Fayyaz Ahamed, Akash Velu, Amanda Schwartz, and Ruzena Bajcsy. "Toward Real-Time Muscle Force Inference and Device Control via Optical-Flow-Tracked Muscle Deformation". In: *IEEE Transactions on Neural Systems and Rehabilitation Engineering (TNSRE)* (2021, under review).

[51] Laura A. Hallock, Akash Velu, Amanda Schwartz, and Ruzena Bajcsy. "Muscle deformation correlates with output force during isometric contraction". In: *IEEE RAS/EMBS International Conference on Biomedical Robotics and Biomechatronics (BioRob)*. IEEE. 2020, pp. 1188–1195. DOI: 10.1109/BioRob49111.2020.9224391.

[52] Mo Han, Sezen Yağmur Günay, Gunar Schirner, Taşkın Padır, and Deniz Erdoğmuş. "HANDS: a multimodal dataset for modeling toward human grasp intent inference in prosthetic hands". In: *Intelligent Service Robotics* 13.1 (2020), pp. 179–185.

[53] Adrian P Harrison. "A more precise, repeatable and diagnostic alternative to surface electromyography–an appraisal of the clinical utility of acoustic myography". In: *Clinical Physiology and Functional Imaging* 38.2 (2018), pp. 312–325.

[54] Rudolf Hellmuth. *Hill Muscle Model*. 2010. URL: https://commons.wikimedia.org/wiki/File:Hill_muscle_model.svg (visited on 08/11/2021).

[55] Elwood Henneman, George Somjen, and David O. Carpenter. "Excitability and inhibitibility of motoneurons of different sizes". In: *Journal of Neurophysiology* 28.3 (1965), pp. 599–620.

[56] Walter Herzog. "Skeletal muscle mechanics: questions, problems and possible solutions". In: *Journal of NeuroEngineering and Rehabilitation* 14.1 (Dec. 2017), p. 98. DOI: 10.1186/s12984-017-0310-6.

[57] Walter Herzog, Krysta Powers, Kaleena Johnston, and Mike Duvall. *A new paradigm for muscle contraction*. June 2015. DOI: 10.3389/fphys.2015.00174.

[58] Nalinda Hettiarachchi, Zhaojie Ju, and Honghai Liu. "A New Wearable Ultrasound Muscle Activity Sensing System for Dexterous Prosthetic Control". In: *IEEE International Conference on Systems, Man, and Cybernetics (SMC)*. IEEE, Jan. 2016, pp. 1415–1420. DOI: 10.1109/SMC.2015.251.

[59] A. V. Hill. "The heat of shortening and the dynamic constants of muscle". In: *Proceedings of the Royal Society of London B: Biological Sciences* 126.843 (1938), pp. 136–195.

[60] Marilyn Hinson and Joel Rosentswieg. "Comparative electromyographic values of isometric, isotonic, and isokinetic contraction". In: *Research Quarterly, American Association for Health, Physical Education and Recreation* 44.1 (1973), pp. 71–78.

[61] Hoa X. Hoang, Laura E. Diamond, David G. Lloyd, and Claudio Pizzolato. "A calibrated EMG-informed neuromusculoskeletal model can appropriately account for muscle co-contraction in the estimation of hip joint contact forces in people with hip osteoarthritis". In: *Journal of Biomechanics* 83 (2019), pp. 134–142.

[62] P. W. Hodges, L. H. M. Pengel, R. D. Herbert, and S. C. Gandevia. "Measurement of muscle contraction with ultrasound imaging". In: *Muscle & Nerve* 27.6 (2003), pp. 682–692.

[63] John A. Hodgson, Sheng-Wei Chi, Judy P. Yang, Jiun-Shyan Chen, Victor R. Edgerton, and Shantanu Sinha. "Finite element modeling of passive material influence on the deformation and force output of skeletal muscle". In: *Journal of the Mechanical Behavior of Biomedical Materials* 9 (2012), pp. 163–183.

[64] Katherine R. S. Holzbaur, Wendy M. Murray, Garry E. Gold, and Scott L. Delp. "Upper limb muscle volumes in adult subjects". In: *Journal of Biomechanics* 40.4 (2007), pp. 742–749.

[65] Andrew F. Huxley and Rolf Niedergerke. "Structural changes in muscle during contraction: interference microscopy of living muscle fibres". In: *Nature* 173.4412 (1954), pp. 971–973.

[66] Hugh Huxley and Jean Hanson. "Changes in the cross-striations of muscle during contraction and stretch and their structural interpretation". In: *Nature* 173.4412 (1954), pp. 973–976.

[67] *Innovator X Post Op Elbow Brace Support – Ossur*. URL: https://www.orthotape.com/innovator-x-elbow-brace.html (visited on 08/11/2021).

[68] Siddarth Jain and Brenna Argall. "Probabilistic human intent recognition for shared autonomy in assistive robotics". In: *ACM Transactions on Human-Robot Interaction (THRI)* 9.1 (2019), pp. 1–23.

[69] Marie-Ange Janvier, Samir Merouche, Louise Allard, Gilles Soulez, and Guy Cloutier. "A 3-D ultrasound imaging robotic system to detect and quantify lower limb arterial stenoses: in vivo feasibility". In: *Ultrasound in Medicine & Biology* 40.1 (2014), pp. 232–243.

[70] Noémie Jaquier, Mathilde Connan, Claudio Castellini, and Sylvain Calinon. "Combining Electromyography and Tactile Myography to Improve Hand and Wrist Activity Detection in Prostheses". In: *Technologies* 5.4 (2017), p. 64.

[71] F. Jørgensen and O. Lammert. "Accelerometermyography (AMG)". In: *Biomechanics V: Proceedings of the Fifth International Congress of Biomechanics, Jyväskylä, Finland*. Vol. 1. University Park Press. 1976, p. 159.

[72] B. Juul-Kristensen, F. Bojsen-Møller, E. Holst, and C. Ekdahl. "Comparison of muscle sizes and moment arms of two rotator cuff muscles measured by Ultrasonography and Magnetic Resonance Imaging". In: *European Journal of Ultrasound* 11.3 (June 2000), pp. 161–173. DOI: 10.1016/S0929-8266(00)00084-7.

[73] Konstantinos Kamnitsas, Christian Ledig, Virginia F. J. Newcombe, Joanna P. Simpson, Andrew D. Kane, David K. Menon, Daniel Rueckert, and Ben Glocker. "Efficient multi-scale 3D CNN with fully connected CRF for accurate brain lesion segmentation". In: *Medical Image Analysis* 36 (2017), pp. 61–78.

[74] Daisuke Kaneishi, Robert Peter Matthew, and Masayoshi Tomizuka. "Optimal Control Parameterization for Active/Passive EXoskeleton with Variable Impedance Actuator". In: *IEEE RAS/EMBS International Conference on Biomedical Robotics and Biomechatronics (BioRob)*. IEEE, Oct. 2018, pp. 713–719. DOI: 10.1109/BIOROB.2018.8487719.

[75] Matthias Keidel and Wolf-Dieter Keidel. "The Computer-Vibromyography as a Biometric Progress in Studying Muscle Function – Die Computer-Vibromyographie - Ein biometrischer Ansatz zur Messung mechanischer Muskelaktivität". In: *Biomedizinische Technik* (1989).

[76] Kyle B. Kiesel, Tim L. Uhl, Frank B. Underwood, Donald W. Rodd, and Arthur J. Nitz. "Measurement of lumbar multifidus muscle contraction with rehabilitative ultrasound imaging". In: *Manual Therapy* 12.2 (2007), pp. 161–166.

[77] Sung-Min Kim, Ju-Hwan Lee, Seung-Gyu Roh, and Sung-Yun Park. "The Study of Pre-Processing Algorithm for Improving Performance of Optical Flow Techniques on Ultrasound Image". In: *Journal of the Institute of Electronics Engineers of Korea SC* 47.5 (2016), pp. 24–32. DOI: 10.2316/p.2010.728-006.

[78] Yoon Sang Kim, Sooyong Lee, Changhyun Cho, Munsang Kim, and Chong-Won Lee. "A new exoskeleton-type masterarm with force reflection based on the torque sensor beam". In: *IEEE International Conference on Robotics and Automation (ICRA)*. Vol. 3. IEEE. 2001, pp. 2628–2633.

[79] Diederik P. Kingma and Jimmy Ba. *Adam: A method for stochastic optimization*. 2014. arXiv: 1412.6980.

[80] Joseph J. Knapik, James E. Wright, Roberta H. Mawdsley, and Joanne Braun. "Isometric, isotonic, and isokinetic torque variations in four muscle groups through a range of joint motion". In: *Physical Therapy* 63.6 (1983), pp. 938–947.

[81] Jan Wiebe H. Korstanje, Ruud W. Selles, Henk J. Stam, Steven E. R. Hovius, and Johan G. Bosch. "Development and validation of ultrasound speckle tracking to quantify tendon displacement". In: *Journal of Biomechanics* 43.7 (May 2010), pp. 1373–1379. DOI: 10.1016/j.jbiomech.2010.01.001.

[82] AI Kostyukov. "Muscle hysteresis and movement control: a theoretical study". In: *Neuroscience* 83.1 (1998), pp. 303–320.

[83] Lisa S. Krivickas, David J. Dorer, Julien Ochala, and Walter R. Frontera. "Relationship between force and size in human single muscle fibres". In: *Experimental Physiology* 96.5 (2011), pp. 539–547.

[84] Korupalli V. Rajesh Kumar and Susan Elias. "Smart Neck-Band for Rehabilitation of Musculoskeletal Disorders". In: *International Conference on COMmunication Systems & NETworkS (COMSNETS)*. IEEE. 2020, pp. 47–52.

[85] Prashant Kumar, Piyush Sharma, Harish Kumar Banga, Parveen Kalra, and Rajesh Kumar. "Performance Evaluation of Ankle Foot Orthosis on Lower Extremity Disabled Persons while Walking using OpenSim". In: *Operations Management and Systems Engineering*. Springer, 2021, pp. 299–310.

[86] O. Lammert, F. Jorgensen, and N. Einer-Jensen. *Accelerometermyography (AMG) I: method for measuring mechanical vibrations from isometrically contracted muscles*. 1976.

[87] Janna Brit Langholz, Gunnar Westman, and Magnus Karlsteen. "Musculoskeletal modelling in sports-evaluation of different software tools with focus on swimming". In: *Procedia Engineering* 147 (2016), pp. 281–287.

[88] Andras Lasso, Tamas Heffter, Adam Rankin, Csaba Pinter, Tamas Ungi, and Gabor Fichtinger. "PLUS: Open-source toolkit for ultrasound-guided intervention systems". In: *IEEE Transactions on Biomedical Engineering* 10 (Oct. 2014), pp. 2527–2537. DOI: 10.1109/TBME.2014.2322864.

[89] Hyunglae Lee, Elliott J. Rouse, and Hermano Igo Krebs. "Summary of human ankle mechanical impedance during walking". In: *IEEE Journal of Translational Engineering in Health and Medicine* 4 (2016), pp. 1–7.

[90] Robson Lemos, Marcelo Epstein, Walter Herzog, and Brian Wyvill. "Realistic skeletal muscle deformation using finite element analysis". In: *Proceedings of XIV Brazilian Symposium on Computer Graphics and Image Processing*. IEEE. 2001, pp. 192–199.

[91] Robson R. Lemos, Jon Rokne, Gladimir V. G. Baranoski, Yasuo Kawakami, and Toshiyuki Kurihara. "Modeling and simulating the deformation of human skeletal muscle based on anatomy and physiology". In: *Computer Animation and Virtual Worlds* 16.3-4 (2005), pp. 319–330.

[92] Geert Litjens, Thijs Kooi, Babak Ehteshami Bejnordi, Arnaud Arindra Adiyoso Setio, Francesco Ciompi, Mohsen Ghafoorian, Jeroen Awm Van Der Laak, Bram Van Ginneken, and Clara I Sánchez. "A survey on deep learning in medical image analysis". In: *Medical Image Analysis* 42 (2017), pp. 60–88.

[93] Chang Liu, Jessica B. Hamrick, Jaime F. Fisac, Anca D. Dragan, J. Karl Hedrick, S. Shankar Sastry, and Thomas L. Griffiths. *Goal inference improves objective and perceived performance in human-robot collaboration*. 2018. arXiv: 1802.01780.

[94] T. Liu, Y. Inoue, K. Shibata, Y. Hirota, and K. Shiojima. "A mobile force plate system and its application to quantitative evaluation of normal and pathological gait". In: *IEEE/ASME International Conference on Advanced Intelligent Mechatronics*. IEEE. 2010, pp. 272–277.

[95] Tao Liu, Yoshio Inoue, Kyoko Shibata, and Rencheng Zheng. "Wearable sensor system for human dynamics analysis". In: *Mechatronic Systems Applications* (2010), pp. 117–136.

[96] Bruce D. Lucas and Takeo Kanade. "An Iterative Image Registration Technique with an Application to Stereo Vision". In: *Proceedings of Imaging Understanding Workshop*. 1981, pp. 121–130.

[97] Constantinos N. Maganaris. "Force–length characteristics of in vivo human skeletal muscle". In: *Acta Physiologica Scandinavica* 172.4 (2001), pp. 279–285.

[98] Kasper Marstal, Floris Berendsen, Marius Staring, and Stefan Klein. "SimpleElastix: A user-friendly, multi-lingual library for medical image registration". In: *IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*. IEEE. 2016, pp. 574–582.

[99] Jack A. Martin, Scott C. E. Brandon, Emily M. Keuler, James R. Hermus, Alexander C. Ehlers, Daniel J. Segalman, Matthew S. Allen, and Darryl G. Thelen. "Gauging force by tapping tendons". In: *Nature Communications* 9.1 (Dec. 2018), p. 1592. DOI: 10.1038/s41467-018-03797-6.

[100] Martín Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S. Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Ian Goodfellow, Andrew Harp, Geoffrey Irving, Michael Isard, Yangqing Jia, Rafal Jozefowicz, Lukasz Kaiser, Manjunath Kudlur, Josh Levenberg, Dandelion Mané, Rajat Monga, Sherry Moore, Derek Murray, Chris Olah, Mike Schuster, Jonathon Shlens, Benoit Steiner, Ilya Sutskever, Kunal Talwar, Paul Tucker, Vincent Vanhoucke, Vijay Vasudevan, Fernanda Viégas, Oriol Vinyals, Pete Warden, Martin Wattenberg, Martin Wicke, Yuan Yu, and Xiaoqiang Zheng. *TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems*. Software available from tensorflow.org. 2015. URL: https://www.tensorflow.org/.

[101] Kim Mathiassen, Jørgen Enger Fjellin, Kyrre Glette, Per Kristian Hol, and Ole Jakob Elle. "An Ultrasound Robotic System Using the Commercial Robot UR5". In: *Frontiers in Robotics and AI* 3 (2016), p. 1. DOI: 10.3389/frobt.2016.00001.

[102] Robert Peter Matthew, Sarah Seko, and Ruzena Bajcsy. "Fusing motion-capture and inertial measurements for improved joint state recovery: An application for sit-to-stand actions". In: *International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*. IEEE. 2017, pp. 1893–1896.

[103] Thomas A. McMahon. *Muscles, Reflexes, and Locomotion*. Princeton University Press, 1984.

[104] J. M. McMeeken, I. D. Beith, D. J. Newham, P. Milligan, and D. J. Critchley. "The relationship between EMG and change in thickness of transversus abdominis". In: *Clinical Biomechanics* 19.4 (2004), pp. 337–342. DOI: 10.1016/j.clinbiomech.2004.01.007.

[105] Roberto Merletti, Alberto Botter, and U. Barone. "Detection and conditioning of surface EMG signals". In: *Surface Electromyography: Physiology, Engineering, and Applications*. Wiley Online Library, 2016, pp. 54–90.

[106] Andrzej P. Mierzwa, Sean P. Huang, Kristen T. Nguyen, Martin O. Culjat, and Rahul S. Singh. "Wearable ultrasound array for point-of-care imaging and patient monitoring". In: *Studies in Health Technology and Informatics*. Vol. 220. 2016, pp. 241–244. DOI: 10.3233/978-1-61499-625-5-241.

[107] Ivana Mikic, Slawomir Krucinski, and James D. Thomas. "Segmentation and tracking in echocardiographic sequences: Active contours guided by optical flow estimates". In: *IEEE Transactions on Medical Imaging* 17.2 (1998), pp. 274–284. DOI: 10.1109/42.700739.

[108] Laura C. Miller, Christopher K. Thompson, Francesco Negro, C. J. Heckman, Dario Farina, and Julius P. A. Dewald. "High-density surface EMG decomposition allows for recording of motor unit discharge from proximal and distal flexion synergy muscles simultaneously in individuals with stroke". In: *International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*. IEEE. 2014, pp. 5340–5344.

[109] Joseph Mizrahi. "Mechanical impedance and its relations to motor control, limb dynamics, and motion biomechanics". In: *Journal of Medical and Biological Engineering* 35.1 (2015), pp. 1–20.

[110] Pim Moeskops, Jelmer M. Wolterink, Bas H. M. van der Velden, Kenneth G. A. Gilhuijs, Tim Leiner, Max A. Viergever, and Ivana Išgum. "Deep learning for multi-task medical image segmentation in multiple modalities". In: *International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI)*. Springer. 2016, pp. 478–486.

[111] Ferdinando A. Mussa-Ivaldi, Neville Hogan, and Emilio Bizzi. "Neural, mechanical, and geometric factors subserving arm posture in humans". In: *Journal of Neuroscience* 5.10 (1985), pp. 2732–2743.

[112] Kiisa Nishikawa. "Eccentric contraction: unraveling mechanisms of force enhancement and energy conservation". In: *Journal of Experimental Biology* 219.2 (2016), pp. 189–196.

[113] Yonatan Nozik*, Laura A. Hallock*, Daniel Ho, Sai Mandava, Chris Mitchell, Thomas Hui Li, and Ruzena Bajcsy. "OpenArm 2.0: Automated Segmentation of 3D Tissue Structures for Multi-Subject Study of Muscle Deformation Dynamics". In: *International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*. IEEE. 2019, pp. 982–988. DOI: 10.1109/EMBC.2019.8857669. *Equal contribution.

[114] Claudio Orizio. "Muscle sound: Bases for the Introduction of a Mechanomyographic Signal in Muscle Studies". In: *Journal of Critical Reviews in Biomedical Engineering* 21.3 (1993), pp. 201–43.

[115] G. Oster and J. S. Jaffe. "Low frequency sounds from sustained contraction of human skeletal muscle". In: *Biophysical Journal* 30.1 (1980), pp. 119–127. DOI: 10.1016/S0006-3495(80)85080-6.

[116] Ece Ozkan, Christine Tanner, Matej Kastelic, Oliver Mattausch, Maxim Makhinya, and Orcun Goksel. "Robust motion tracking in liver from 2D ultrasound images using supporters". In: *International Journal of Computer Assisted Radiology and Surgery* 12.6 (June 2017), pp. 941–950. DOI: 10.1007/s11548-017-1559-8.

[117] Sang Il Park and Jessica K. Hodgins. "Data-driven modeling of skin and muscle deformation". In: *ACM Transactions on Graphics (TOG)*. Vol. 27. 3. ACM. 2008, p. 96.

[118] Elyse Passmore, Adrian Lai, Morgan Sangeux, Anthony G. Schache, and Marcus G. Pandy. "Application of ultrasound imaging to subject-specific modelling of the human musculoskeletal system". In: *Meccanica* 52.3 (2017), pp. 665–676.

[119] Guido G. Pena, Leonardo J. Consoni, Wilian M. dos Santos, and Adriano A. G. Siqueira. "Feasibility of an optimal EMG-driven adaptive impedance control applied to an active knee orthosis". In: *Robotics and Autonomous Systems* 112 (2019), pp. 98–108.

[120] Rita A. Popat, David E. Krebs, John Mansfield, Donald Russell, Edward Clancy, Kathleen M. Gill-Body, and Neville Hogan. "Quantitative assessment of four men using above-elbow prosthetic control". In: *Archives of Physical Medicine and Rehabilitation* 74.7 (1993), pp. 720–729.

[121] M. H. Pope, R. Crowninshield, R. Miller, and R. Johnson. "The static and dynamic behavior of the human knee in vivo". In: *Journal of Biomechanics* 9.7 (1976), pp. 449–452.

[122] Bérangère Ravary, Philippe Pourcelot, Claude Bortolussi, Stanislas Konieczka, and Nathalie Crevier-Denoix. "Strain and force transducers used in human and veterinary tendon and ligament biomechanical studies". In: *Clinical Biomechanics* 19.5 (2004), pp. 433–447. DOI: j.clinbiomech.2004.01.008.

[123] Marc Roig, Kelly O'Brien, Gregory Kirk, Ryan Murray, Patrick McKinnon, Babak Shadgan, and W. D. Reid. "The effects of eccentric versus concentric resistance training on muscle strength and mass in healthy adults: a systematic review with meta-analysis". In: *British Journal of Sports Medicine* 43.8 (2009), pp. 556–568.

[124] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. "U-Net: Convolutional networks for biomedical image segmentation". In: *International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI)*. Springer. 2015, pp. 234–241.

[125] Michael Rosenblum, Georgy Firsov, Robert Kuuz, and Bernd Pompe. "Human postural control: Force plate experiments and modelling". In: *Nonlinear Analysis of Physiological Data*. Springer, 1998, pp. 283–306.

[126] P. J. Rowe, C. M. Myles, S. J. Hillmann, and M. E. Hazlewood. "Validation of flexible electrogoniometry as a measure of joint kinematics". In: *Physiotherapy* 87.9 (2001), pp. 479–488.

[127] Mikel Sagardia, Katharina Hertkorn, David Sierra González, and Claudio Castellini. "Ultrapiano: A novel human-machine interface applied to virtual reality". In: *IEEE International Conference on Robotics and Automation (ICRA)*. IEEE. 2014, pp. 2089–2089.

[128] Septimiu E. Salcudean, Wen Hong Zhu, P. Abolmaesumi, Simon Bachmann, and Peter D. Lawrence. "A robot system for medical ultrasound". In: *Robotics Research*. Springer, 2000, pp. 195–202.

[129] Gaspare Santaera, Emanuele Luberto, Alessandro Serio, Marco Gabiccini, and Antonio Bicchi. "Low-cost, fast and accurate reconstruction of robotic and human postures via IMU measurements". In: *IEEE International Conference on Robotics and Automation (ICRA)*. IEEE. 2015, pp. 2728–2735.

[130] Massimo Sartori, Monica Reggiani, Dario Farina, and David G. Lloyd. "EMG-driven forward-dynamic estimation of muscle force and joint moment about multiple degrees of freedom in the human lower extremity". In: *PLOS ONE* 7.12 (2012).

[131] Verena J. M. M. Schrier, Stefanie Evers, Johan G. Bosch, Ruud W. Selles, and Peter C. Amadio. "Reliability of ultrasound speckle tracking with singular value decomposition for quantifying displacement in the carpal tunnel". In: *Journal of Biomechanics* 85 (Mar. 2019), pp. 141–147. DOI: 10.1016/j.jbiomech.2019.01.022.

[132] Ali Shafti, Ahmad Ataka, B. Urbistondo Lazpita, Ali Shiva, Helge A. Wurdemann, and Kaspar Althoefer. "Real-time robot-assisted ergonomics". In: *IEEE International Conference on Robotics and Automation (ICRA)*. IEEE. 2019, pp. 1975–1981.

[133] Yu She, Deshan Meng, Hongliang Shi, and Hai-Jun Su. "Dynamic modeling of a 2D compliant link for safety evaluation in human-robot interactions". In: *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE. 2015, pp. 3759–3764.

[134] Jianbo Shi and Carlo Tomasi. "Good features to track". In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, 1994, pp. 593–600. ISBN: 0818658274. DOI: 10.1109/cvpr.1994.323794.

[135] Jun Shi, Yong-Ping Zheng, Xin Chen, and Qing-Hua Huang. "Assessment of muscle fatigue using sonomyography: Muscle thickness change detected from ultrasound images". In: *Medical Engineering & Physics* 29.4 (2007), pp. 472–479. DOI: 10.1016/S0929-8266(00)00084-7.

[136] Patrice Y. Simard, Dave Steinkraus, and John C. Platt. "Best practices for convolutional neural networks applied to visual document analysis". In: *IEEE International Conference on Document Analysis and Recognition (ICDAR)*. IEEE. 2003, p. 958.

[137] M. Hongchul Sohn, J. Lucas McKay, and Lena H. Ting. "Defining feasible bounds on muscle activation in a redundant biomechanical task: practical implications of redundancy". In: *Journal of Biomechanics* 46.7 (2013), pp. 1363–1368.

[138] Satoshi Suzuki and Keiichi A. Be. "Topological structural analysis of digitized binary images by border following". In: *Computer Vision, Graphics and Image Processing* 30.1 (1985), pp. 32–46. DOI: 10.1016/0734-189X(85)90016-7.

[139] Jan Swammerdam. "The Book of Nature II". In: *London (UK): Seyffert* (1758), pp. 122–132.

[140] Andras Lasso Tamas Ungi and Gabor Fichtinger. "Open-source platforms for navigated image-guided interventions". In: *Medical Image Analysis* 33 (Oct. 2016), pp. 181–186. DOI: 10.1016/j.media.2016.06.011.

[141] C. Tomasi and R. Manduchi. "Bilateral filtering for gray and color images". In: *IEEE International Conference on Computer Vision (ICCV)*. 1998, pp. 839–846. DOI: 10.1109/iccv.1998.710815.

[142] Bruno M. Trindade, Yuu Ono, Edward D. Lemaire, and Ibrahim AlMohimeed. "Development of a wearable ultrasonic sensor and method for continuous monitoring of mechanical properties of plantar soft tissue for diabetic patients". In: *Proceedings of the IEEE International Ultrasonics Symposium (IUS)*. IEEE. 2014, pp. 2112–2115.

[143] Ursula Trinler, Hermann Schwameder, Richard Baker, and Nathalie Alexander. "Muscle force estimation in clinical gait analysis using AnyBody and OpenSim". In: *Journal of Biomechanics* 86 (2019), pp. 55–63.

[144] Yoichiro Tsutsui, Takayuki Tanaka, Shun'ichi Kaneko, and Maria Q. Feng. "Duplex ultrasonic muscle activity sensor". In: *Proceedings of IEEE SENSORS*. IEEE. 2005, p. 4.

[145] Gijs van Tulder. *Elastic deformations for N-dimensional images (Python, SciPy, NumPy, TensorFlow)*. 2018. URL: https://github.com/gvtulder/elasticdeform.

[146] Tamas Ungi, Gabrielle Gauvin, Andras Lasso, Caitlin T. Yeo, Padina Pezeshki, Thomas Vaughan, Kaci Carter, John Rudan, C. Jay Engel, and Gabor Fichtinger. "Navigated breast tumor excision using electromagnetically tracked ultrasound and surgical instruments". In: *IEEE Transactions on Biomedical Engineering* 63.3 (2016), pp. 600–606.

[147] Marco A. Vaz. "Mechanism of Muscle Vibrations During Stimulated and Voluntary Isometric Contractions of Mammalian Skeletal Muscle". PhD thesis. University of Calgary, 1996. URL: https://www.lume.ufrgs.br/handle/10183/1485.

[148] Salvatore Virga, Rüdiger Göbl, Maximilian Baust, Nassir Navab, and Christoph Hennersperger. "Use the force: deformation correction in robotic 3D ultrasound". In: *International Journal of Computer Assisted Radiology and Surgery* 13.5 (2018), pp. 619–627.

[149] Jing-jing Wan, Zhen Qin, Peng-yuan Wang, Yang Sun, and Xia Liu. "Muscle fatigue: general understanding and treatment". In: *Experimental & Molecular Medicine* 49.10 (2017).

[150] Gil Weinberg, Mason Bretan, Guy Hoffman, and Scott Driscoll. *Robotic Musicianship: Embodied Artificial Creativity and Mechatronic Musical Expression*. Springer Nature Switzerland AG, 2020. DOI: 10.1007/978-3-030-38930-7.

[151] Jack M. Winters. "Hill-based Muscle Models: A Systems Engineering Perspective". In: *Multiple Muscle Systems*. Springer, 1990, pp. 69–93.

[152] William Hyde Wollaston. "On the duration of muscular action". In: *Philosophical Transactions of the Royal Society of London* (1810), pp. 1–5.

[153] Hang Xu, Donald Bloswick, and Andrew Merryweather. "An improved OpenSim gait model with multiple degrees of freedom knee joint and knee ligaments". In: *Computer Methods in Biomechanics and Biomedical Engineering* 18.11 (2015), pp. 1217–1224.

[154] Taojin Xu, Zhongwei Jiang, Jongyeob Jeong, Minoru Morita, and Hongbin Xu. "Integrated System for Monitoring Muscular States during Elbow Flexor Resistance Training in Bedridden Patients". In: *Journal of Healthcare Engineering* 2019 (2019).

[155] Xingchen Yang, Zhenfeng Chen, Nalinda Hettiarachchi, Jipeng Yan, and Honghai Liu. "A wearable ultrasound system for sensing muscular morphological deformations". In: *IEEE Transactions on Systems, Man, and Cybernetics: Systems* 51.6 (2021), pp. 3370–3379. DOI: 10.1109/TSMC.2019.2924984.

[156] Paul A. Yushkevich, Joseph Piven, Heather Cody Hazlett, Rachel Gimpel Smith, Sean Ho, James C. Gee, and Guido Gerig. "User-Guided 3D Active Contour Segmentation of Anatomical Structures: Significantly Improved Efficiency and Reliability". In: *Neuroimage* 31.3 (2006), pp. 1116–1128.

[157] George Ireneus Zahalak. "A distribution-moment approximation for kinetic theories of muscular contraction". In: *Mathematical Biosciences* 55.1-2 (1981), pp. 89–114.

[158] Jeffrey Zhang, Sravani Gajjala, Pulkit Agrawal, Geoffrey H. Tison, Laura A. Hallock, Lauren Beussink-Nelson, Mats H. Lassen, Eugene Fan, Mandar A. Aras, ChaRandle Jordan, Kirsten E. Fleischmann, Michelle Melisko, Atif Qasim, Sanjiv J. Shah, Ruzena Bajcsy, and Rahul C. Deo. "Fully automated echocardiogram interpretation in clinical practice: feasibility and diagnostic accuracy". In: *Circulation* 138.16 (2018), pp. 1623–1635. DOI: 10.1161/CIRCULATIONAHA.118.034338.

[159] Qiang Zhang, Kang Kim, and Nitin Sharma. "Prediction of Ankle Dorsiflexion Moment by Combined Ultrasound Sonography and Electromyography". In: *IEEE Transactions on Neural Systems and Rehabilitation Engineering (TNSRE)* 28.1 (Jan. 2020), pp. 318–327. DOI: 10.1109/TNSRE.2019.2953588.

[160] Yong-Ping Zheng, M. M. F. Chan, Jun Shi, Xin Chen, and Qing-Hua Huang. "Sonomyography: Monitoring morphological changes of forearm muscles in actions with the feasibility for the control of powered prosthesis". In: *Medical Engineering & Physics* 28.5 (2006), pp. 405–415.

[161] Xiaodong Zhong, Frederick H. Epstein, Bruce S. Spottiswoode, Patrick A. Helm, and Silvia S. Blemker. "Imaging two-dimensional displacements and strains in skeletal muscle during joint motion by cine DENSE MR". In: *Journal of Biomechanics* 41.3 (2008), pp. 532–540.

[162] Guang Quan Zhou and Yong Ping Zheng. "Automatic Fascicle Length Estimation on Muscle Ultrasound Images with an Orientation-Sensitive Segmentation". In: *IEEE Transactions on Biomedical Engineering* 62.12 (2015), pp. 2828–2836. DOI: 10.1109/TBME.2015.2445345.

[163] Qing-Hong Zhu, Yan Chen, and Arie Kaufman. "Real-time Biomechanically-based Muscle Volume Deformation using FEM". In: *Computer Graphics Forum*. Vol. 17. 3. Wiley Online Library. 1998, pp. 275–284.