Movement Intent and Its Construction

by

Justin Ryan Horowitz
B.S. (Kalamazoo College) 2005
B.S., M.S. (Washington University in St. Louis) 2008

THESIS

Submitted as partial fulfillment of the requirements
for the degree of Doctor of Philosophy in Bioengineering
in the Graduate College of the
University of Illinois at Chicago, 2016

Chicago, Illinois

Defense Committee:
James L. Patton, Bioengineering, Chair and Adviser
Daniel Corcos, Kinesiology
John Hetling, Bioengineering
Sandro Mussa-Ivaldi, Northwestern University
Milos Zefran, Electrical and Computer Engeering

To my wife,

Theresa Horowitz,

without whose love and support this would not have been possible.

ACKNOWLEDGMENTS

To all of my coauthors, editors, and colleagues, thank you for putting up with me.

CONTRIBUTIONS OF AUTHORS

Chapter 1 places my work in a broader scientific context and was written solely by me to facilitate reading this manuscript. Chapter 2 is a published manuscript (1) for which I developed the conceptual framework, designed and carried out the experiment, and performed the analysis. My research mentor, Dr. James Patton, provided oversight and guidance throughout this process. I was the primary author of the paper, but the final wording of the manuscript owes more to Dr. Patton. Chapter 3 is a published manuscript (2) for which I developed the conceptual framework, designed the experiment, performed the analysis, and wrote the manuscript. The summer interns I supervised, Tejas Madhavan and Christine Massie, performed the experiment, assisted the analysis throughout, and assisted in writing and editing the rough paper. Dr. Patton identified the need for this paper, provided oversight and guidance throughout the process, and contributed heavily to revising the manuscript. Chapter 4 is a manuscript accepted for publication (3) for which I developed the conceptual framework, designed and carried out the experiment, performed the analysis, and wrote the manuscript. My co-author, Yazan Adbel Majeed, aided in designing and performing the analysis, generating polished figures, and generating and editing the manuscript. Dr. Patton provided oversight and guidance throughout the process as well as editing the manuscript. Chapter 5 is an unpublished manuscript for which I developed the conceptual framework, performed the analysis, and performed most of the writing. Yazan aided in editing the manuscript, verifying the conceptual framework, designing and performing the analysis, and designing informative and clear

CONTRIBUTIONS OF AUTHORS (Continued)

graphics. Dr. Patton also verified the conceptual framework, performed some analysis, and heavily edited the manuscript. Chapter 6 extends the implications of the previous works to a broader context and was written solely by me.

TABLE OF CONTENTS

TABLE OF CONTENTS (Continued)

CHAPTER **PAGE**

LIST OF TABLES

LIST OF FIGURES

LIST OF FIGURES (Continued)

LIST OF FIGURES (Continued)

LIST OF FIGURES (Continued)

FIGURE **PAGE**

LIST OF ABBREVIATIONS

MUE	Mean Unsigned Error
IF	Intent Feedback

SUMMARY

Humans constantly construct intentions and act on them, but neither the content of intention nor the method of its construction are well-understood. We develop a series of tools that enable us to inspect intention and its construction even when our subjects are confronted with experimental disturbances. We then formed a predictive model of intent as a state trajectory constructed by a stochastic process that minimizes the cost of action while maximizing both reward and the rate of reward. This model appears to predict and explain the statistical distributions of both the pieces composing intent and and whole motions that we observe.

CHAPTER 1

INTRODUCTION

?? Human motion has been described as "repetition without repetition": even as we repeatedly meet a goal subject to some constraints, other facets of the movement vary (7). While this matches the intuition that many different movements can accomplish the same goal, the brain appears to tune movement in various respects. Movements tend to minimize total jerk, the rate at which acceleration changes (8). Moreover, the many joints and muscles of the arm are redundant, which allows variability to be concentrated into some aspects of movement and not others (9). Finally, movement can be described as minimizing an expected cost (10). While these insights widely describe the movements humans produce, they do not make any falsifiable predictions that have not been falsified already.

Despite lingering questions about how movement is planned and controlled, its physiology is well-understood. The brain receives information about the state of the body and its environment through vision, proprioception, and tactile sensation. After processing this information, the brain can transmit signals that cause muscles to change their resting length. The elasticity of muscle combined with its shortening can generate torques that in turn bring about structured movement. Shortening muscles that oppose one another has the effect of stiffening the body, but comes at the cost of having contracted addition muscles. Accurate modeling of the mechanics underlying bone and muscle requires some mathematical complexity, but modern computer hardware is capable of performing the necessary calculations in real-time.

We extend these physical models to understand not the capacity of the brain to control movement but rather its limitations. To successfully execute a plan, the plan must reflect the body's physiology and obey the laws of physics. Together, these restrictions allow us to invert the planning process and examine each of a series of stages possibly used by the brain to construct and control reaching movements. Understanding these steps is critical to understanding movement disorders – such as stroke – as well as to potentially replicating human-like movement planning in robotics.

Our inversion of the movement planning process begins with an observation by Feldman (11): the elasticity of muscle implies the presence of a dynamic equilibrium underlying movement. While Feldman's muscle-centric formulation was eventually discredited (12), others took notice of the arm's dynamic equilibrium (13) and formulated an end-effector-centric model that enjoyed great success (6). Chapter 2 introduces this model in more detail before showing how it can be inverted and why the inversion is successful. Chapter 3 demonstrates that this dynamic equilibrium is immediately recognizable to a person as their own intent and that they can readily control it. These chapters allow us to conclude that movement planning takes the form of a trajectory: a desired state that evolves in time.

Recovering an intent trajectory does not reveal by what process that arose. Early movement research found evidence that movement is composed of more elemental pieces (14). These pieces were later found to be evident during infant development (15) and following stroke (16) as well. In chapter 4, we recover these pieces and describe their statistics. In chapter 5, we explain those statistics by deriving them from first principles. In the process of this derivation, we uncover

a potential-like quantity that bears a striking similarity to the concept of money. Much of the discussion in chapter 7 focuses on the implications of the structures uncovered in chapter 6.

CHAPTER 2

DETERMINING MOVEMENT INTENT

??

2.1 Overview

This work, written in collaboration with my adviser, was published in the peer-reviewed journal PLoS ONE on September 1st, 2015 (1). Subsequently, it received international press coverage primarily due to the nickname I had given the project: "psychic robot." While this algorithm cannot read minds or predict the future, it does leverage physical and mathematical relationships to reveal in real-time or after-the-fact what a person means to do even if they are prevented from doing it. We built on a famous paper by a member of the committee who demonstrated that if the desired trajectory of a movement and its mechanical equilibrium point were the same this could explain how humans react to unexpected force disturbances. We extend his result using a finding from nonlinear control theory to recover that intended equilibrium point from force-disturbed movements.

2.2 Introduction

Disturbances, distractions, and pathologies can interfere in many situations and prevent our actions from matching our intent. A pilot, for example, may fail to complete a maneuver because of turbulence. Such challenges are broadly present in many human-machine interactions. One may speculate on how we might use these very same machines to elucidate the underlying

intention. Such a possibility would be broadly useful in any area where intended actions might be thwarted by disturbances.

It is not necessary for the brain to represent intent explicitly to generate action. Instead the nervous system might simply learn the relationship between muscle activations and the accomplishment of goals. For example, one might simply evaluate a motion outcome and adjust descending signals to the muscles. At the same time, one may not deny the fact that there are tasks that require an explicit representation of a trajectory, such as performing a dance, drawing a picture, or conducting an orchestra. For these tasks, intent is likely represented at some level in the nervous system. Nevertheless, the mechanics of the body dictate the existence of an equilibrium whether intent is explicitly represented or not. Here our goal was to answer a modest question, "*when* movement intent is a changing equilibrium, can we recover it despite external disturbances?"

The words *intent* and *equilibrium* are contentious because of the various definitions that exist in the literature. Our simple definition here is that the intent is the path that would have been taken had there been no external disturbances. In other words, action will match intent in the absence of any unexpected disturbances. We take intent to be a dynamic *equilibrium* when action approaches intent over time following any sufficiently small disturbance.

Dynamic models that strive to understand intent are wide ranging, and can include mathematical models that range in application from swarm prediction to athletic performance. The plant (such as the actions of a crowd of people or the motions of an opposing team) may be too complex to model with simple linear transformations as we present above. However, by

substituting any approximation or lookup table of input-outcome tendencies, it is possible in a variety of applications to obtain the intent, even when disturbed. In other words, it may be possible to intervene before a car crashes or an opposing team scores by knowing the intent behind their action.

Here we propose a general method, *intent extraction*, which calculates the intended trajectory even in the face of disturbances by relating environmental interaction and state to intent. The method is a class of filters that can infer the intended trajectory from the disturbance (turbulence), process (dynamic equations), actuator (muscles), and motion. The method affords new ways to study models of motor control, understand the timing and composition of intent, and how it might be altered by disturbances and injuries. Below we first present the mathematics underlying intent extraction. Then, for simplicity in this initial study, we create synthetic reaching data to examine the quality and uncertainty of this method. Finally, we demonstrate its effectiveness extracting intended action from real human reaches that have been disrupted by unpredictable disturbances. We demonstrate this effectiveness by testing the hypothesis that disturbance does not change intended action for at least 120 milliseconds. This early experiment also provides some new insights on how intended trajectories can change in reaction to disturbances.

2.3 Methods

The sections below describe the theoretical method used and the means by which we evaluated it. First, we present an idealized system using synthetic data that demonstrated the concept and provided an understanding of the computational process. Next, we present an

experimental study on humans that evaluated success in a simple test of the approach's ability to extract straight intentions even when disturbed.

2.3.1 Intent Extraction: Deducing the Desired Trajectory

We first demonstrate the *intent extraction* approach in human motor control, but we later show that the process is applicable to any controlled process with an invertible highest order plant term. The process begins by presuming a model of the controller. Here, we choose the well-known motion control structure of Shadmehr and Mussa-Ivaldi (6) where the feedforward aspect of the controller perfectly predicts the plant and linearizes the system through cancellation. Additionally, linear feedback rejects position and velocity error. This model was chosen to help illuminate the approach, but later we present how to generalize this approach to more complex models without loss of generality. The equation governing the passive planar dynamics (plant) of musculoskeletal structure is of the form,

$$\underbrace{\overbrace{M(q)\ddot{q}}^{\text{Inertia}} + \overbrace{G(q,\dot{q})}^{\text{Coriolis, Centripetal}}}_{\text{Plant}} + E = 0 \tag{2.1}$$

where M is the mass matrix, q is the joint angles, $\dot{q}$ is joint angular velocity, $\ddot{q}$ is joint angular acceleration, G contains both Coriolis and centripetal effects, and E is any externally-applied

torque. The motion behavior changes with the addition of feedforward and/or feedback controllers,

$$\underbrace{\overbrace{M(q)\ddot{q}}^{\text{Inertia}} + \overbrace{G(q,\dot{q})}^{\text{Coriolis, Centripetal}}}_{\text{Plant}} + E = \underbrace{\overbrace{\hat{M}(q_d)\ddot{q}_d}^{\text{Inertia}} + \overbrace{\hat{G}(q_d,\dot{q}_d)}^{\text{Coriolis, Centripetal}} + \hat{E}}_{\text{Feedforward Controller}} + \underbrace{K_p(q_d - q) + K_d(\dot{q}_d - \dot{q})}_{\text{Impedance, Feedback Controller}} \tag{2.2}$$

where terms with hats over them ($\hat{M}$, $\hat{G}$, or $\hat{E}$) indicate that they represent the nervous system's best estimate of the forces and dynamics it will encounter, which is also known as an internal model (6). This portion of the system serves as an inverse-dynamics feedforward controller that cancels out the dynamics of the arm in the torque balance. If the nervous system has sufficient experience and is expecting E, it is included as part of the internal model, $\hat{E}$; otherwise $\hat{E}$ is set to zero. K_p and K_d are the lumped impedance and feedback terms that employ a moving state equilibrium to accomplish the desired trajectory, q_d. This q_d has a dual meaning in that it signifies both the unknown desired trajectory that we seek to discover and also the moving equilibrium trajectory of the arm.

For typical dynamic simulations, in order to determine the trajectory of the system (i.e., the forward dynamics problem), this second-order differential equation is solved by numerical integration to determine the solution to the initial value problem in time. This entails algebraic manipulation to solve for $\ddot{q}$, followed by integration to determine the state trajectory. Intent determination takes the novel approach of instead solving for $\ddot{q}_d$,

$$\ddot{q}_d = \hat{M}(q_d)^{-1}\left\{M(q)\ddot{q} + G(q,\dot{q}) + E - [\hat{G}(q_d,\dot{q}_d) + \hat{E} + K_p(q_d - q) + K_d(\dot{q}_d - \dot{q})]\right\} \tag{2.3}$$

such that $\ddot{q}_d$ can then be integrated numerically using a differential equation solver to determine the intended state trajectory $q_d(t)$. This relies on many assumptions: The model of plant and controller must be accurate and precise. The initial conditions must be available and accurate. The mass matrix estimate, $\hat{M}$, must be invertible. Externally-applied force must be precisely and accurately measured. If all of these conditions are met, then the system yields an accurate estimate of the intent.

To explore the conditions under which this procedure for recovering the intent trajectory can be generalized, we take a series of steps. First, we distinguish between joint space and generalized space by introducing a generalized state variable, x. Second, we separate the plant (arm) into its process (dynamics), P, and actuator (muscles), A, components. Third, we model these components in a very general sense as operations,

$$\underbrace{\overbrace{\sum_{n=1}^{N} P_n x^{(n)}}^{\text{Process}} + \overbrace{\sum_{m=1}^{M} A_m (x - x_e)^{(m)}}^{\text{Actuator}}}_{\text{Plant}} + E = 0 \tag{2.4}$$

where x_e is the equilibrium trajectory of the actuator. Fourth, we describe a control law for the actuator using the same expansions,

$$\overbrace{\sum_{n=1}^{N} \hat{P}_n x_d^{(n)} + \hat{E}}^{\text{Internal Model}} + \overbrace{\sum_{m=1}^{M} A_m (x_d - x_e)^{(m)}}^{\text{Actuator}} = 0 \tag{2.5}$$

where an internal model, $\hat{P}$ and $\hat{E}$, predicts system dynamics and external disturbances in order to determine the actuator equilibrium x_e such that x will track a desired path x_d. The

feedforward component must use the actuator and hence it shares the actuator's equilibrium with the plant. This allows us to finally relate the physical system and its control law by solving for x_e and substituting into Equation 2.4,

$$\overbrace{\sum_{n=1}^{N} P_n x^{(n)}}^{\text{Process}} + E = \overbrace{\sum_{n=1}^{N} \hat{P}_n x_d^{(n)} + \hat{E}}^{\text{Internal Model}} + \overbrace{\sum_{m=1}^{M} A_m (x_d - x)^{(m)}}^{\text{Feedback (Actuator)}} \tag{2.6}$$

wherein x_e vanishes recovering our familiar model. While a proper choice of x_e is perhaps essential for control, our approach does not require a model relating this equilibrium and intended trajectory. Note also that the actuator's equilibrium x_e is not the process's equilibrium unless all derivatives of x_d are zero.

The fact that the actuator's equilibrium x_e is not the process's equilibrium is important as the actuator state (λ) has been claimed to be the intended movement (11) (see Section 2.6). This upholds Gomi and Kawato's finding that the arm muscles' equilibrium does not represent reaching intent (12). If the highest order coefficient of the dynamic model of the process $\hat{P}_N$ can be inverted, the impedances can be modeled, and the system is stable (e.g., the stiffness and damping are positive), it is possible to solve for $x_d^{(N)}$ and integrate for x_d, revealing the intended trajectory.

While we treat both the process and actuator as dynamic models in our derivation, it is important to note that both the derivation and the technique are agnostic to the process model, simply because any hypothesis can be modeled and used by our approach to determine where the model would have gone had it not been disturbed.

Note that instead of a dynamic model, any model may be possible, such as lookup table. For instance, a switch is well-modeled as a lookup table or threshold without any consideration of its underlying mechanism. In this case, our technique could be used to determine a person's intention to flip the switch and how it changes in the face of disturbance. As long as some bidirectional relationship exists between state and outcome (even if determined empirically), intended outcome can be determined. This allows determination of intent in many situations of interest, even where the process is otherwise irreducibly complex.

2.3.2 Experimental Design

We chose 15 centimeter long simulated reaches of the right arm, beginning at 38 centimeters out from and 5.7 centimeters left of the right shoulder and ending at 38 centimeters out from and 10.7 centimeters right of the right shoulder. We used the following two types of perturbing forces for each combination of distance and direction:

1. Pulse forces applied in one of the the two directions perpendicular to the direction of movement began when the subject had moved either 10% or 50% of the distance to the target and lasted for 150ms.
2. Noise forces began once the subject moved 3 millimeters, and lasted for the duration of the motion. The forces were drawn from a white noise generator at 1000 Hz with flat power spectral density of 1N, and then passed through a 4th order low-pass Butterworth filter with cutoff 10π rad/s.

Our error metric was deviation in position, where unsigned error was calculated as the mean magnitude of deviation from the straight-line nominal trajectory to the target. This was eval-

uated across each of 150 bins measuring 1 millimeter in width and spaced evenly throughout the 15 centimeter reach. Mean unsigned error (MUE) was also used to summarize the overall error in each movement for sensitivity analysis. Mean and maximum signed error relative to the direction of disturbance, also called perpendicular deviation, were used to measure reaching accuracy in order to elucidate direction of any corrections.

2.3.3 Dynamic simulation of arm and intended trajectories

While the derivation of intent determination is general, testing its application to human reaching requires choosing plant and actuator models, which themselves require physical parameters. Appropriate plant modeling is well-understood: measured arm segment lengths and self-reported body mass are converted into the plant's inertial, centripetal, and Coriolis terms. Anatomical landmarks and values from Dempster (17) and Winter (18) relate body mass to limb mass, limb length to limb center of mass, and limb mass and length to moment of inertia. For the actuator, we choose the model and parameters of Burdet et al. (4). We fit a subject-specific constant scale factor for impedance relative to this model to account for any task-dependence of the subjects' impedance (19) as described in the next section. Desired trajectory in time was idealized as a typical minimum jerk, 5th order polynomial of duration 700 ms starting and ending with zero velocity and acceleration (8). In this model, musculoskeletal stiffness was linearly related to torque. Like Burdet et al. (4), we approximated muscle torque by looking backward in time 500μs. Note that this backward look led to negligible (sub-micrometer mean unsigned error) discrepancies in subsequent intent estimates.

2.3.4 Torque Calculations for Intent Extraction

The impedance properties of the human arm are known to be task-dependent (19), but we need to model them in the context of our task without assuming our conclusions in the premises. We do so by scaling an established model (4) using an intermittently and unexpectedly presented secondary task. This ensures that the model parameters are not in any sense tuned to our primary task. Additionally, unpredictability enables us to leverage the ergodic assumption: our tasks should at least initially share the same motor plan and impedance properties.

Calibrating the model requires several steps. First, we fit some expectation of the motor plan, x_d, by using a Gaussian weighted average ($\sigma = 22$ milliseconds) of each subject's hand trajectory as it evolves in time during undisturbed movement. Separate expectations were fit for each pair of movement start and end points. We then use inverse kinematics to convert these expected hand trajectories to expected joint angle trajectories, q_d. Taking time derivatives of q_d then allows us to form an expectation of the feedforward torque, τ_{ff}, that will be present early in movement as in Equation 2.2,

$$\tau_{ff} = \hat{M}(q_d)\ddot{q}_d + \hat{G}(q_d, \dot{q}_d) \tag{2.7}$$

which then allows us to construct the full torque balance for the joints,

$$\tau_P + E = \tau_{ff} + c\tau_{fb} \tag{2.8}$$

where E is the measured force on the hand converted to joint torques via the Jacobian, τ_P is calculated from the arm's dynamics, τ_{fb} is calculated from the impedance model of Burdet et al. (4), and c is a scalar constant to be fit. We calculate τ_P using measured limb segment lengths, L_1 and L_2 and self-reported body mass, m_g. Terms followed by a subscript 1 indicate the upper arm or shoulder joint while terms followed by a subscript 2 indicate the forearm or elbow joint. We converted from these gross measurements to specific parameters using the nominal ratios provided in Table II,

$$m_1 = 0.028 m_g \qquad L_{m1} = 0.426 L_1 \qquad J_1 = m_1 (0.322 L_1)^2 \tag{2.9}$$

$$m_2 = 0.022 m_g \qquad L_{m2} = 0.682 L_2 \qquad J_2 = m_2 (0.468 L_2)^2 \tag{2.10}$$

where m_1 and m_2 are segment masses, L_{m1} and L_{m2} are limb centers of mass, and J_1 and J_2 are mass moments of inertia. We then calculate M, G, and τ_P using these subject-specific

parameters and measured values for the joint trajectory, q, and its time derivatives from reaches that received the white noise force disturbance,

$$M = \begin{bmatrix} J_1 + J_2 + m_1 L_{m1}^2 + m_2(L_1^2 + L_{m2}^2 + 2L_1 L_{m2}\cos q_2) & J_2 + m_2(L_{m2}^2 + L_1 L_{m2}\cos q_2) \\ J_2 + m_2(L_{m2}^2 + L_1 L_{m2}\cos q_2) & J_2 + m_2 L_{m2}^2 \end{bmatrix} \tag{2.11}$$

$$G = \begin{bmatrix} m_2 L_1 L_{m2}\dot{q}_2(2\dot{q}_1 + \dot{q}_2)\sin q_2 \\ m_2 L_1 L_{m2}\dot{q}_1^2 \sin q_2 \end{bmatrix} \tag{2.12}$$

$$\tau_P = M(q)\ddot{q} + G(q, \dot{q}) \tag{2.13}$$

for the first 150 milliseconds following the onset of movement. τ_{ff} was calculated in the same fashion τ_{ff} except that $\hat{M}$ and $\hat{G}$ were calculated by replacing all instances of q with q_d. We then calculated τ_{fb} as in Burdet et al. (4) while additionally leveraging the dual identity of the muscle torque, τ_m,

$$|\tau_m| = |\tau_{ff} + \tau_{fb}| = |\tau_P + E| \tag{2.14}$$

$$K = \begin{bmatrix} 10.8 + 3.18|\tau_{m1}| & 2.83 + 2.15|\tau_{m2}| \\ 2.51 + 2.34|\tau_{m2}| & 8.67 + 6.18|\tau_{m2}| \end{bmatrix} \tag{2.15}$$

$$\tau_{fb}(t) = K\big(q_d(t) - q(t) + \frac{\dot{q}_d(t) - q(t)}{12} + \frac{q_d(t-\phi) - q(t-\phi) + 2(\dot{q}_d(t-\phi) - q(t-\phi))}{50}\big) \tag{2.16}$$

where ϕ is a delay of 60 milliseconds. Finally, constrained optimization is performed to find the value of c that minimizes the expression,

$$\tau_{ff} + c\tau_{fb} - \tau_P - E \tag{2.17}$$

for each subject as the sum of all torques acting on the joints should ideally be zero. The optimization is constrained such that $c > 0.15$ both to avoid unphysiological values and numerical instabilities during simulation. The value of c, reported in Table I, is then used when extracting intended trajectories from pulse-disturbed movements as in Equation 2.3,

$$\ddot{q}_d = \hat{M}(q_d)^{-1} \left\{ M(q)\ddot{q} + G(q,\dot{q}) + E - [\hat{G}(q_d,\dot{q}_d) + c\tau_{fb}] \right\} \tag{2.18}$$

and integrated numerically as described. By scaling an established model using a force disturbance other than our disturbance of interest, we avoided presuming the intent trajectory we hoped to recover.

2.3.5 Indices of variance-based sensitivity

To ensure a properly spaced and efficient evaluation of sensitivity, we employed Sobol-distributed matrices, generated using MATLAB's *sobolset()* function and converted to parameter distributions (see Table II) using inverse cumulative probability density. Nominal parameter values were mostly taken from Dempster (17) and Burdet et al. (4) with typical values for subject-specific parameters (height and weight) and direct measurements for device-specific parameters (force sensor noise and drift). Standard deviations reflect expected variation in

Subject Number	1	2	3	4	5	6	7	8
Gross Body Mass (kg), m_g	86.18	54.43	95.25	86.18	94.35	86.18	63.50	72.57
Upper Arm Length (cm), L_1	28	28	31	33	34	37	28	29
Forearm Length (cm), L_2	32	30	34	34.5	34	38	31	33
Shoulder Parallel Coordinate (cm)	0	-3	0	0	-2	-3	-5	0
Shoulder Perpendicular Coordinate (cm)	53	48	51	49	55	51	48	50
Feedback Torque Gain, c	0.15	0.15	0.36	0.15	0.94	0.77	0.48	0.86

TABLE I

SUBJECT-SPECIFIC PARAMETERS. GROSS BODY MASS WAS SELF-REPORTED. UPPER ARM LENGTH, FOREARM LENGTH, AND SHOULDER POSITION IN THE ROBOT'S COORDINATE SYSTEM WERE MEASURED *IN SITU*. FEEDBACK TORQUE GAIN WAS FIT AS DESCRIBED IN OUR METHODS SECTION USING THE INTERMITTENTLY PRESENTED WHITE NOISE FORCE DISTURBANCE. THIS GAIN IS RELATIVE TO THE TOTAL FEEDBACK TORQUE OUTPUT (STIFFNESS, DAMPING, AND REFLEXES) OF THE MODEL OF BURDET ET AL. (4).

repeated measurement (ie. height and weight) or measured variance (ie. trial-to-trial variation in sensor noise or uncertainty in mass ratios reported by Dempster (17)). Distributions were clamped to the range of ± 3 standard deviations in order to limit parameters to a realistic range. Combinations and calculations were made as prescribed by Saltelli et al. (5) to arrive at direct and total sensitivity indices.

2.3.6 Human Subjects

The human data trajectories analyzed here are drawn from eight subjects who gave informed consent in accordance with Northwestern University Institutional Review Board, which specifically approved this study and follows the principles expressed in the Declaration of Helsinki. Five male and three female right-handed subjects (ages 24 to 30) performed the reaches with

their right arm and were not compensated. Subjects' arm segment lengths were directly measured *in situ* while body mass was self-reported.

2.3.7 Apparatus

A planar manipulandum (described in Patton and Mussa-Ivaldi (20)) was programmed to compensate and minimize any friction or mass. The MATLAB XPC-TARGET package (21) was used to render this force environment at 1000 Hz and data were collected at 1000 Hz. Visual feedback of hand position was performed at 60 Hz using OpenGL. Closed-loop data transmission time (position measurement to completed rendering to recognition of rendering by the position measurement system) was less than 8 milliseconds, ensuring a visual delay less than one 60 Hz frame. Because force sensors tend to drift, we performed a linear re-zeroing procedure between each motion to assure unbiased measurements.

2.3.8 Protocol

Subjects made 730 reaches in total, along a line parallel to their coronal plane and approximately 45 centimeters from their shoulder. Reaches were either 15 or 30 centimeters long, starting and ending at one of three points spaced 15 centimeters apart on the line. To prevent any learning effect, forces were presented intermittently with random frequency, but never less than 5 reaches apart. Any effects of exposure to pulse forces cannot be detected 5 reaches later (22). Forces were chosen pseudorandomly such that each type, direction, and distance combination mentioned above was presented 5 times. We used these disturbed, intermittent trials for the analysis.

2.3.9 Statistical Analysis

Our principal goal was to determine if and when disturbed trajectories departed from undisturbed trajectories. To detect this departure, we use the one-tailed Student's t-Test ($\alpha = .05$) at 5 millisecond intervals to determine whether or not the deviation of the disturbed trajectories had exceeded the maximum deviations of the undisturbed trajectories. The median time span between departure of the hand and departure of the estimated intent was compared against 120 milliseconds using the nonparametric Sign Test. The MATLAB statistics toolbox package (21) was used for all comparisons.

2.4 Results

As expected, the model was able to recover the original intended trajectory even when disturbed (Figure 1). However this idealized analysis cannot reveal any vulnerabilities to unexamined model parameters or to inaccuracy in the structure of the model itself as discussed in the next sections.

2.4.1 Noise Robustness

While some inversion processes might be highly sensitive to noise, our process for recovering intent from action did not appear to be vulnerable. Our position sensors are very accurate (no detectable drift, RMS noise 1.1 micrometers), but the force sensors are less so. Even after correcting for drift in our force measurements, the force measurements reported during reaching can be expected to have noticeable bias and further noise on the order of that bias as reported in Table II. Taken together, force measurement errors account for 12% of all sensitivity. We

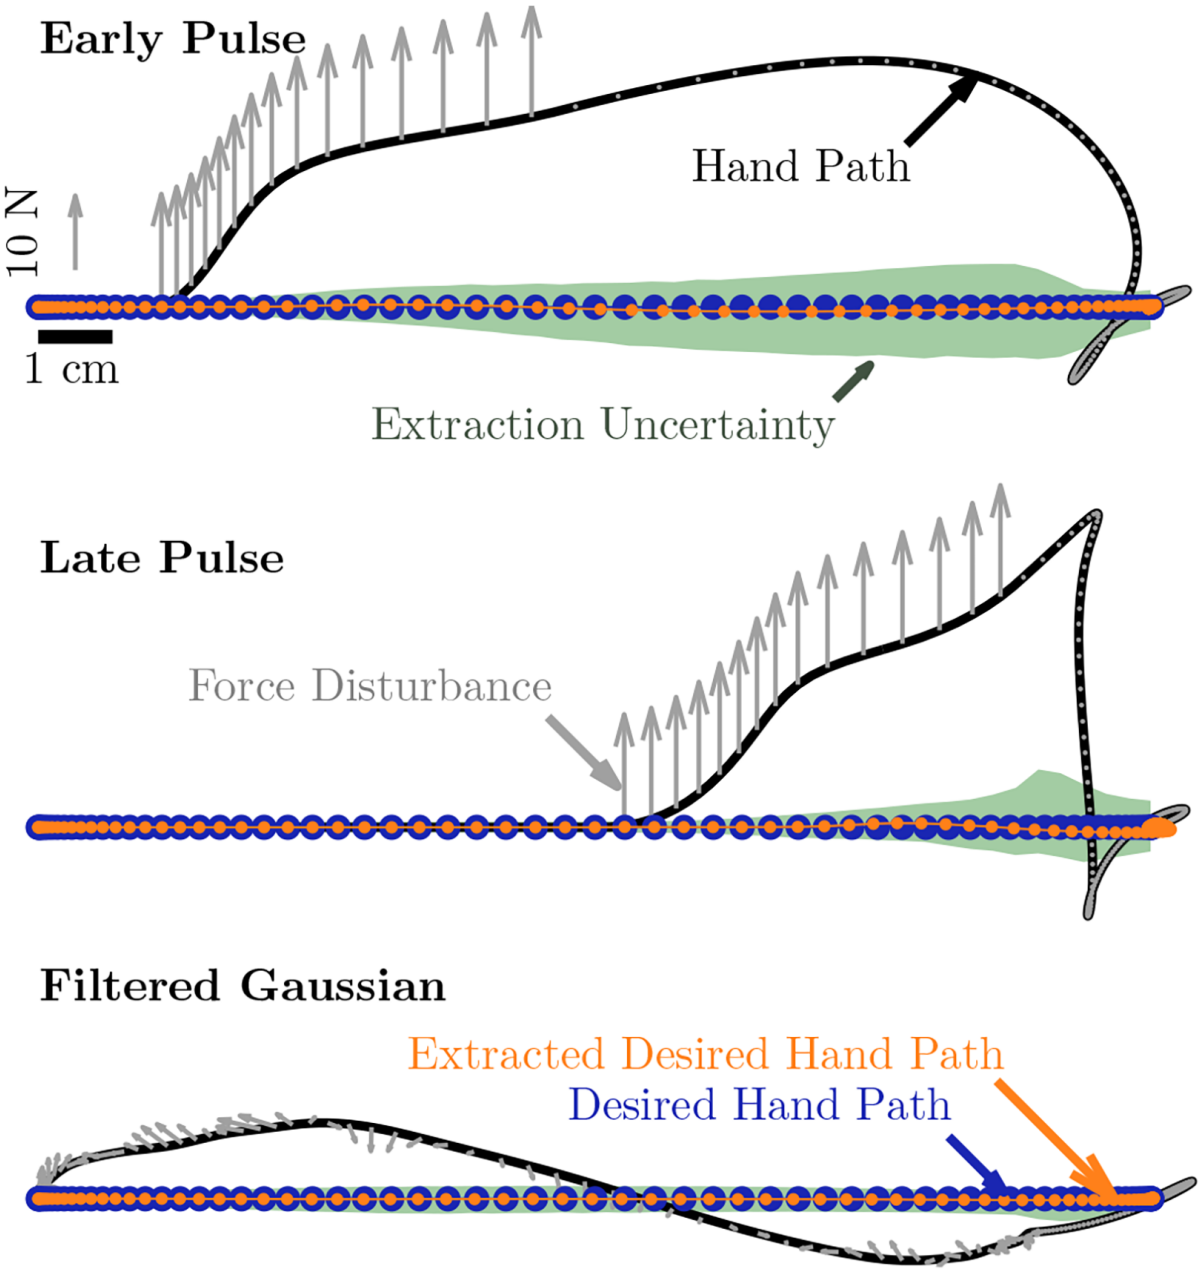

Figure 1. Simulated data illustrating tautology of extraction across pulse and filtered Gaussian noise disturbance types. Intent is modeled as a minimum jerk, 5th order polynomial. Forces experienced are combined with intent via Burdet et al.'s (4) model to produce the simulated arm trajectory. Extraction to recover intention from arm and force trajectory follows. Parameter errors are introduced into the extraction and varied according to Table II to estimate sensitivity. Examining the distribution of deviation in trajectories from this analysis reveals that sensitivity is not uniformly distributed throughout a reach.

also examined noise in sensor signals explicitly in the more comprehensive sensitivity analysis, discussed below.

2.4.2 Parameter Sensitivity

Variance in intended trajectories due to estimated uncertainty in model parameters was lower than the natural variation in undisturbed motions. Simulated point-to-point reaches disturbed by either filtered white noise forces or a pulse force perpendicular to the direction of the reach (Figure 1) were extracted in the presence of 220,000 variations upon the parameters (Table II) according to the methods of Saltelli et al. (5). Expected variance due to parameter uncertainty was 2.24mm^2 for pulse forces and .3mm^2 for filtered white forces. Variance in recorded undisturbed point-to-point reaching under the same time and reach distance conditions was 2.93mm^2. These variances describe mean unsigned error, but this uncertainty is not distributed evenly in space or time.

Simulated error due to direct parameter uncertainty (Figure 2) reached the order of millimeters and revealed particular sensitivity to mis-estimation of stiffness and changes in shoulder position from trial to trial. Direct sensitivity, which reveals the proportion of error that could be removed by correcting inaccuracy of a parameter, was an order of magnitude lower than total sensitivity, which reveals how much error would remain if all other parameters were accurate. Both implicated stiffness estimation inaccuracy as a primary cause of extraction uncertainty. Note that this sensitivity analysis does not provide information on whether models or parameter estimates are accurate, only how sensitive the model would be if they were inaccurate.

Figure 2. Model sensitivity testing (from the variance-based method of Saltelli et al. (5)) reveals that the model is mostly sensitive to stiffness. Parameter values are varied in a Sobol-distributed fashion according to Table II. First order sensitivity, in purple, is the error that would be removed if the parameter was fixed at its nominal value. Total sensitivity, in teal, shows error that would remain if all other parameters were fixed at their nominal values.

Parameter Name	Units	Nominal	Source	SD	Source	S	S_T
Upper Arm Length (L_1)	m	0.353	(17)	0.017	(17)	0.03	0.17
Forearm Length (L_2)	m	0.363	(17)	0.011	(17)	0	0.03
Upper Arm Center of Mass Ratio ($\frac{L_{m1}}{L_1}$)	1	0.436	(18)	0.0695	15%	0	0.02
Forearm Center of Mass Ratio ($\frac{L_{m2}}{L_2}$)	1	0.682	(18)	0.0431	15%	0.01	0.07
Gross Body Mass (m_g)	kg	88.4	(23)	3.1	(24)	0	0.02
Upper Arm Mass Ratio ($\frac{m_1}{m_g}$)	1	0.028	(18)	0.0029	(17)	0	0.01
Forearm Mass Ratio ($\frac{m_2}{m_g}$)	1	0.022	(18)	0.0025	(17)	0	0.06
Upper Arm Radius of Gyration Ratio	1	0.322	(18)	0.0161	15%	0.01	0.06
Forearm Radius of Gyration Ratio	1	0.468	(18)	0.0234	15%	0	0.04
Shoulder Parallel Coordinate	m	-0.057	RM	0.02	RM	0.01	0.14
Shoulder Perpendicular Coordinate	m	0.88	RM	0.02	RM	0.02	0.06
Force Sensor Miscalibration x-axis	N	0	RM	0.231	RM	0.01	0.05
Force Sensor Miscalibration y-axis	N	0	RM	0.1067	RM	0.03	0.08
Force Sensor Gaussian Noise SD x-axis	N	0	RM	0.1653	RM	0	0.01
Force Sensor Gaussian Noise SD y-axis	N	0	RM	0.2869	RM	0	0.01
Torque-Invariant Impedance Mis-estimation Ratio	1	1	(4)	0.15	15%	0.05	0.19
Torque-Varying Impedance Mis-estimation Ratio	1	1	(4)	0.15	15%	0.05	0.19
Damping-to-Stiffness Ratio (k_d)	sec^{-1}	0.0833	(4)	0.0125	15%	0.01	0.08
Reflex Impedance Scale Factor	1	0.02	(4)	0.003	15%	0.01	0.04
Reflex Damping to Stiffness Ratio (g_d)	sec^{-1}	2	(4)	0.3	15%	0	0.04

TABLE II

SYNTHETIC MODEL PARAMETERS AND THEIR ASSOCIATED MEAN (NOMINAL) VALUES AND STANDARD DEVIATIONS (SD) USED TO DETERMINE THE SENSITIVITY INDICES OF SALTELLI ET AL. (5). ALSO SHOWN ARE THE RESULTING SENSITIVITY INDICES. SENSITIVITY INDICES LESS THAN ONE-THOUSANDTH WERE REPORTED AS ZERO. TOTAL VARIANCE WAS $1.56mm^2$, SHOWING VERY SMALL AVERAGE DEVIATIONS WHEN USING THIS APPROACH. FIRST-ORDER SENSITIVITY, S, CAN BE INTERPRETED AS FRACTION OF VARIANCE REMOVABLE BY PERFECTLY CORRECTING A FACTOR. TOTAL SENSITIVITY, S_T, CAN BE INTERPRETED AS THE FRACTION LEFT AFTER CORRECTING ALL OTHER FACTORS. MOST PARAMETER VALUES COME FROM THE LITERATURE, BUT INSTRUMENT-SPECIFIC PARAMETERS WERE DETERMINED EMPIRICALLY THROUGH REPEATED MEASUREMENTS (RM). WHERE PARAMETER VALUES WERE UNAVAILABLE, 15% OF THE MEAN WAS USED AS A CONSERVATIVE ESTIMATE.

2.4.3 Human Intent Trajectories

Extraction of intent from human point-to-point reaching revealed straight-line movement that persisted for hundreds of milliseconds after the onset of disturbing forces (Figure 3). If intent can change following disturbance, it can only do so after some period of time due to processing and communication delays in the sensorimotor system that produces new motor commands. We hypothesized that intent should not diverge from the undisturbed intent path even if the actual hand was disturbed within this window of delay. In both pulse timing conditions, the intended trajectory remained straight for longer than the hypothesized 120 milliseconds ($p < 0.001$), even though the actual hand was dramatically deflected, supporting our hypothesis and the approach.

These results suggested a reproducible approach for inspecting how individuals might alter their intent in response to disturbances. Actual hand paths could no longer be explained ($p = 0.05$) by deviation in undisturbed movement within 145 milliseconds after the onset of disturbance. Despite the disturbance, three of the eight subjects' intents did not depart significantly from baseline (Figure 3, panel A). The remaining five subjects' intents deviated between 150 and 255 milliseconds after the hand deviated (Figure 3, panel B). When disturbances occurred early in reaching, intent showed some signs of correction mid-reach. When disturbances occurred later, the intent continued to the target, but interestingly, it then corrected. Intention showed systematic deflections that counteracted the direction of disturbance beginning about 200 milliseconds after the onset of the force disturbance; however, this counteraction did not immediately depart from the range of undisturbed movements (Figure 3, bottom of panel B).

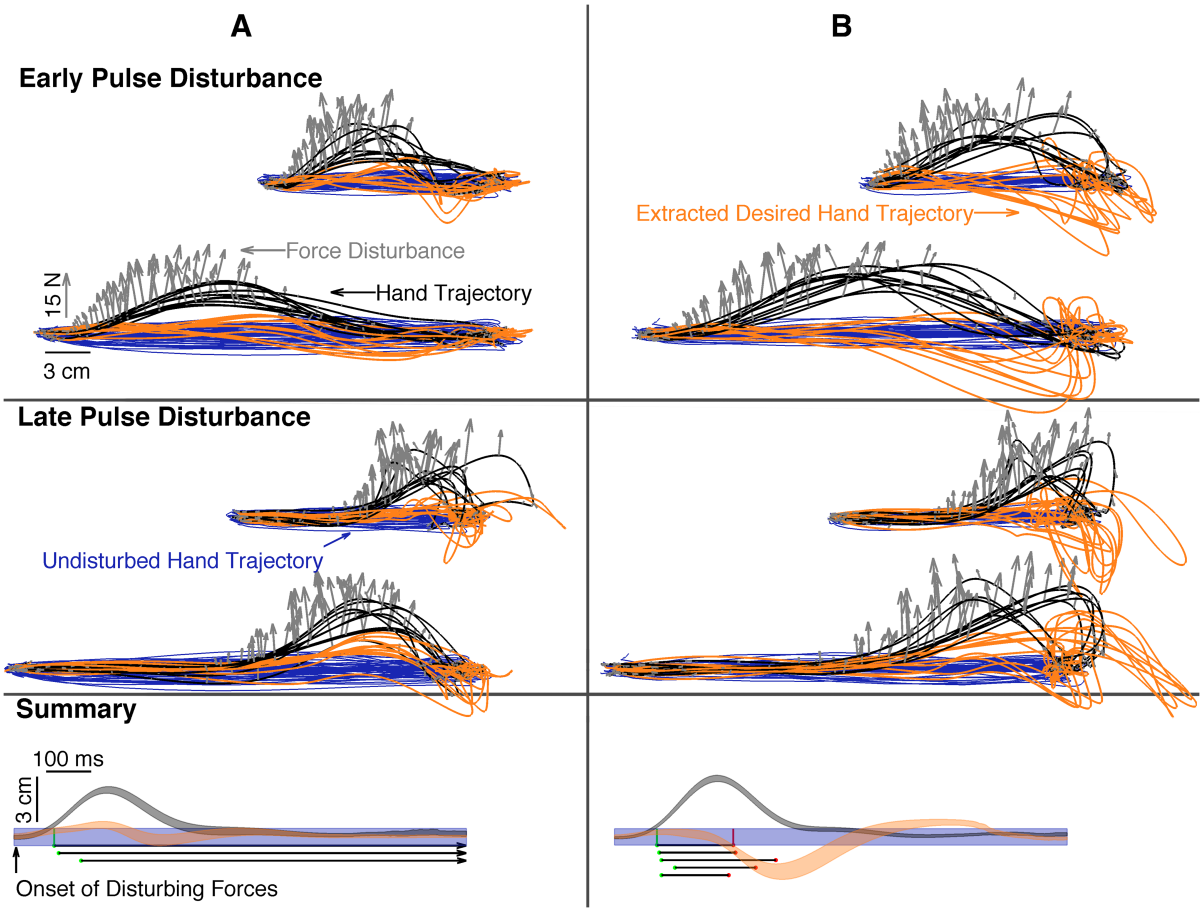

Figure 3. Subjects' disturbed hand trajectories and extraction of intended trajectories from them reveal that both variance and invariance of the motor plan can occur even in response to very large disturbances. The hand path, in black, deviates from the blue baseline, as a force pulse, gray arrows, is applied. Three of the eight subjects' intent trajectories (orange) did not significantly deviate from undisturbed movement (panel A) as observed by the intent region not departing the undisturbed movement (blue) region. Five of the eight subjects' intent trajectories did significantly deviate from undisturbed movement (panel B). The summary plots display whole-subject statistics as shaded regions (mean ± 95% confidence interval on the mean). The time at which the hand region (black) first departed from undisturbed movement (blue) is marked with a green dot for each subject. If the intent region (orange) departed from undisturbed movement, the onset of departure is marked with a red dot for each subject. Panel A subjects' intent did not depart while panel B subjects' intents did depart.

2.5 Discussion

We sought to test the suitability and robustness of an algorithm that determines a person's intent during reaching motion, even if there are disturbances. Sensitivity analysis on synthetic data revealed that errors in response to parameter variations were smaller than the trial-to-trial variance commonly observed in human reaching. Mis-estimation was largest if the stiffness was inaccurately modeled. We tested this on human reaching and found that in spite of the hand being disturbed by forces, the intended movement remained on-course to the target, possibly changing after 150 milliseconds. This represents a new and accurate method for viewing intent and how it changes when faced with a force disturbance.

While the derivation of intent determination required many steps, the outcome is a filter whose use is straight-forward. Many filters for converting assumed intent into simulated hand trajectories have been proposed and validated (6; 25; 4). Demonstrating the validity and plausibility of mathematically transforming those filters to instead convert measured hand trajectories to intent adds a novel method to the arsenal of analytic methods for exploring human motor control.

Interestingly, the formulation leads to some implications on the expected behavior of this moving equilibrium. Even if the muscle equilibrium jumped abruptly, this would be insufficient to move the intended trajectory abruptly due to the presence of the mass term in the intention formulation – reaching intent cannot jump abruptly. The intent trajectories certainly follow this pattern, smoothly responding to disturbances at a latency of about 200 milliseconds. While this work examines the tendencies of the onset of change, statistical testing cannot be used to

detect the absence of change in an individual movement. Still, Figure 3 shows reaches that are more compatible with intermittent control than with continuous optimal feedback control. In classic optimal control, the system is always updating intent in response to error (10). While most of our subjects displayed a latent response consistent with sensory feedback latencies (26), three of our eight subjects showed no detectable response. The absence of response to a large force disturbance is instead consistent with adherence to a trajectory planned (perhaps entirely) before the disturbance. Control that is intermittent (27) with limited opportunities to change intent can explain such observations. It remains to be seen if and how change in intent might be triggered.

Determination of intent may facilitate motor training and stroke recovery. Error augmentation, which presently relies on dictating the reaching intent, has demonstrated the capacity to increase and speed up learning in healthy patients (20) and following stroke (28). Augmentation of difference from static, dictated intent could be replaced with scaling of the magnitude of the difference between the desired and realized trajectory. This makes error augmentation during undirected reaching and exploration possible, including both error reduction and error magnification. Partial error cancellation would allow learning to take place without harm to task goals. As demonstrated by the success of the challenge-point framework (29), dynamic variation of augmentation as task learning progresses can be beneficial. With this extraction, error can be measured and augmented in real-time even without an explicit task, potentially enabling wider utility and application.

This novel determination of intended trajectories from force-disturbed movements allows hypothesis testing that was not previously possible. This approach can take **any** "candidate model" of the human arm and create a filter that estimates the intended trajectory. By using tests such as the comparison with undisturbed straight line reaching employed in this paper, one can evaluate which candidate produces the most plausible intended trajectory, supporting one model over another. Instead of comparing two generative models to see which better approximates movement, it is now possible to perform element-by-element fitting. Direct access to the intended trajectory allows optimal feedback control models to better understand and fit hypothesized cost functions and their parameters. In addition, mid-movement replanning became falsifiable: we were able to show that intended trajectories from disturbed motions differ from those of undisturbed motions and how they differed.

Nevertheless, there are some limitations to these methods. Accurate interface force measurement (within a few tenths of a newton) is needed. Such instrumentation is available, but must be used with great care in order to preserve the accuracy of the estimate. Our sensitivity analysis on synthetic data revealed that slight error can lead to a trajectory that accumulates a drift due to bias caused by the force term. More importantly, while modern force sensors can be highly accurate with very high signal to noise ratios, they tend to drift over time, leading to error that can grow if the device is not periodically tared. Our empirical approach used several methods to mitigate these effects, and hence the errors became small. It remains to be seen if there are methods to further reduce the mis-estimation due to force error or eliminate the

need for force estimation altogether. It will be also interesting to determine whether relaxing constraints and operating in 3D would produce the same levels of accuracy found here.

Since trajectories were not controlled experimentally, our tests rely heavily on the assumption that disturbed intentions initially mirror undisturbed intentions. This study examined straight line motions because it is well understood that people tend to repetitively attempt to reach in a straight line towards a target (30). Also, even when subjects are persistently perturbed by a force field, they recover their motions to a straight line with repetitive experience (6). Future studies might employ tasks with a more specific and/or explicit trajectory to facilitate comparison with the estimate intent's path. Unfortunately, even such an intended trajectory may not be strictly invariant during external disturbance because some displacements may cause intent to be recalculated by the nervous system.

Another key assumption in this study is that the joint impedance (stiffness) is linearly related to torque. While other approaches tend to construct paradigms that indicate or presume that the intention is static in order to derive impedance – the so called family of system identification applications, this study makes a constrained assumption of stiffness in order to determine intent. In a sense, this approach widens the scope to include movement. Sensitivity analyses showed that approximation of stiffness is the most critical factor to accuracy in estimating intent. While serious trial-to-trial stiffness mis-estimation is unlikely, co-contraction in response to disturbance might plausibly double stiffness mid-reach. Even as the model is insensitive to normal variations, a human can halve or double their stiffness (31) significantly reducing determination accuracy. Our results also point to the possibility of the impedance changing during motion, perhaps as

a function of time. This stiffening may even explain the reported deviation of intent from the baseline trajectories. Of particular concern are the large hooks at the end of the motions. While subjects might have preemptively stiffened in response to the possibility of disturbance, the many reaches in between disturbances may have induced them to relax instead. These changes in stiffness might be substantiated (or even corrected for) by detecting muscle co-contraction using electromyography or using additional modeling assumptions to allow real-time estimation of stiffness.

The simple linear scaling of Burdet et al.'s stiffness model (4) facilitated hypothesis testing, but was inherently limiting. However, more accurate approaches would provide us with many free parameters through which to determine the outcome of the eventual extraction and obtain a spurious favorable result. Like others who have examined large force disturbances, we find lower stiffnesses (our c parameter) than might be expected (32) (Table I). We explored the possibility that both the hooks and the low stiffness resulted from underestimating the contribution of reflex impedance, but using values more similar to those of Crevecouer and Scott(32) exacerbated the hooks without increasing total stiffness. Differences in tasks studied (19) and/or the dependency of muscle impedance on perturbation size (33) may explain these discrepancies.

The general intent extraction approach presented here reveals the conditions needed to solve for the moving intent in an arbitrary dynamic process. Sensitivity analysis for this approach applied to the human arm demonstrated favorable conditions for determination of intent during force-disturbed reaching. Reaching intent became significantly different from baseline reaching only after enough time had elapsed for the disturbance to be processed and descending motor

commands to change. This window into intent should allow advances in arm modeling, motor training, and human-machine interaction.

2.6 What is Intent

The word *intent* can be contentious and lead to confusion; and hence may be best placed in the context of various literature that use related terminology. It is important to distinguish intended action from *motivation* (34; 35), *cost* (10; 8), or *goal selection decisions* (36). In typical motor control studies, subjects are motivated to complete an experiment in a timely fashion, and are usually explicitly provided targets to reach to. Classification of intent is prevalent in both lower (37; 38) and upper limb (39; 40) prosthetics where hybrid control algorithms select from among a set of discrete actions (walking/standing/flexion/extension/etc). While subjects may be *motivated* to complete experiments with minimal effort/cost and their *goal* may be to reach a target, here we use intent to describe the course of action (i.e., the trajectory of the arm) taken in service of goals and motives and not the goals nor the motives themselves. Particularly of interest is the intended course of action (i.e, the intended trajectory), even when the actual movement is disturbed and hence no longer matches the intent. In other words, we operationally define *intent* as a subjunctive – where would the motion have gone had it not been disturbed?

Attempts to deduce motor intent in the past have focused on the assumed spring-like properties of human muscles. Springs produce a force according to their impedance and stretch. By measuring force, impedance, and position, Gomi and Kawato (12) were able to deduce stretch and thereby infer the muscle's equilibrium point. Supporters of the "λ model" (11) hypothe-

sized that this muscle equilibrium point, and not the equilibrium of the whole arm, represented the intent of a movement even though it did not compensate for the dynamics of the arm as a feedforward controller would. Upon Gomi and Kawato's inspection of the muscle equilibrium point as it evolved in time, it was clear that it was highly complex and often not anatomically realizable. Therefore, it could not well-represent the intent of a simple reaching movement. Unlike the equilibrium point of muscle, the a pre-planned equilibrium point of the entire arm is the path the arm will follow in the absence of disturbance (6). In the presence of disturbance, the arm might be deflected from its equilibrium. We explored if or how this arm equilibrium point might change due to disturbances.

Our findings could explain the discovery by Bizzi et al. (13) of a "virtual trajectory" (α) in deafferented monkeys that progressed smoothly from the initial to final positions and how that virtual trajectory is different from λ, the equilibrium point of muscles – intent and λ are equivalent only when intent is unchanging; therefore, this muscle equilibrium cannot be the intent of a movement. By deriving the technique in a general form, we discovered that our intent is equivalent to the virtual trajectory and can be used to determine the muscle equilibrium of a λ model (See Methods section).

Acknowledgements

Research supported by US NIH NINDS1 R01 NS053606. We thank Felix Huang, Jagan Jimmy, Christine Massie, Tejas Madhavan, Mohi Ahmed, Dayna Larson, and Yazan Abdel Majeed for their review of preliminary manuscript. We also thank the community of the

Robotics Lab at the Rehabilitation Institute of Chicago for critical commentary during the genesis of this paper and the work leading up to it.

CHAPTER 3

REALTIME FEEDBACK OF MOVEMENT INTENT

??

3.1 Overview

This work, written in collaboration with my adviser and two summer interns, was published in the peer-reviewed journal Frontiers in Behavioral Neuroscience on January 12th, 2016 (1). While the previous paper established that *if* a path is intended it must be the path we recover, it remained unclear whether that path was explicitly planned or merely arose from the physics of the arm. To address this question, we allowed subjects to interact with their intent in real-time to verify that they could recognize and control it.

3.2 Introduction

Humans often interact with machines in uncertain and complicated environments, such as crowds and traffic, where they must contend with turbulence, moving obstacles, distractions, and disturbances. Despite our capacity to learn and adapt, some environments evolve too quickly or with too much uncertainty for meaningful learning. Human and animal nervous systems intelligently solve many problems by planning ahead (41) and suppressing suboptimal actions (42), yet in the face of uncertainties we often cannot adequately prevent errors. There is the possibility, however, to exploit additional information from instruments – particularly fast and accurate force sensors – that can measure human machine interactions. Combining

sensors with filtering techniques makes it possible to determine a person's underlying *intent*, operationally defined as the motion they would have made had they not been disturbed. While other components of a movement, such as its goal, are also intended (42), our work here addresses only the intended trajectory. This intent provides new ways to understand the nature of control and provide novel feedback.

Recent work in our laboratory has attempted to outline a methodology for obtaining estimates of intent (1). This method assumes a model of the dynamics and control of the human arm. Following manipulations of the equations of motion, the method integrates to find a unique estimate of intent. The algorithm recovers the trajectory a person intended to take, even if they were forced away from it due to environmental disturbances. This analysis has enabled us to show how some subjects sometimes alter their intent following exposure to unexpected force pulses.

A new question that arises is whether seeing one's own intent, rather than what actually happens, may be useful. The intent extraction method can be streamlined to allow for real-time estimations of intent that can be presented to the subject as a cursor. Estimated intention may outperform the movement accuracy in the presence of unexpected disturbances. If so, such a method holds great promise in any situation where humans and machines interact as it enables the machine to give the human operator what they want. This human-machine collaboration could outperform what a person can do alone.

Displaying anything other than what truthfully happens is a distortion and a deceit. Like many other visual distortion experiments (43; 44), intent feedback (IF) introduces a visuomo-

tor discrepancy that may be confusing to the nervous system. Preplanning a specific route may not be necessary, and instead the system might continuously react to any environmental disturbances until it reaches the goal. If people try to achieve a goal while minimizing some measure of cost, it is possible to compute a set of rules for reacting to the environment (10). No specific intended route is needed when using optimal feedback control. While this modeling strategy has been very successful at explaining data, it fundamentally assumes that corrective actions will be taken in response to relevant disturbances. Goals can also be reached at minimal *expected* cost by constructing – and possibly updating – a specific intended route. If no particular trajectory is intended, the nervous system could be unable to recognize IF.

While performance is the best indicator of IF's worth, changes in arm stiffness can provide supporting evidence that subjects are actually getting what they want. Arm stiffness is known to increase during exposure to instability (25) and uncertainty (45). We anticipated that these changes might also be modulated by the presence of IF. Reducing the effects of environmental instability and presenting the subject with a signal already known to them should relax their arm and make it more compliant. We hypothesized that any elevated arm stiffness from noisy disturbances would decrease while subjects received IF.

In this paper, we describe our streamlined method for real-time IF to which we exposed subjects in an unstable and unpredictable environment. Their ability to perform goal-directed reaching using visual feedback of their hand position was compared against IF. We hypothesized that IF should lead to better performance in the presence of force-based disturbances.

Accordingly, we hypothesized that during random disturbances, the intent trajectory should deviate less than the hand trajectory.

3.3 Materials & Methods

3.3.1 Intent Extraction

The well-known motion control structure of Shadmehr and Mussa-Ivaldi (6) relates arm trajectory, q, to desired arm trajectory, q_d, and any external disturbance, E using physical parameters of the arm. To show how this model can be algebraically inverted to instead describe desired arm trajectory as a function of arm trajectory and external disturbance, we write it as a torque balance:

$$\underbrace{\overbrace{M(q)\ddot{q}}^{\text{Inertia}} + \overbrace{G(q,\dot{q})}^{\text{Coriolis, Centripal}}}_{\text{Plant}} + E = \underbrace{\overbrace{\tau_{ff}}^{\text{Feedforward}} + \overbrace{\tau_{fb}}^{\text{Feedback}}}_{\text{Controller}} \tag{3.1}$$

Where M is the mass matrix, q is the joint angles, $\dot{q}$ is joint angular velocities, $\ddot{q}$ is joint angular accelerations, and G contains both Coriolis and centripetal effects. Typical applications solve this torque balance for $\ddot{q}$ and use a numerical differential equation solver to predict arm trajectory in the context of a disturbance of interest, a feedback model, and feedforward torques determined by inverse dynamics. Rather than test hypotheses regarding the learning, production, or composition of this feedforward torque, we instead solved for it:

$$\tau_{ff} = M(q)\ddot{q} + G(q,\dot{q}) + E - \tau_{fb} \tag{3.2}$$

Then we noted that feedforward torque can have a one-to-one correspondence with desired acceleration, $\ddot{q}_d$:

$$\hat{M}(q_d)\ddot{q}_d + \hat{G}(q_d, \dot{q}_d) + \hat{E} = \tau_{ff} \tag{3.3}$$

Hats ($\hat{}$) denote the nervous system's best estimate of a physical quantity. Combining these expressions, suppressing state dependencies, and solving for $\ddot{q}_d$:

$$\ddot{q}_d = \hat{M}^{-1}\{M\ddot{q} + G - \hat{G} + E - \hat{E} - \tau_{fb}\} \tag{3.4}$$

In this form, a differential equation solver can determine q_d as it evolves in time if a few assumptions are made and conditions are met. First, $\hat{E}$ must be modeled or assumed, so we chose $\hat{E} = 0$. In the presence of a zero mean white noise force disturbance, its mean should be zero, but it is unlikely to be exactly zero and may reflect an average of only the last few exposures (46). Next, the matrix $\hat{M}(q_d)$ must be invertible, but we ensured this through our choice of workspace. Finally, feedback torque requires a model of arm impedance, which is known to be task-dependent (19) and may vary over the course of a reach (47). With no prior knowledge of arm impedance for this task-disturbance combination, we presumed the feedback torque model of (6) anticipating that it is sufficiently accurate or easy to learn (experiment 1). The experiment was repeated with a lower stiffness estimate (experiment 2) to explore any dependence on this assumption.

3.3.2 Apparatus

A planar manipulandum (described in (20) and depicted in Figure 4) was programmed to minimize any friction or mass. The MATLAB XPC-TARGET package (21) was used to render this force environment at 1000 Hz and data was collected at 1000 Hz. Visual feedback of cursor position was performed at 60 Hz using OpenGL. Closed-loop data transmission time (position measurement to completed rendering to recognition of rendering by the position measurement system) was less than 8 milliseconds, ensuring a visual delay less than one 60 Hz frame. Numerical simulation was performed in real-time using the GNU Scientific Library's odeiv2 driver with Runge-Kutta-Fehlberg (4,5) stepping (48). Visual feedback was given using an opaque screen that prevented subjects from seeing their arm during movement.

3.3.3 Human Subjects

The human data trajectories analyzed here are drawn from sixteen subjects who gave informed consent in accordance with Northwestern University Institutional Review Board, which specifically approved this study and followed the principles expressed in the Declaration of Helsinki. Fourteen male and two female right-handed subjects (ages 21 to 30) performed the reaches with their right arm and were not compensated. Subjects' arm segment lengths were directly measured *in situ* while body mass and handedness were self-reported.

3.3.4 Experimental Design

Subjects performed center-out reaches of the right arm to one of three visually-presented targets 15 centimeters from the center and chosen at 120° intervals. These targets were represented as a red circle with a radius of 1 centimeter. Target selection was carried out pseudorandomly

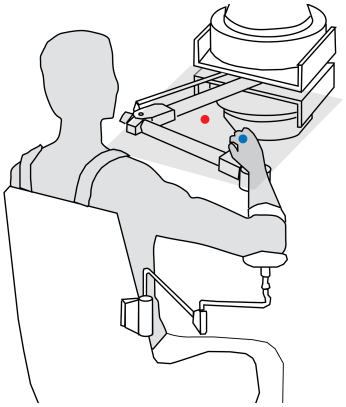

Figure 4. Subjects were seated at a planar manipulandum capable of measuring position and force as well as rendering forces. The subject's hand was positioned below an opaque screen so the subject could not see their hand as they reached towards the targets. On the screen, a red circle (the target to reach to) appeared on the screen and subjects were shown a blue circle that either represented the actual position of their hand or their estimated intent as they moved toward the target depending on the movement block.

such that each outer target was visited 16 times in all five blocks of 96 reaches each. During blocks 2 through 4, subjects experienced filtered white noise forces drawn from a white noise generator at 1000 Hertz with flat power spectral density of 1 Newton. Forces were then passed through a 4th order low-pass Butterworth filter with cutoff 10π radians per second. In all blocks, except block 3, cursor position (represented as a blue circle with a radius of 1 centimeter) indicated the subjects' actual hand positions. In block 3, the cursor position indicated the subjects' estimated intents. Once the cursor (blue circle) made contact with the target (red circle), a new target was immediately presented. Both actual hand position and intent were recorded at all times, even though at any given moment only one was visible to the subject.

3.3.5 Dynamic simulation of arm and intended trajectories

Anatomical landmarks and values from (17) and (18) were used to estimate relationships between body mass, limb mass, limb length, limb center of mass, and moment of inertia. Viscosity parameters, K_d, were taken from (6). Stiffness parameters were either taken from (6) (K_{P1}, experiment 1) or estimated (K_{P2}, experiment 2). Expressed in Newton-meters per radian:

$$K_{P1} = \begin{bmatrix} 15 & 6 \\ 6 & 16 \end{bmatrix} \quad K_{P2} = \begin{bmatrix} 8 & 2 \\ 2 & 5 \end{bmatrix} \tag{3.5}$$

To estimate this reduced stiffness, a pilot subject was asked to intend to remain still on the center target, q_d, while co-contracting as little as possible. K_{P2} was then calculated from one minute of white noise forces, E, and joint angle traces, q as the least squares solution to the system:

$$K_{P2}\big(q_d(t) - q(t)\big) = M\big(q(t)\big)\ddot{q}(t) + G(q(t), \dot{q}(t)) + E(t) + K_D\dot{q}(t) \tag{3.6}$$

where K_{P2} is a 2-by-2 matrix while the state difference and torque are 2-by-60000 matrices.

Feedback torque was calculated as the sum of viscous and elastic impedances

$$\tau_{fb} = K_D(\dot{q}_d - \dot{q}) - K_P(q_d - q) \tag{3.7}$$

with K_P chosen as described and K_D taken from (6) as expressed below in Newton-meters per radian-second.

$$K_D = \begin{bmatrix} 2.3 & 0.09 \\ 0.09 & 2.4 \end{bmatrix} \tag{3.8}$$

3.3.6 Metrics and Statistical Analysis

Trajectories and forces were rotated such that movement and force parallel to the line connecting the previous target (the reach origin) and the presented target were along a *progress* axis, while perpendicular movement and force were along an *error* axis. Reach onset was detected as the moment the cursor's distance from the center of the previous target first exceeded 1 centimeter. Maximum perpendicular error for a trajectory was the largest error magnitude within 250 milliseconds of reach onset. A scalar stiffness, k, was calculated for the *error* axis during this same 250 milliseconds time span by linear regression:

$$F_e = m\ddot{e} + b\dot{e} + ke + F_O \tag{3.9}$$

Force (F_e) and state ($\ddot{e}$, $\dot{e}$, e) were known. Mass (m), viscosity (b), and stiffness offset (F_0) terms were calculated, but discarded. While joint stiffness is usually described as a matrix, instantaneous endpoint stiffness in only the error direction is a scalar. This *effective stiffness* metric isolated stiffness in the error direction and facilitated statistical comparison between treatments and blocks. The paired t-test was used to detect differences in maximum perpendicular error and stiffness between blocks and treatments at the 5% significance level using the MATLAB statistics toolbox package (49).

3.4 Results

As expected, the model was able to deduce an intended trajectory and all subjects were able to use this estimate of their intent to perform targeted reaching while experiencing turbulent forces (Figure 5). There were no obvious changes in performance using intent over time. We also observe no after-effects of either IF or the forces. In block 3 where IF replaced visual feedback of the hand, the intent significantly ($p = 0.02$) outperformed the hand itself. For this intent estimate to be useful, subjects should perform better when using IF than when using feedback of the hand. Comparing the hand's performance in block 2 (Figure 6, panel B) to IF's performance in block 3, subjects performed significantly ($p = 0.02$) better using their estimated intent. While subjects also performed better using IF in block 3 than using their hand in block 4, this difference was not significant ($p = 0.07$).

Interestingly, our measure of effective stiffness changed dramatically across the experimental conditions. Subjects stiffened significantly ($p < 0.001$) in response to white noise forces (compare stiffness in blocks 1 and 2 of Figure 7, panel B). Next, subjects' effective stiffness significantly ($p = 0.01$) decreased when IF replaced the hand location as their cursor (compare blocks 2 and 3 in Figure 7). This decrease did not return stiffness to undisturbed levels, and it remained even after feedback of the hand resumed (comparing blocks 3 and 4). Subjects appear to have adjusted their stiffness in response to the stiffness used by IF. Subject's stiffness's did not significantly ($p = 0.1$) differ from the stiffness of the standard model when it was used by IF. Finally, subjects reported that this IF treatment of the cursor feedback made it easier and they felt more relaxed. Detailed statistics of these differences are summarized in Table III.

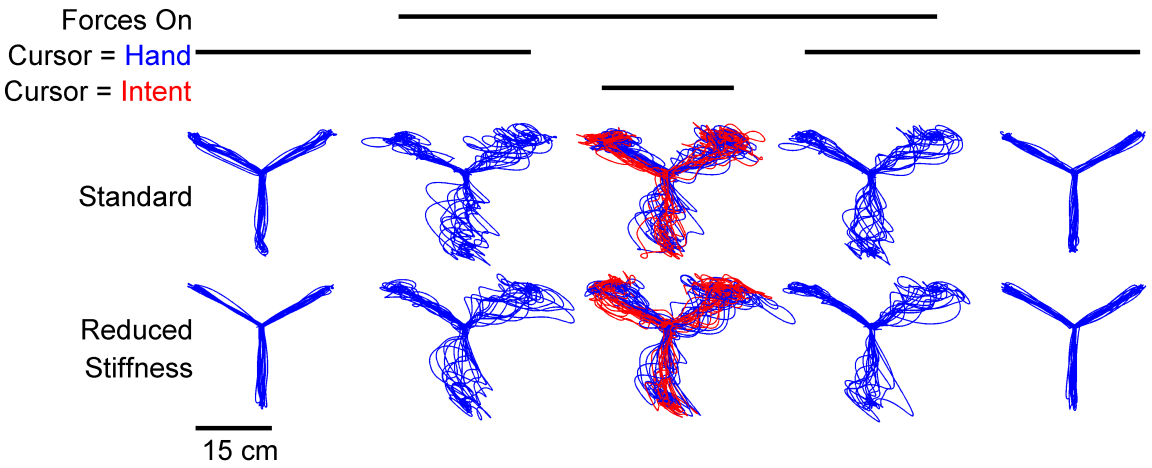

Figure 5. Typical subjects (one from each experiment) made center-out targeted reaching motions under experimentally varied force and feedback conditions. Subjects used feedback of either hand motions (blue lines) or estimated intent (red lines) to complete these reaches. Shown also are the measured hand motions in the third block, which were recorded even though they were not visible to the subject (blue) to be compared to intent (red). Intent was estimated using either the standard stiffness model of (6) or a reduced stiffness model to explore any dependence of reaching stiffness or accuracy on this assumption. The white noise force disturbance was designed to be unpredictable in order to minimize any effect of learning.

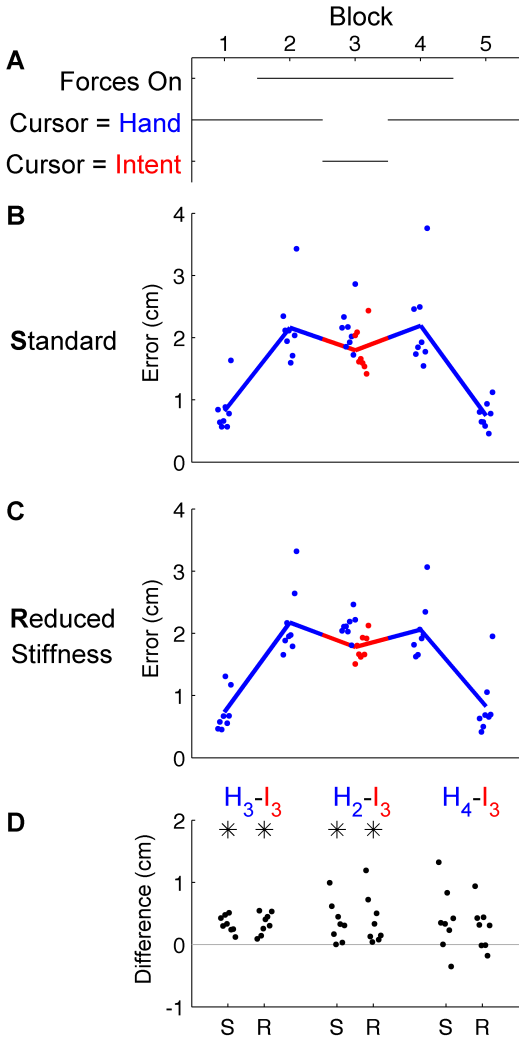

Figure 6. Subjects' reaching accuracy depended on the presence of force disturbance and the contents of visual feedback (**A**). (**B** and **C**) Maximum deviation from straight-line reaching calculated during the first 250 milliseconds after the onset of movement revealed that turbulent force disturbance degraded reaching performance. Comparison across feedback modalities revealed that IF (red) in block 3 alleviated performance error relative to hand performance (blue) in blocks 2 and 4. Note that in block 3 we show blue dots indicating the hand's performance, although it was not visible to the subject. (**D**) Several comparisons showing pairwise performance differences amongst blocks 2, 3, and 4. Comparisons between hand performance in block N and IF performance in block 3 are abbreviated as $H_N - I_3$ (asterisks denote t-test significance at $\alpha = 0.05$ level). Performance did not appear to depend on the choice between the standard stiffness model of (6) (panel B, and "S" labels in **D**) or a reduced stiffness (panel C, and "R" labels in **D**) to determine intent.

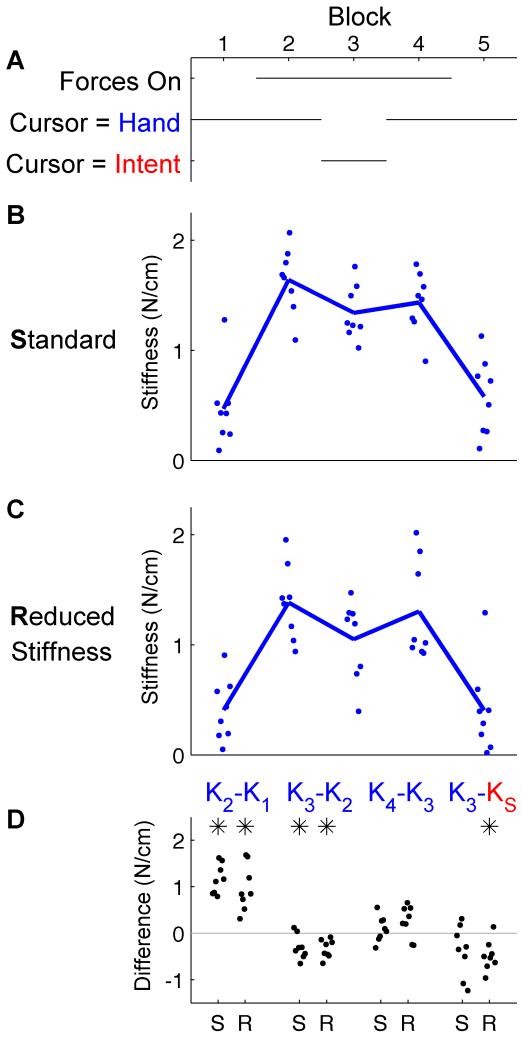

Figure 7. Subjects' effective stiffness depended on the presence of force disturbance and the contents of visual feedback (**A**). (**B** and **C**) Effective stiffness, K, calculated by linear regression during the first 250 milliseconds after the onset of movement revealed that turbulent force disturbance increased this stiffness. (**D**) Comparisons between treatment conditions revealed that exposure to turbulent forces caused significant stiffening, but IF could significantly alleviate arm stiffness. As in Figure 6, the estimated arm stiffness in block 3 significantly depended on our choice of either the classic stiffness model of (6) (panel B, "S" labels) or a reduced stiffness (panel C, "R" labels) to determine intent. Significant differences were determined by paired t-test at the $\alpha = 0.05$ significance level and is denoted by an asterisk.

As we found evidence of IF dramatically reducing effective arm stiffness, we next examined whether we could reduce effective arm stiffness even more. We asked a second group of subjects to perform the same experiment, except that the arm model used by IF was adjusted to assume a greatly reduced stiffness of the arm. Surprisingly, we found the same performance benefits (Figure 6, panel C), but effective stiffness reduced even more (Figure 7, panel C). Performance using IF in block 3 again significantly outperformed the hand in block 3 ($p < 0.001$) and the hand in block 2 ($p = 0.02$). We found the same beneficial effect relative to block 4, but this difference was not significant ($p = 0.06$). Moreover, this reduced stiffness IF allowed subjects to relax to approximately two-thirds the effective stiffness of (6) during block 3 ($p = 0.004$). As before, stiffness increased following exposure to white noise forces ($p < 0.001$) that was significantly ($p = 0.01$) alleviated by IF. Removal of IF was still associated with an increase in arm stiffness, but this increase was not significant ($p = 0.08$). Subjects who received this lower stiffness IF also reported that IF was easier and allowed them to greatly relax. In other words, subjects reduced arm stiffness to accommodate the model.

3.5 Discussion

The work presented here highlights the use of a novel visual distortion of the cursor that leads to superior performance in a hand-eye coordination task in the presence of random disturbances. This real-time distortion marks the estimated intent of the subject rather than the hand location in order to make movements easier. Although exposure to random forces hindered subjects reaching accuracy and increased their arm stiffness, replacing the veridical feedback with intent feedback (IF) improved accuracy and decreased stiffness. While other

	$T(7)$	P	Mean	SEM
Error Comparisons, Experiment 1: Standard Stiffness			cm	cm
Hand (Block 3) - Intent (Block 3)	7.11	<0.01	0.33	0.05
Hand (Block 2) - Intent (Block 3)	3.00	0.02	0.36	0.13
Hand (Block 4) - Intent (Block 3)	2.15	0.07	0.40	0.15
Error Comparisons, Experiment 2: Reduced Stiffness			cm	cm
Hand (Block 3) - Intent (Block 3)	5.69	<0.01	0.34	0.06
Hand (Block 2) - Intent (Block 3)	2.80	0.03	0.40	0.15
Hand (Block 4) - Intent (Block 3)	2.26	0.06	0.28	0.13
Stiffness Comparisons, Experiment 1: Standard Stiffness			N/cm	N/cm
Hand Stiffness (Block 2) - Hand Stiffness (Block 1)	10.2	<0.01	1.17	0.12
Hand Stiffness (Block 3) - Hand Stiffness (Block 2)	-3.24	0.01	-0.30	0.10
Hand Stiffness (Block 4) - Hand Stiffness (Block 3)	0.98	0.36	0.09	0.10
Hand Stiffness (Block 3) - Model Stiffness (K_{P1})	-1.92	0.10	-0.37	0.21
Stiffness Comparisons, Experiment 2: Reduced Stiffness			N/cm	N/cm
Hand Stiffness (Block 2) - Hand Stiffness (Block 1)	5.53	<0.01	0.97	0.19
Hand Stiffness (Block 3) - Hand Stiffness (Block 2)	-4.75	<0.01	0.33	0.07
Hand Stiffness (Block 4) - Hand Stiffness (Block 3)	2.03	0.08	0.25	0.13
Hand Stiffness (Block 3) - Model Stiffness (K_{P2})	3.69	<0.01	0.43	0.12

TABLE III

ERROR AND STIFFNESS CHANGE WITH FEEDBACK TYPE AND PRESENCE OF DISTURBANCE.

visual distortions typically degrade performance and require an adaptation period to overcome, IF immediately enhanced performance. This type of feedback may be a new method for enhancing performance in human-machine interactions, and also sheds light on how the nervous system uses visual feedback.

The most striking result is that although the nervous system sees an untruth about where the end-effector is, it appears to be effective for improving performance. The IF presentation is one of many visuomotor discrepancies in hand-eye coordination tasks, yet this one does not degrade performance and does not require adaptation. Not all perceptual lies appear to be unwanted. Two possibilities explain this result: either the central nervous system was able to adapt to this new feedback within a single reach, or the means to make use of this signal were already available. For instance, IF may mimic efference copy. The performance variability inherent in white noise force disturbances complicated our observations of the learning process, but simple examination showed there were no obvious differences in performance between the first and final exposures to IF. There were also no obvious after-effects from exposure to IF (Figure 5, rightmost panels). The simplest explanation for this is that IF approximates a signal already known to the brain: the path planned for the hand. In addition to its promise in performance enhancement, IF represents a novel means of revealing and studying the mechanisms of motor planning and motor control.

While IF alleviated the increased stiffness caused by exposure to random forces, stiffness remained significantly above baseline levels. Many explanations are reasonable. In particular, we hypothesized that subjects would adapt their own arm stiffness to decrease conflict with

the stiffness model used to estimate their intent and thereby increase the accuracy of the estimate, and while the data did support this conclusion the effect was not strong. Alternatively, inaccuracy and incompleteness of the simple models used might have resulted in an less accurate estimate of intent. Since noise and performance inaccuracy can both lead to co-contraction, this may account for the residual stiffness. Finally, as subjects were not cued regarding the onset or removal of IF, the residual co-contraction may have been a precaution against the resumption of veridical feedback.

The ease with which subjects could make use of their estimated intent provides strong preliminary evidence that a specific intended trajectory was computed for the hand even when reaching in a highly variable environment. While recent work has identified kinematic constraints unnecessary for a task (50), this is the first direct evidence that the entire trajectory is controlled even in the absence of specific instructions or constraints. While portions of the intended trajectory are surely computed before the onset of movement, movement intent is not finalized before the onset of movement and is not strictly ballistic. There is mounting evidence of multiple corrective actions formulated after movement onset in recent literature. For example, Mirabella (51) compared onset and movement times during a task that could be countermanded, and found that quicker onset times were compensated by longer movement times, most likely due to the need for on-line proactive adjustments in anticipation of known task demands. Such a context effect has also been recently replicated on Parkinsons patients (52).

It is important to distinguish the signal we are exploring from others that may use the word "intent." One source of complexity is that formation of actions is a multi-step process in which several brain regions contribute. The intent we refer to here is not the goal-oriented intent that might temporally precede the computation of the motor plan (53; 42). The intended trajectory and goal remain malleable and can change (or be suppressed) even after execution (54). Work by (52) showed that inhibition is important for quickly aborting, interrupting, and re-planning motion after its onset. In contrast, our version of intent focuses on the final stage of planning and hence the command at the present-time. This is after the nervous system has completed any further checks and released the plan. In a sense, the intent we provide is in the present and not pre-planned (for the future) or post-processed *a posteriori* (in the past).

It is also important to distinguish this from other methods that attempt to determine the ultimate target of action (the goal). A number of human-computer approaches strive to identify, for example, the final target of a movement (55). In our rather limited task with only three possible reaching targets, identification was trivial. We analyzed this on our task and were able to easily identify the target with 95% visuomotor efficiency (closely related to accuracy (56)) within the first 80 milliseconds of motion, regardless of whether we used the hand or IF. In contrast to these target-prediction methods, our approach allowed for the *instantaneous* determination of the intended hand location. The novelty of such instantaneous detection created the prospect of real-time feedback in human-machine interactions.

Intent Feedback might facilitate human-machine collaboration and artificial performance augmentation by enabling the machine to preserve an operator's intent while canceling un-

expected disturbances from the environment. This should reduce the demand on the human operator and increase performance – especially in environments with rapidly changing conditions. This assistance goes beyond environment cancellation by also accounting for any errors the operator might make based on their expectation of disturbance.

While IF holds promise, it also has strong limitations. IF is entirely dependent on the accuracy of the models used. While we were able to leverage measurements of cadaver anthropometrics, average tendencies do not capture individual variability. Similarly, the model of (6) appears to have accurately captured the mean tendencies (Figure 6, panel D's rightmost comparison) without accounting for variation among individuals or variation over time. Techniques that could estimate changing stiffness, perhaps even in real-time, would greatly increase the accuracy and utility of IF. While IF is limited by the accuracy of the models used, many candidate models are available and may outperform the simple model we investigated here in this preliminary study.

More broadly, this IF approach may be useful in any situation where some model of the dynamics is available and disturbances can be measured. For example, brain-computer interfaces may need to address measurable common-mode electrical artifacts, such as the electromagnetic disturbances that occur from lights being turned on. In cases where a disturbance can be measured due to its similar effect across all sensors, IF allows the interface to respond in a manner congruent with the user's intent. This is especially important when simply canceling the disturbance is not sufficient or practical, such as an exoskeleton's user intending to crouch down during an earthquake rather than trying to remain upright.

In any case, intent feedback led to performance benefits for subjects moving in a changing, uncertain environment. In addition to increasing subjects' accuracy, IF may have allowed subjects to reach their goals with less effort as arm stiffness decreased. IF provides a novel form of feedback that may facilitate new insights into the nature of motor control and allows a machine to collaborate more effectively with a human user.

Disclosure/Conflict-of-Interest Statement

The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Author Contributions

All authors designed the work, analyzed the data, drafted and revised the manuscript, approved the final version, and agree to be held accountable for all aspects. TM and CM collected the data.

Acknowledgments

The authors wish to thank Yazan Abdel Majeed and Adelina Voukadinova for their review of the preliminary manuscript. We also thank the community of the Robotics Lab at the Rehabilitation Institute of Chicago for critical commentary during the genesis of this paper and the work leading up to it. This work was funded By NIH R01-NS053606 and NIDRR H133E120010.

CHAPTER 4

DECOMPOSING INTENT INTO SUBINTENTS

??

4.1 Overview

This work, written in collaboration with my adviser and a lab mate, was submitted to the peer-reviewed International Conference of the IEEE Engineering in Medicine and Biology Society on March 14th, 2016. All of the previous works grew out of a simple question: if humans modify their intended path mid-reach in response to force disturbances, when do they do so? This simple question required two pieces of technology that did not exist. First, we needed the capacity to recover the intended trajectory even though it is obscured by the force disturbance. Second, we needed a tool that could determine when if and intent had changed. While the literature addressing reaching during infant development and following stroke made clear that motion was composed of more elemental subintents, it did not make clear was how to recover those subintents from healthy adult reaches in a convincing way. This paper makes use of our intent determination and large, unexpected forces to cause healthy reaching to break down and expose its makeup. We then characterize that makeup to obtain preliminary insights into the mechanisms of intent production.

4.2 Introduction

Since the outset of human reaching research, it has been speculated that healthy, adult reaching is composed of subintents (14). These subintents are distinct and easily recognizable during infant development, but gradually blend together until their placement is mostly indistinguishable by adulthood (15). Following a stroke, these subintents are again readily recognizable (57) and their degree of blending reveals information about the severity of stroke and progress towards recovery (16). Neural correlates of subintents have also been identified (58). Attempts to identify and characterize subintents in healthy adults are thwarted by the subintents' shape: a sum of any variety of *radial basis functions* can be used to reproduce any shape of interest (59). Attempts to identify subintents in undisturbed reaching are therefore "doomed to succeed" and cannot support a rigorous and falsifiable hypothesis.

In order to claim that subintents have been accurately identified, we must show that they can be accurately recovered from known sources. One test would be to show accurate recovery of subintents from synthetic motions whose composition is known to us. As subintents' shape is not known, we must additionally show accurate recovery even when we intentionally differ the shape of the subintents in the composition from the shape of the subintents in the decomposition.

Here, we seek to address these problems using three important techniques. First, we use very large forces to disturb the limb and allow for easy identification of distinct corrections. Second, because these forces are so large that they obscure the person's contribution to their own reach, we use a technique that can isolate that contribution – *intent determination* (1; 2). This algorithm recovers the intended trajectory of a movement when the actual trajectory differs from

intent due to external forces. This approach reveals how and when subjects modify their movement intent. Third, we cluster data into subintents using the directional component of velocity rather than seeking to maximize a goodness of fit that explicitly depends on subintents' shape. To support the subunit composition hypothesis, we should recover either a consistent pattern of composition or a specific distribution underlying the stochastic composition we observe.

4.3 Methods

As it is not known when or how subintents are selected, it was not obvious what types of disturbances might produce especially unblended intents. We designed an unusually large disturbance consisting of a strong curl field followed by a pulse of force. We hoped that the combination of these unexpected, strong forces with an abrupt, time-dependent transition would interfere with learning (60) and instead cause subjects to adjust their intent. We then recovered that intent (1) and decomposed it into subintents. Next, we verified the accuracy of the decomposition. Finally, we examined the statistical properties of the subintents we recovered in order to make some inferences about how they are planned.

4.3.1 Intent Determination

Intent determination is a filtering algorithm that combines an arm model with measurements of the actual trajectory of the arm and the forces exerted on the arm to estimate the movement intent (1) (Figure 8). First, the physical length of the upper arm and forearm are measured *in situ* and the subject reports their body mass. Second, these physical properties are placed into a model of the arm constructed from cadaver data (17; 18) and experimental observations (4) to estimate the arm's impedances. Third, the arm model is solved for the intended acceleration

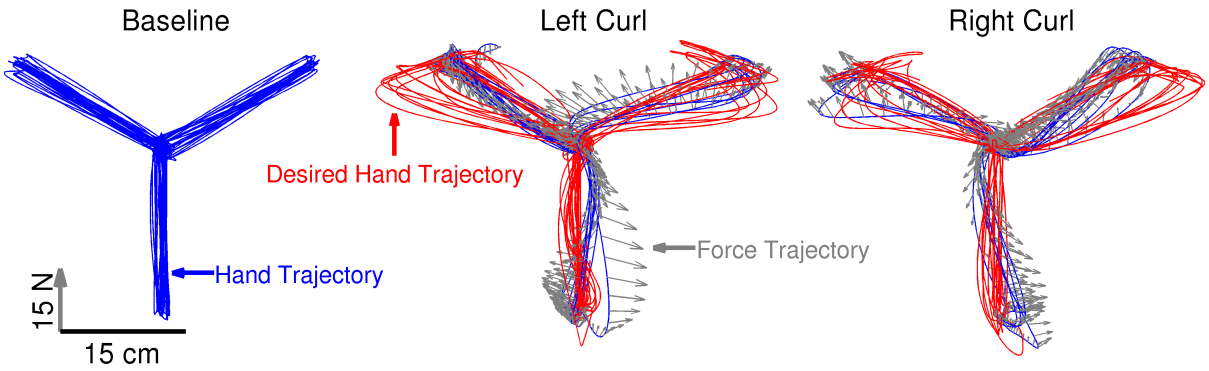

Figure 8. Using intent determination, an algorithm that recovers intent from force-disturbed movements (1), we were to expose subjects to strong, abrupt forces and observe abrupt corrections at harsh angles. This directional separation facilitated decomposition of those intents into subintents.

and integrated numerically to find the intent trajectory: a time series of desired positions. These steps and parameters differ from our earlier work only in that we did not take extra steps in order to customize the stiffness model.

4.3.2 Subintent Modeling

We define subintents as minimum-jerk, bell-shaped speed profiles based on work showing that they best fit human movement data among symmetric subunit shapes (61). The intent trajectory, $\vec{y}$, is then a sum of scaled, translated versions of this profile κ,

$$\vec{y}(t) = \vec{\epsilon} + \vec{y}(0) + \sum_n \vec{L}_n \kappa(t, C_n, S_n) \tag{4.1}$$

$$\kappa(t, C, S) = 10\tau^3 - 15\tau^4 + 6\tau^5 \tag{4.2}$$

$$\tau = \begin{cases} 0 & \text{if } 0 > \frac{t-C}{S} + \frac{1}{2} \\ \frac{t-C}{S} + \frac{1}{2} & \text{if } 0 \leq \frac{t-C}{S} + \frac{1}{2} \leq 1 \\ 1 & \text{if } 1 < \frac{t-C}{S} + \frac{1}{2} \end{cases} \tag{4.3}$$

where t is time, C_n and S_n are the center and duration in time of the nth subunit, $\vec{L}_n$ is the vector change in position due to the nth subunit, and $\vec{\epsilon}$ combines all factors not explained by subintents such as the inaccuracies due to imperfect recording and intent determination. The cases used to define τ have the effect of limiting the contribution of a subunit such that it adds smoothly to the trajectory within its span and has a constant contribution thereafter. In terms of velocity,

$$\dot{\vec{y}}(t) = \dot{\vec{\epsilon}} + \sum_n \vec{L}_n \dot{\kappa}(t, C_n, S_n) \tag{4.4}$$

$$\dot{\kappa}(t, C, S) = (30\tau^2 - 60\tau^3 + 30\tau^4)S^{-1} \tag{4.5}$$

the contribution of a subunit is dramatically simpler as it makes no contribution at all outside its span. It is for this reason that we perform our decomposition in the velocity space.

4.3.3 Subintent Decomposition

A decomposition method must overcome many challenges. The κ component can itself be further decomposed into other radial basis functions. Worse, the decomposition can be dependent on the shape of the basis function used. Others have tried to overcome these challenges by invoking *Occam's Razor*: minimizing the number of subintents used to decompose a movement also minimizes the number of free parameters. They optimize the parameters of n subintents and then increase n if a fit accuracy threshold has not yet been attained (59). Alternatively, some groups have instead stopped at a point of diminishing returns (62). This is not a falsifiable hypothesis as no threshold for n is obvious and quality of fit might be driven by factors other than an insufficient number of subintents. In particular, it is not known how many subintents might compose a typical motion. Subintents may also be very small as evidenced by miniature art: humans are capable of freehand writing on a single grain of rice. Instead of asking how many subintents we need to fit a movement, we instead examine how many directions are evident within a movement. Leveraging the distinctness of the subintents' directions allows us to avoid concerns of spurious decomposition without explicitly minimizing any cost function.

A decomposition borne of subunit distinctness is straightforward: if subintents are distinct then their direction at peak speed is their direction. We need only identify peak speeds, record their directions, and then use the dot product as an expedient means to identify when subintents

start and end. This process is summarized graphically in Figure 9. Stated as a stepwise algorithm that recovers a single subunit:

1. Identify the time at which the maximum speed occurs within a movement, C.
2. Note the movement direction at this time, $\hat{v}_p = \frac{\vec{v}(C)}{|\vec{v}(C)|}$.
3. Construct a rate of change that depends only on direction, $\psi = \frac{d}{dt}\frac{\vec{v}(C)\cdot\vec{v}(t)}{|\vec{v}(t)|}$.
4. Take S as 2 times the distance from C to the nearest local maximum or minimum of ψ.
5. $\vec{L} = 1.875^{-1}S\vec{v}(C)$ because $\dot{\kappa}(t = C) = 1.875S^{-1}$.
6. Subtract $\vec{L}\kappa(t, C, S)$ from the velocity trajectory.
7. Repeat the algorithm on the remainder to recover additional subintents as desired.

Our algorithm terminated once the largest remaining peak in speed did not exceed ten centimeters per second. These steps rely on the fundamentally restrictive assumption that the direction observed at peaks in speed reflects a single subunit and not a sum of subintents. Naively, this assumption is quite poor as it presumes a structure to movement composition that is not founded in previous observations; however, our hypothesized subunit distinctness appears to be upheld. Figure 11 (blue bars) demonstrates that this process does accurately recover subunit parameters from synthetic data.

4.3.4 Synthetic Data

One unique way to test our approach is to artificially construct a motion from known submovements, and then see if our approach can recover the original elements. As the subintent composition of movement is not known, we construct two best estimates: one which adheres

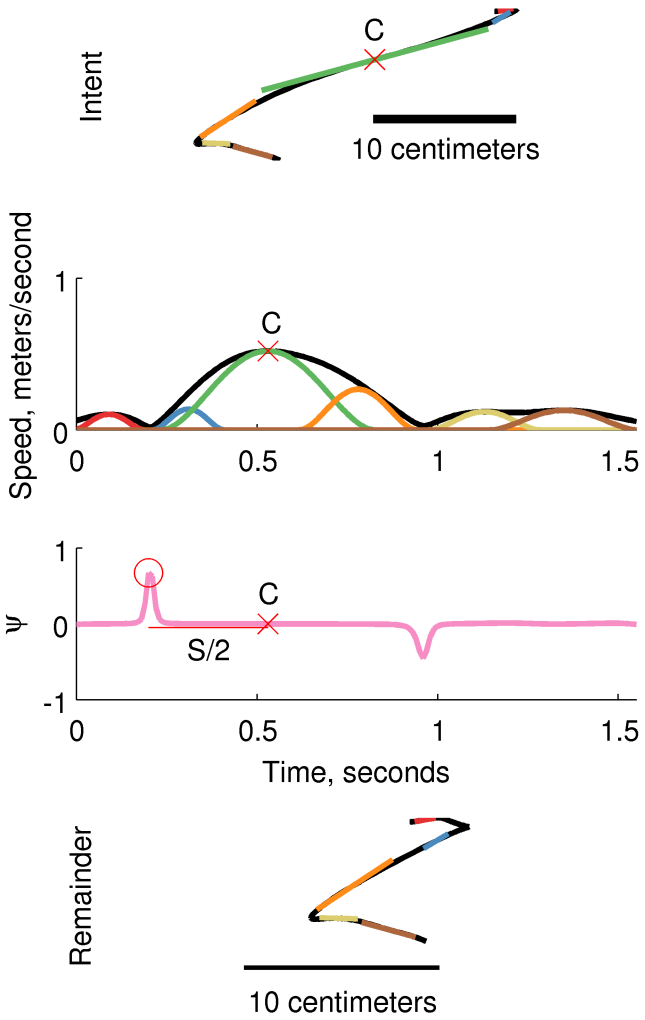

Figure 9. To demonstrate our decomposition method, we begin with a recorded intent (black line, intent panel) and its speed trajectory (black line, speed panel). We have additionally marked pastel-colored subintents in the intent, speed, and remainder panels to illustrate typical properties. The largest subintent's peak, C, is identified and marked with a red X. The time derivative of velocity's dot product with the direction at C is computed as ψ. Submovement duration, S, is twice the distance in time from the red X to the red circle that marks a maximum in the rate of change of direction. The velocity at C and the subintent's duration also define its magnitude L. We can then calculate the subintent's contributions to position and velocity and subtract them from the intent trajectory.

to our assumptions regarding subintent shape and one which does not. Subject 1's recovered movement intent was decomposed by the algorithm described above, without any refinement of S, and recomposed according to equation Equation 4.1. This recomposition was our naive best estimate of the underlying subunit composition of healthy adult reaching. We additionally recompose movement using the sin function as an alternative basis in order to test the robustness of our extraction to our assumption that submovements are minimum jerk in shape. In this case $\dot{\kappa} = \frac{\pi}{2S}\sin(\pi\tau)$.

4.3.5 Experimental Design

Subjects made center-out reaching motions in three directions that were sometimes disturbed by a large and difficult-to-learn *curl-kick* force. This disturbance was time-varying,

$$\vec{F} = \begin{cases} \pm \begin{bmatrix} 0 & 50 \\ -50 & 0 \end{bmatrix} \vec{v} & \text{if } t_0 \leq t \leq t_0 + 400ms \\ 15\frac{\vec{x}_f - \vec{x}}{|\vec{x}_f - \vec{x}|} & \text{if } t_0 + 400ms \leq t \leq t_0 + 800ms \\ 0 & \text{elsewhere} \end{cases} \tag{4.6}$$

and depended on velocity $\vec{v}$, position $\vec{x}$, and the position of the reach's target $\vec{x}_f$. The direction of the curl field is pseudo-random, which makes it very difficult to correct for even if it is detected early. The kick portion was programmed to push the subject towards the target to induce replanning. The onset of movement, t_0, was detected as the moment the subject left the previous reach's target. These forces were not present during the first 100 reaches so subjects

could familiarize themselves with the robot and any adjustments could be made. After the hundredth reach, disturbances were presented pseudo-randomly such that each direction was disturbed fifteen times. A total of 810 reaches were performed.

4.3.6 Human Subjects

Our human reaching trajectories were produced by eight healthy adults, two females and six males, aged 21 to 34 years old who gave informed consent in accordance with Northwestern University Institutional Review Board, which specifically approved this study and follows the principles expressed in the Declaration of Helsinki. Subjects self-reported right-handedness, performed the reaches with their right arm, and were compensated for their time. Subjects' arm segment lengths were directly measured *in situ* while body mass was self-reported.

4.3.7 Apparatus

Subjects held the handle of the planar manipulandum depicted in Figure 10, which was programmed to compensate and minimize any friction or mass. The MATLAB XPC-TARGET package (49) was used to render the force force environment described above and collect data at 1000 Hz. Visual feedback of target and hand locations was performed at 60 Hz using OpenGL. Force sensor drift was addressed using a linear re-zeroing procedure.

4.3.8 Statistics

In order to test for bias in our parameter recovery, we performed a Student's T-Test to determine whether or not the mean of the difference of the known composition's parameters and the decomposition's parameters differed significantly from zero. Tests were conducted at the $\alpha = 0.05$ significance level.

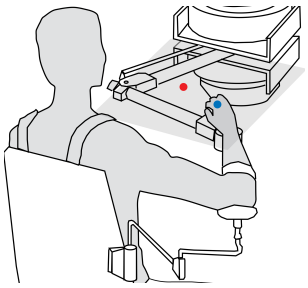

Figure 10. Subjects held the handle of a robotic device that rendered programmed forces and recorded position and force information. Position was indicated to the subject with a blue cursor and the target of a reach was indicated in red. The opaque, horizontal screen occluded subjects' vision of their hand.

The cumulative distribution function of the exponential distribution is defined as the probability that the random variable A takes on a value less than or equal to a: $P(A \leq a) = 1 - e^{-\lambda a}$. Therefore we expect that if an empirical cumulative distribution, $C(a)$, is generated by this distribution, $\ln(1 - C)$ should be linearly related to a as $\ln(1 - C) = -\lambda a$. We examine this linearity using the coefficient of determination, R^2.

4.4 Results

Parameter recovery from synthetic data was highly accurate. Student's T-Test detected no significant bias ($n = 1352$, $p > 0.49$) between the parameters recovered and their known values. While peak location and position change had lower variance, duration and subintent count were detected with much higher variance. Decomposition accuracy did decrease somewhat when sin-based subintents were used to construct synthetic movement even though minimum jerk-based subintents were still used in the decomposition, but this decrease in accuracy was not significant ($n = 1352$, $p > 0.09$).

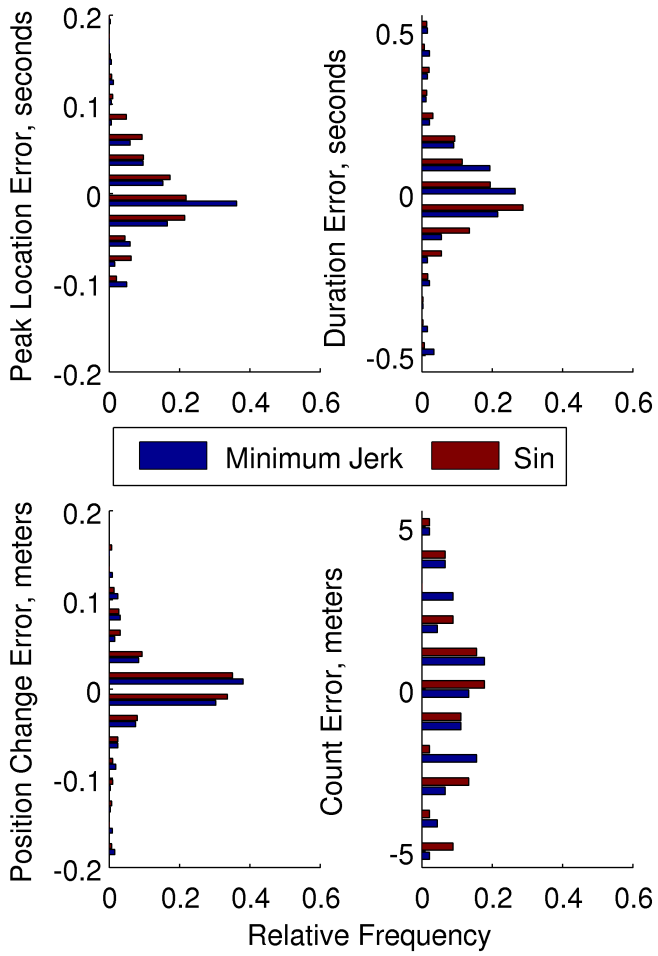

Figure 11. Error in recovered subintent properties was not dependent on the subintent shape used to compose synthetic data. Parameter recovery was accurate, despite a high variance, as evidenced by mean error not significantly differing from zero as detected by Student's T-Test.

We failed to identify any structure or motifs in our subintent decompositions. Lacking evidence of a recurring structure, we instead sought to identify the statistical properties of subintents. The first pattern we noticed was that our residuals, which have units of velocity squared, appeared to be exponentially distributed ($R^2 = 0.94$, Figure 12 upper-left plot). One explanation for this observation is that the largest remaining subintent accounts almost entirely for the residual and the subintents themselves are exponentially distributed in velocity squared. This hypothesis is supported as subintent peak kinetic energy, $1.875^2 L^2 S^{-2}$, is indeed well-accounted for by an exponential distribution ($R^2 = 0.95$, Figure 12 upper-right plot). Seeking other exponentially distributed quantities, we also discovered that the duration between subintent peaks and the number of subintents composing a movement also appear to be exponentially distributed ($R^2 = 0.91$ and $R^2 = 0.93$ respectively, Figure 12 lower plots).

4.5 Discussion

This study provided support for a novel approach to decompose measured actions into subintents in an improved manner that considers the underlying intent and does not depend on assumptions of shape or number of elements. By capitalizing on a new intent extraction technique (1) and exposing healthy adults to strong and lasting forces, midmovement replanning became obvious. We found that our method could accurately recover subintent parameters and that this accuracy was robust even when an incorrect subintent shape was assumed. This method was not necessarily precise, however. While this imprecision is problematic, future work might improve upon it it by using our unbiased estimate of subunit timing and duration to recover subintents' shape empirically by averaging a sufficient number of subintents.

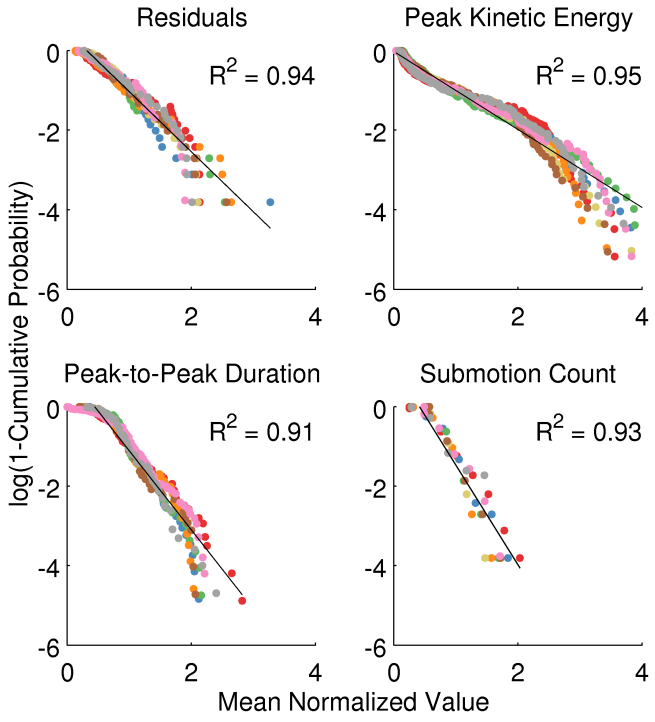

Figure 12. After discovering that the residuals of our decomposition appeared exponentially distributed in velocity squared (upper-left), we tested other quantities for this distribution and found that peak kinetic energy, subintent count, and the duration in between successive subintent peaks was also well-explained by an exponential distribution. Colors represent different subjects while dots represent individual subintents or reaches. Quantities were normalized by their mean, which causes all regressions to have the same slope, in order to facilitate comparison across subjects.

Our finding that subintent counts and peak kinetic energy were exponentially distributed means that subintent count will occasionally be very high and submovement magnitude will occasionally be very small. This contrasts with previous work that explicitly avoids finding small subintents (59; 62). Cost functions are usually constructed such that they maximize quality of fit while minimizing the number of free parameters used to do so. Based on the statistical distributions in our findings, we recommend that model likelihood be maximized instead.

Finally, we found an interesting distribution of subintent parameters: most were exponentially distributed. Interestingly, such distribution tendencies arise naturally as the *maximum entropy* distribution supporting only positive numbers for a given mean. In physics, these are termed Boltzmann or Gibbs Distributions and they describe many important phenomena such as the motions of classical gases and the separation of charged particles by a semi-permeable membrane. Since the statistical properties we observe for subintents are so similar to the statistical properties of particles in the physics of gases, this suggests the enticing possibility that the same tools used to derive and understand the statistical mechanics of gases will also apply to human reaching motions.

Acknowledgments

The authors wish to thank Theresa Horowitz for her review of the preliminary analysis and manuscript. We also thank the community of the Robotics Lab at the Rehabilitation Institute of Chicago for critical commentary during the genesis of this paper and the work leading up to it.

CHAPTER 5

A STATISTIC MECHANICS OF SUBINTENTS

??

5.1 Overview

This work, written in collaboration with my adviser and a lab mate, will be submitted to Nature as a presubmission inquiry. The decompositions of the previous work were completed years ago and presented at a conference, but abandoned for a time as they had no obvious value. During the time that the Patton lab was beginning to research an unconstrained movement paradigm nicknamed "free exploration" (63), I was enrolled in a nanotechnology course that included quantum and particle physics components. By luck, I noticed that hand speed during this paradigm appeared to have the same statistics as noble gases. We noticed several other superficial similarities and then turned to our then-abandoned subintent decomposition results. Indeed, individual subintents behaved like the non-interacting particles of a noble gas. Though we will derive an algorithm that explains and predicts human motion, little here is truly original. Instead, this is a product of happenstance enabled by standing on the shoulders of giants.

5.2 Introduction

Although a body of evidence has shown that reaching movements are comprised from separate subintents, conclusions regarding how many subintents compose a movement, how they might be distributed within the movement, or how they are chosen have been elusive (59).

Subintents are distinct during childhood (15), but coalesce over the course of development into the typical single, smooth, bell-shaped speed profile observed in healthy adults (14). Brain injuries such as stroke may cause movement to regress back into separate, jerky subintents (57) that again coalesce into a single smooth unit during recovery (16). One particularly important confound to exploring the properties of subintents is that there is no unique decomposition, since there are potentially multiple valid ways to decompose any individual movement (59). One novel strategy for overcoming this problem might be to turn to the statistics of repetitive movements.

We took a three-pronged approach to articulating and testing this subunit planning hypothesis: movement is a series of overlapping subintents selected by a winner-takes-all process that statistically pursues reward and avoids cost. First, we use energy balance and other first principles to model how a system might behave if influenced by three quantities associated with the transition to a new state: reward gained (64), rate of reward (65), and cost-to-go (10). Second, we tested specific predictions of subintents from human reaches that were disturbed by robotic forces (66). Third, we tested specific predictions regarding the statistics of whole-motion peak speed across three different experiments. These statistical findings support our hypothesized model of arm motion composition and planning.

5.3 Results

As an overview of what is to come, we begin with the simple assumption that subintents exist and have a stochastic component. Consequently, we want to explore their statistical properties, make predictions from first principles, and make use of energy arguments to ground

our findings in the laws of physics. Once we assume that the system is Markov (memoryless) we learn that this implies a potential energy component. Further, this potential component allows us to explore distributions as a tool for understanding subintents. Because of this, we show that it is possible to develop conditional relations describing state transitions, leading to a hypothesis about the characteristics of distance and duration in subintents. We are also able to make predictions about the how submovements are chained together, and can hypothesize a specific distribution for the number of subintents that might make up a discreet action. Finally, we make statistical predictions on the peak speed of whole, undisturbed motions by combining our previous insights. Using these tools, predicting the statistical tendencies of any given motion reduces to simple sampling from the distributions we have obtained. The sections below describe these steps in detail.

5.3.1 Describing actions using the mechanics of energy

We begin by examining how the selection of a particular action might result in changes in energy. The total work balance equation is $\Delta R = \Delta J + \Delta T$, which conserves the total energy among internal energy J, kinetic energy T, and work done by the environment ΔR. We typically think of such balance equations in terms of an instantaneous interplay of energy, but if movement is compromised of subintents we would not benefit from this extra analysis. Instead, we might discretize time and only inspect the start and end of each subunit. Then we could consider only the work across each subunit, the total changes in energy over the course of the subunit, rather than the conservations of energy that occurred throughout. Note that this aspect reflects the fact that all organisms might invest energy temporarily for a later return.

Interestingly, if we assume that subintents make no contribution to motion at their endpoints, the ΔT term vanishes and we can simplify the conservation to $\Delta R = \Delta J$. Internal energy now entirely describes the net uptake in energy.

While the principles of muscle energetics are well-understood, measuring them (using oxygen consumption, calorimetry, or other available tools) is noisy and slow relative to the time-scale of most individual actions. In addition, it can give only global measures of whole body energy consumption. Some research has successfully related behavior directly to effort (67; 68; 69). Here, we model effort simply as proportional to the amount of kinetic energy converted to heat (and hence lost) by the end of an subunit. This thermodynamic model is clearly incomplete as the amount of energy consumed by movement depends on many factors such as the resting length and velocity of muscle (70). Nonetheless, we hypothesize that the nervous system uses this type of calculation to plan movement.

In the most abstract sense, we specify a subunit of motion as some function that maps time, t, to state, x, and has a finite span in both: $0 \leq t \leq S$ and $0 \leq x \leq L$. Only during this span of time can velocity and acceleration be non-zero. Using this definition for a subunit, we can then show that ΔJ scales with $L^2 S^{-2}$ when no other energy is exchanged. We begin by defining a *reference subunit*, x_1(t), for which $S = L = 1$. We can then define a version of this function scaled in time and space, $x(\frac{\tau}{S}) = Lx_1(t)$, where $\tau = St$. The chain rule relates how velocity scales with these scaling factors:

$$\frac{dx}{d\tau} = \frac{dx}{dt}\frac{dt}{d\tau} = \frac{L}{S}\frac{dx_1}{dt} \tag{5.1}$$

Energy scales with the square of velocity and thus cost-to-go scales with L^2S^{-2} relative to the cost of $x_1(t)$, Q_1. This homogeneous description of energy consumption is dependent on a coordinate system's inertial reference frame, but is otherwise independent of any choice of coordinate system(s). This allows abstraction of planning to coordinate systems in which more accurate muscle models may be prohibitively complicated to represent.

Next, our model abstracts the concept of potential energy as state appraisal, ΔU, defining it as the portion of change in internal energy not accounted for by the cost-to-go, $\Delta U = \Delta J + L^2S^{-2}Q_1 + Q_0$. $L^2S^{-2}Q_1$ is the cost of transitioning from one state to another while Q_0 corresponds to a cost common to all outcomes. This might include factors like cognitive effort, resting metabolism, or finite opportunity to act. State appraisal is a conserved, state-based change in *inferred* energy. It may not seem intuitive that state appraisal should be a potential-like energy, but chemical energy expended is already accounted for by the cost-to-go and kinetic energy is not present. State-based, conservative energy is all that remains. An organism might have many options for its next behavior at any given point in time, each with an associated net uptake composed of a cost-to-go state transition-defined component and a conservative state-defined component.

5.3.2 First prediction: the distribution of subintent peak kinetic energy

Even the most over-practiced actions have variability (7), and we seek a model of that variability that is purely dependent on internal energy. As change in internal energy depends

only on the previous state and not on the history of states that proceeded it (i.e., Markov), this implies that the distribution of states visited will satisfy an exponential probability of state,

$$p(x) \propto e^{-\beta U(x)} \tag{5.2}$$

where β is a constant. Such a relation is known as a *Gibbs measure* (71), which describes the distribution of states occupied over long spans of time if state appraisal is constant and the environment does not intervene. As neither of these conditions are often satisfied *in vivo*, we take additional steps to relate the Gibbs measure to reaching motions: given that we know the endpoint of the prior subintent, where might we expect the next subintent to go? The probability of transitioning from a given state to any other is expressed as a conditional probability that depends on both cost-go-go and potential (state-based) energy. To relate this conditional probability to the unconditional probability (eq. Equation 5.2), we can leverage Bayes Theorem as a ratio of probabilities,

$$\frac{p(x_n|x_{n-1})}{p(x_{n-1}|x_n)} = \frac{p(x_n)}{p(x_{n-1})} = e^{-\beta\left(U(x_n)-U(x_{n-1})\right)} \tag{5.3}$$

where subscripts denote the sequence of subintent end points. Since the ratio of the conditional probabilities is an exponential distribution, this suggests that the conditional probability should itself take that form,

$$p(x_n|x_{n-1}) \propto e^{\alpha\Delta J} = e^{\alpha(\Delta U - L^2 S^{-2} Q_1 - Q_0)} \tag{5.4}$$

where α is a constant related to β. As the Q_1-based term is statistically independent of the ΔU and Q_0-based terms, this yields our first testable hypothesis: we expect an exponential distribution of subintent distance squared over duration squared. A previous work of ours, which examined subintent decomposition of intents from force-disturbed reaching, does indeed support this prediction ($R^2 = 0.95$) (66) (Figure 12, top right panel).

5.3.3 Second prediction: the distribution of the number of subintents within a goal-directe reach

We next attempt to specify subintent duration utilizing the insight that humans adjust the duration of their actions to maximize their rate of reward (65). The mean rate of reward, $E[\dot{J}]$, is given by

$$E[\dot{J}] = (\Delta U - Q_0)S^{-1} - L^2 S^{-3} Q_1 \tag{5.5}$$

which is maximized when

$$S^{-2} = \frac{\Delta U - Q_0}{3L^2 Q_1} \tag{5.6}$$

This deterministic relationship between change in state appraisal and subintent duration allows us to predict a relationship between two tangible measures: the ratio of the probability of a subintent's start and end points and the duration of a subintent between those points in space. Unfortunately, this ratio of probabilities is only valid at steady state and our sole method of investigating subintent properties requires disturbing movement. Therefore, we take an additional step to relate this prediction to the properties of whole reaches. Consider a reach that starts at some point x_0, ends at some other point x_f such that after N subintents $x_N = x_f$,

and is otherwise not constrained to follow any particular path between those endpoints. The probability of such a reach following a specific path can be written as a Markov chain defined by the subintent endpoints that describe the path,

$$P(x_N|x_{N-1})P(x_{N-1}|x_{N-2})P(x_{N-2}|x_{N-3})...P(x_1|x_0) = \prod e^{\alpha\Delta J} = e^{\alpha\sum\Delta J} \tag{5.7}$$

and then simplified by substituting the expression for S that we hypothesized earlier (Equation 5.6),

$$\sum_{n=1}^{N}\Delta J = \frac{2}{3}(\Delta U_{0N} - NQ_0) \tag{5.8}$$

This result reveals that the probability of a particular sequence of subintents depends only on the total change in state appraisal and the number of subintents comprising the chain. Since the change in state appraisal is roughly constant across all reaches that share the same origin and target, we arrive at a strong prediction regarding the distribution of the number of subintents within a reach: $\ln p(n) - \ln p(n-1) = \frac{2}{3}\alpha Q_0$. In our prior work, we did indeed find that $p(n)$ was well fit by a line on a log scale ($R^2 = 0.93$) (66) (Figure 12, bottom right panel).

5.3.4 Third prediction: the distribution of peak speeds during point-to-point reaching

While the above predictions are consistent with measured subintents, we also wish to make predictions about the statistics of whole, undisturbed movements independent of any need to recover intent or decompose it. Not only are the statistics of whole, undisturbed reaches of much wider interest, identifying the peak speed of a reach is simple and common. We begin

formulating our prediction by noting that fixed start and end points imply a fixed total change in potential. The properties of the individual subintents are therefore no longer statistically independent as their sum is fixed. We cannot simply examine the expected maximum of n unrelated subintents. Rather, we must generate n subintents whose total change in potential reflects this sum,

$$p(\Delta U_n|\Delta U_{0N}) = \Delta U_{0N}\frac{e^{\alpha\Delta U}}{\sum_n e^{\alpha\Delta U_n}} \tag{5.9}$$

By repeating this calculation for many simulated movements, we can generate an empirical estimate of the probability of the maximum change in potential among the subunits composing an N-subunit reach. N must also be generated for each simulated movement using Equation 5.7 and Equation 5.8. Through Equation 5.6 we learned that change in potential is proportional to peak kinetic energy, which allows us to finally fit the cost-to-go terms Q_1 and Q_0 through this relationship. The fitting procedure is straight-forward as simulated change in potential and measured peak kinetic energies share a one-to-one relationship through their cumulative density functions (Figure 13, top panels). After performing this fitting, we examined the quality of the fit (Figure 13, bottom left panel) and its accuracy relative to the Gaussian distribution (Figure 13, bottom right panel). We find across these experiments that fits are generally of good quality ($R^2 = 0.9169 \pm 0.1031$, mean±standard deviation) and significantly outperform the Gaussian distribution (sign test: $n = 60$ subjects, $p < 0.0001$).

5.4 Discussion

If human motion indeed arises from the process model we describe, human motion planning is elegantly simple. Every parameter can be encoded using only an interspike interval. Every

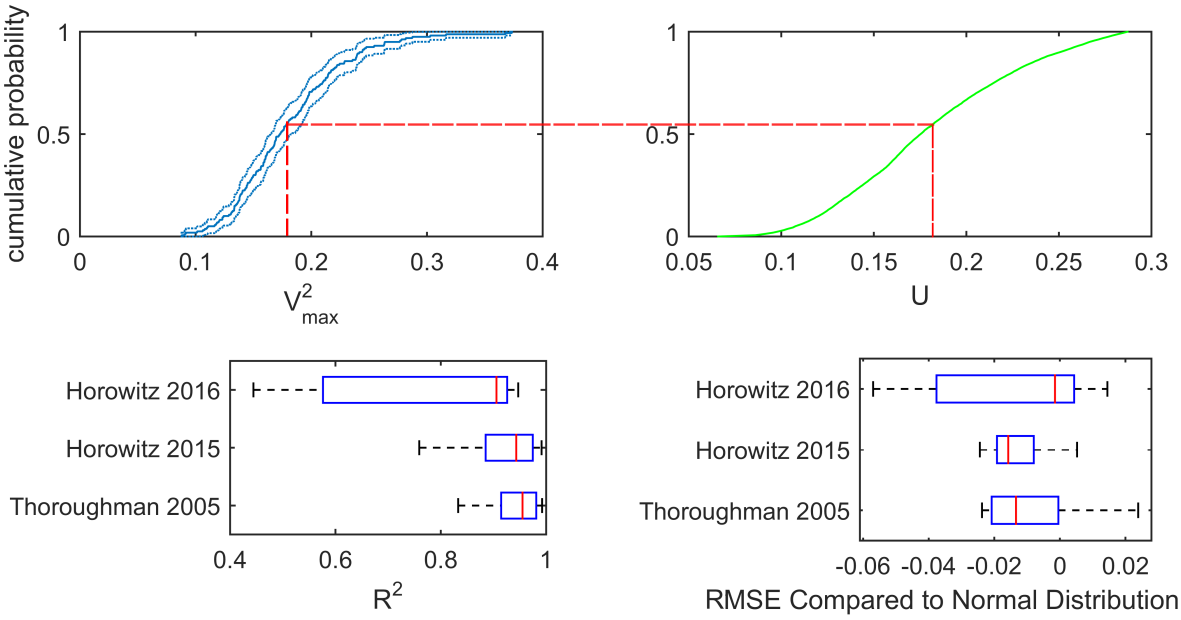

Figure 13. We relate the simulated cumulative probability of the maximum of a movement's subintents' changes in potential energy (top left) to the measured cumulative probability of the maximum of a movement's kinetic energy (top right). These cumulative probabilities provide a one-to-one relationship between the quantities (dashed red line). As we hypothesize a linear relationship between them that is governed by key model parameters, we demonstrate the linearity of this fit (bottom left) and its superiority to modeling these peaks as Gaussians (bottom right).

selection can be performed by a first-to-fire winner-takes-all. The behavior that results efficiently spans the space of exploration while still statistically acting to minimize expected cost, maximize expected reward, and maximize the rate of reward. No look-ahead is performed. No accurate recollection of the past is needed. The environment and state transitions are represented only as heuristics, which makes the selection of these heuristics critically important to performance.

Perhaps the most striking feature of this process is the central role of state appraisal. As the probability of occupying a state depends logarithmically on its appraisal, we can hypothesize that general satisfaction should depend logarithmically on the abstract capacity to change states. Surprisingly, this hypothesis is supported by behavioral economics research showing many measures of stress and well-being depend logarithmically on the most abstract form of appraisal known to humankind: money (72). Conceptually, money does obey the key limitation of a conservative potential: it is fungible. A dollar is a dollar no matter how it is earned or spent. Moreover, taking an action for pay generally enables one to use that pay to obtain the same action from others. Restated, money is a potential-like, state-dependent form of effort storage that impacts state satisfaction logarithmically just as we would expect of our state appraisal. These findings about the economy of movement may therefore extend to the economy of money.

While our findings may prove to be broad, our conclusions are limited by the nature of our experiments. We examined particularly low-dimensional tasks wherein we could reasonably assume that our subjects possess a complete and accurate understanding of their environment. Day-to-day life occurs in a markedly high-dimensional social and physical environment that

has evaded complete or accurate understanding. Moreover, we examined only point-to-point reaching and not walking, tool-use, or verbal communication. Examining these contexts to determine if our findings generalize to them is an important next step.

A process model for human movement may enable better understanding of movement disorders and replication of human-like movement in robotics. Characterization in terms of the quantities identified here holds the promise of identifying movement deficits and addressing them directly. Similarly, robots have many advantages over the human body: more-accurate sensors, faster processors, and more-powerful actuators. Despite these, we have struggled to produce human-like handling of complex loads. Replication of the human movement strategy may rectify this. Fortunately for the enterprise of science, markedly more study is needed.

CHAPTER 6

GENERAL DISCUSSION

All of the previous results flow from an attempt to address a deceptively simple question: what is error? My original motivation came from the error augmentation learning enhancement technology pioneered by my adviser (20). In order to use his technology, it is necessary to first define an error signal such as the difference between a person's movement and an ideal movement. While the model of Flash and Hogan (8) suggests that a straight-line reach to the target is ideal, very large disturbances might cause a subject to replan mid-movement. We therefore used the model of Shadmehr and Mussa-Ivaldi (6) to produce many different possible responses to a mid-movement force disturbance to determine which best-fit actual motions. This method was computationally demanding, and worse, it did not provide any clear answers.

While presenting on the method and its results at the yearly Neuroscience conference in 2011, I realized that optimization could be used to determine the best equilibrium trajectory to describe a particular motion rather than merely guessing at the forms that replanning might take. This optimization step naively seemed necessarily: systems as nonlinear as the human arm typically cannot be analytically inverted. As I examined the problem and made programming improvements, I found that the arm model could be analytically inverted and this inverse could be solved directly. My first step was to verify the tautology of this process, shown in Figure 1: if I simulate movement with a specified equilibrium trajectory and disturbance, I should recover the equilibrium through my inversion. Ultimately, the inversion also proved to be stable (Figure 2)

and provided reasonable estimates of the arm's equilibrium (Figure 3). Surprisingly, not all subjects were modifying their equilibrium in response to a very large disturbance. This meant that no general approximation could be used to replace the straight-line reach.

If a person intends any trajectory, it must be the equilibrium trajectory of the arm. However, it was not clear that humans intend their trajectories. For example, the control scheme of Todorov and Jordan only intends to reach a goal state while minimizing a measure of cost (10). In order to firmly establish that the arm's equilibrium is specifically intended, we next presented subjects with their arm's equilibrium as a cursor that they could control in realtime. We found that all of our subjects could immediately recognize and control their equilibrium (Figure 5). In the presence of disturbance, reaching with their intent caused our subjects' performance to increase (Figure 6) even as they were able to relax their arm (Figure 7). Taken together these results strongly support the claim that humans explicitly plan their arm's equilibrium trajectory and that it is their intent.

To answer our original question, we still need to predict when and how intent might change. We knew that infants (15) and stroke patients (16) show strong evidence of segmentation in their movements. Traditionally, these submovement segments have been interpreted as potential building blocks of movement. Our insight is that intent and movement can differ as seen in Chapters 2 and 3. On this basis, we examine intent specifically when it is different from movement to recover subintents. In addition, we use a clustering algorithm rather than explicit optimization as it is not clear what should be optimized. The clustering produced accurate, but

not precise, outcomes (Figure 11) that appeared to reveal exponential distributions in subintent parameters.

The exponential distribution frequently arises in physics where it relates the energy of a state to its probability. Since one of the exponentially distributed quantities has units of velocity squared, which is proportional to kinetic energy, this was a tantalizing hint that perhaps a relationship to statistical physics might be present. Ultimately, we were able to derive a statistical mechanics from assumptions of conservation of energy and memorylessness that could predict and explain the distributions of parameters we observed. As the laws of physics conserve energy and are memoryless, it is unsurprising that biology would take advantage of this. Of the quantities we find, perhaps the most interesting is the existence of a state-based energy term that mirrors potential energy and seems to possess many of the same properties as money. In addition, these quantities could be abstracted to predict distributions of whole-movement measures such as duration and rectified work. Together, these chapters demonstrate a means of viewing human intent, examining its construction, and predicting future intents.

6.1 Wide Applicability

I have no reason to believe that the algorithms of the previous paper are limited to reaching movement. Many of the quantities known to psychology such as preference and values may be understood in the framework of a potential field. This mirrors the findings of Weber's Law (73) that sensation is logarithmic as well as the findings of Kahneman that life satisfaction depends logarithmically on income (72). On this basis, it appears possible that higher-order human faculties are an abstraction of the mechanisms underlying motor control. Much of

this discussion will focus on the implications *if* the equations of Chapter 5 were somehow a complete description of human planning. This includes clinical applications of influencing human behavior.

6.2 Clinical Implications

Characterizing deficits following stroke and other neurological injuries might be as simple as taking the natural log of the probability of movement behaviors over some span of time. This is straight-forward to move into clinical and even home settings as cell phones incorporate motion censors and GPS. This makes the necessary tracking possible using only commodity hardware. In addition, cell phone processors are fast enough to compute the equations we have uncovered in real-time. This may enable real-time feedback in the spirit of Huang et al. (63) that can address movement deficits and irregularities.

In addition to identifying spacial deficits that might be linked to state appraisal or the lack of such appraisal, this approach may also be able to characterize differences in other parameters. For example, the increased number of subintents following stroke (16) indicates an decrease in product of the two parameters governing the distribution of the number of subintents, which we model in Equation 5.7: αq_0. We can relate this to the increased variability of movement that follows stroke (74) as variability is proportional to α^{-2}. An increase in α^{-2} implies a decrease in α, which in turn implies a decrease in αq_0 and an increase in the number of subintents. This link not only provides a testable hypothesis that these two features should be correlated, but in addition identifies a common mechanism underlying two key features of movement following stroke.

We can now design an intervention. As we should expect the spatial deficits following stroke to arise from a difference in $U(x)$ relative to healthy subjects, stroke patients may benefit from experiencing either the healthy potential field of the difference between the healthy potential field and their own rendered as a haptic environment. In addition, the high variability that implies a high α^{-2} also implies a high mean subintent magnitude, which is proportional to α^{-1}. We can readily reduce subintent magnitude simply by directing subjects to keep their speed below a threshold. This suggestion mimics the finding that slower movement at the outset of training has been found to be correlated with a better change in clinical scores (75).

6.3 A Theory of Learning

For any state appraisal, $U(x)$, we can re-represent this function without loss of generality as a kernel inner product $U(\vec{x}) = \sum_n a_n \kappa(\vec{p_n}, \vec{x})$ where a is an appraisal and p is a preferred stimulus. Stated as a physiologically reasonable example, a rat's appraisal of a position in a water maze will be calculated in layers. First, state will be computed as the generalized inner product of a preferred stimuli with an experienced stimulus by a place cell neuron's firing. Second, this firing rate will multiplied by a connection strength that depends on the number of post-synaptic receptors. This representation is not only very general, it is also physiologically realizable in an especially plain way.

We need not compute a new learning rule to see how the appraisals should adapt in the face of new state-outcome pairs as changing the weights assigned to each kernel is already understood (76). The change in a preferred stimulus's appraisal is proportional to that state's contribution to an appraisal error: $\Delta a_n \propto U_e a_n \kappa(p_n, x_e)$ where U_e denotes an error in appraisal and x_e

denotes the state at which at the error occurred. For large Δa_n, a new preferred stimulus is instead created such that $p_{n+1} = x_e$ and $a_{n+1} = U_e$. This structure allows us to set the appraisal of an experience or even erase it. To set the value of an experience, we need only create new experiences completely surrounding it whose appraisal we control.

This strategy is evident in advertisement where experiences that include a product – but primarily focus on extraordinary components – are used to modify the appraisal of the product. These extraordinary components are needed to avoid conflicting first-hand experience. For example, a basketball star might endorse a sports drink as tasty and energizing. As the viewer of such a commercial is very unlikely to possess contradictory experience, such as viewing the star selecting a different beverage when offered the drink in an actual game, this allows the sports drink company to modify a person's future appraisal of the product. While this effect gradually diminishes with first-hand experience of the product, new advertisements that possess small inner products with one another and first-hand experience can maintain a bias. This effect is dependent on the ability of large discrepancies between experience and advertisement to create new preferred stimuli instead of a mild adaption of existing appraisal. This allows us to make a straight-forward prediction regarding advertisement effectiveness. The least effective form of advertisement should be exposing the consumer directly to the product and informing them that they like it. The most effective form of advertisement should be exposing the consumer to a depiction of someone they idolize, but have not met, using and enjoying the product.

6.3.1 Application to Rehabilitation Robotics

If this theory of learning is correct, it explains a central finding of rehabilitation robotics. In a typical study, a chronic stroke survivor whose motor performance has plateaued is exposed to a robotic device for a relatively short amount of time. Following this exposure, a small but statistically significant improvement in both their performance and clinical assessments are detected with the study's intervention producing a somewhat larger improvement (20). Why should every exposure produce approximately the same consequence? This is readily explained if stroke has the effect of destroying preferred stimuli. Over time, experiences more or less orthogonal to the remaining stimuli fill in many gaps, but not all. The novelty of robotic rehabilitation and the increased novelty of experimental interventions create an opportunity for additional learning. This implies a testable hypothesis: wearing an upper-arm exoskeleton that does nothing more than compensate its own mass should cause improvement in chronic stroke patients. Moreover, performance on clinical assessments should be higher while the exoskeleton is worn as the boost comes from the novel context. If upheld, this hypothesis suggests a strategy for using orthgonality to compute optimal novelty-based interventions from measured movement deficits.

6.4 Conclusions

This thesis spans the construction and observation of human intent. We began by inverting an established model to recover a person's intent from their actions. Subjects could even control their intent if presented with it. We were able to break apart intent into its building blocks and find physics-based reasons for their properties. Finally, I extended these properties to a

provisional theory of learning and some of its potential consequences for rehabilitation. Any technology that can modify human behavior for medical purposes can likely be used to other ends as well.

APPENDICES

Appendix A

COPYRIGHT POLICIES

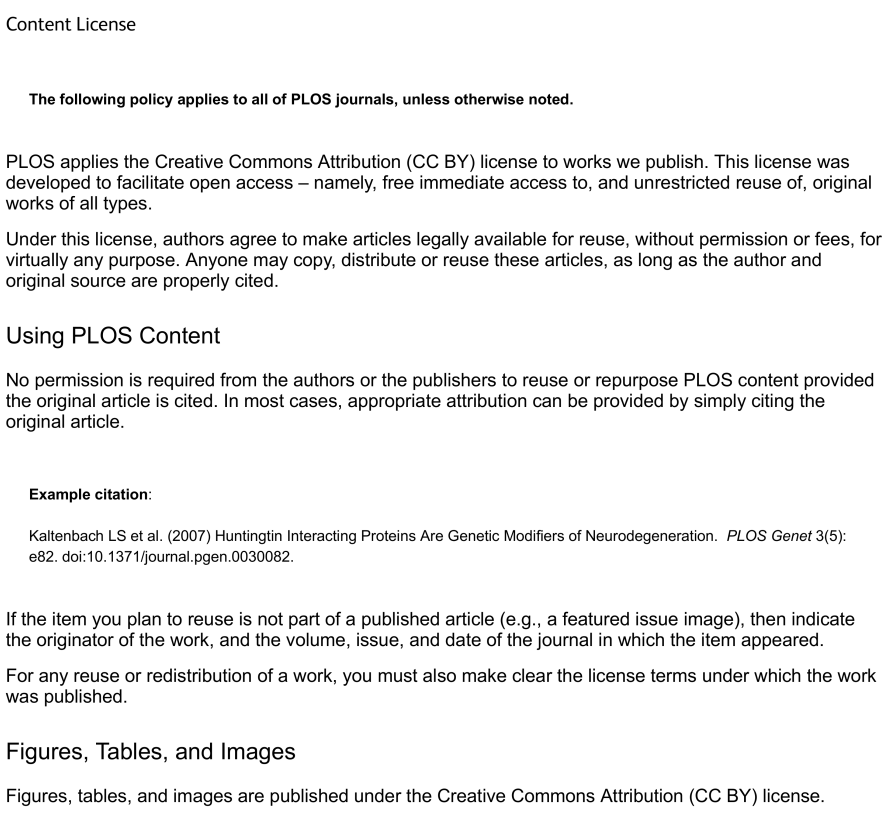

Content License

The following policy applies to all of PLOS journals, unless otherwise noted.

PLOS applies the Creative Commons Attribution (CC BY) license to works we publish. This license was developed to facilitate open access – namely, free immediate access to, and unrestricted reuse of, original works of all types.

Under this license, authors agree to make articles legally available for reuse, without permission or fees, for virtually any purpose. Anyone may copy, distribute or reuse these articles, as long as the author and original source are properly cited.

Using PLOS Content

No permission is required from the authors or the publishers to reuse or repurpose PLOS content provided the original article is cited. In most cases, appropriate attribution can be provided by simply citing the original article.

Example citation:

Kaltenbach LS et al. (2007) Huntingtin Interacting Proteins Are Genetic Modifiers of Neurodegeneration. *PLOS Genet* 3(5): e82. doi:10.1371/journal.pgen.0030082.

If the item you plan to reuse is not part of a published article (e.g., a featured issue image), then indicate the originator of the work, and the volume, issue, and date of the journal in which the item appeared.

For any reuse or redistribution of a work, you must also make clear the license terms under which the work was published.

Figures, Tables, and Images

Figures, tables, and images are published under the Creative Commons Attribution (CC BY) license.

Figure 14. My first paper (1) was published in PLOS ONE under a Creative Commons license, which requires only that it be properly attributed and a link to the license (http://creativecommons.org/licenses/by/4.0/) be provided.

Appendix A (Continued)

Figure 15. My first paper (2) was published in Frontiers in Behavioral Neuroscience, which requires only that it be properly attributed and a link to the license (http://creativecommons.org/licenses/by/4.0/) be provided.

CITED LITERATURE

1. Horowitz, J. and Patton, J.: I meant to do that: Determining the intentions of action in the face of disturbances. PloS one, 10(9):e0137289, 2015.

2. Horowitz, J., Madhavan, T., Massie, C., and Patton, J.: Reaching is better when you get what you want: Realtime feedback of intended reaching trajectory despite an unstable environment. Frontiers in Behavioral Neuroscience, 9, 2015.

3. Horowitz, J., Madhavan, T., Massie, C., and Patton, J.: Reaching is better when you get what you want: Realtime feedback of intended reaching trajectory despite an unstable environment. EMBC, 9, 2016.

4. Burdet, E., Tee, K. P., Mareels, I., Milner, T. E., Chew, C.-M., Franklin, D. W., Osu, R., and Kawato, M.: Stability and motor adaptation in human arm movements. Biological cybernetics, 94(1):20–32, 2006.

5. Saltelli, A., Annoni, P., Azzini, I., Campolongo, F., Ratto, M., and Tarantola, S.: Variance based sensitivity analysis of model output. design and estimator for the total sensitivity index. Computer Physics Communications, 181(2):259–270, 2010.

6. Shadmehr, R. and Mussa-Ivaldi, F. A.: Adaptive representation of dynamics during learning of a motor task. The Journal of Neuroscience, 14(5):3208–3224, 1994.

7. Bernstein, N. A.: The co-ordination and regulation of movements. 1967.

8. Flash, T. and Hogan, N.: The coordination of arm movements: an experimentally confirmed mathematical model. The journal of Neuroscience, 5(7):1688–1703, 1985.

9. Latash, M. L., Scholz, J. P., and Schoner, G.: Toward a new theory of motor synergies. MOTOR CONTROL-CHAMPAIGN-, 11(3):276, 2007.

10. Todorov, E. and Jordan, M. I.: Optimal feedback control as a theory of motor coordination. Nature neuroscience, 5(11):1226–1235, 2002.

11. Feldman, A. G. and Levin, M. F.: The origin and use of positional frames of reference in motor control. Behavioral and Brain Sciences, 18(04):723–744, 1995.

12. Gomi, H. and Kawato, M.: Human arm stiffness and equilibrium-point trajectory during multi-joint movement. Biological cybernetics, 76(3):163–171, 1997.

13. Bizzi, E., Accornero, N., Chapple, W., and Hogan, N.: Posture control and trajectory formation during arm movement. The Journal of Neuroscience, 4(11):2738–2744, 1984.

14. Woodworth, R. S.: Accuracy of voluntary movement. The Psychological Review: Monograph Supplements, 3(3):i, 1899.

15. von Hofsten, C. and Lindhagen, K.: Observations on the development of reaching for moving objects. Journal of experimental child psychology, 28(1):158–173, 1979.

16. Rohrer, B., Fasoli, S., Krebs, H. I., Volpe, B., Frontera, W. R., Stein, J., and Hogan, N.: Submovements grow larger, fewer, and more blended during stroke recovery. Motor control, 8(4), 2004.

17. Dempster, W. and Center, W. A. D.: Space requirements of the seated operator: Geometrical, kinematic, and mechanical aspects of the body with special reference to the limbs. Technical report, Wright Patterson Air Force Base, 1955.

18. Winter, D.: Biomechanics and Motor Control of Human Movement. Wiley, 2009.

19. Gomi, H. and Osu, R.: Task-dependent viscoelasticity of human multijoint arm and its spatial characteristics for interaction with environments. The Journal of Neuroscience, 18(21):8965–8978, 1998.

20. Patton, J. L. and Mussa-Ivaldi, F. A.: Robot-assisted adaptive training: custom force fields for teaching movement patterns. Biomedical Engineering, IEEE Transactions on, 51(4):636–646, 2004.

21. MATLAB: version 7.6.0.324 (R2008a). Natick, Massachusetts, The MathWorks Inc., 2008.

22. Thoroughman, K. A. and Shadmehr, R.: Learning of action through adaptive combination of motor primitives. Nature, 407(6805):742–747, 2000.

23. Fryar, C. D., Gu, Q., and Ogden, C. L.: Anthropometric reference data for children and adults: United states, 2007–2010. National Center for Health Statistics Vital and Health Statistics, 11:1–40, 2012.

24. Stunkard, A. J. and Albaum, J. M.: The accuracy of self-reported weights. The American journal of clinical nutrition, 34(8):1593–1599, 1981.

25. Franklin, D. W., Osu, R., Burdet, E., Kawato, M., Milner, T. E., et al.: Adaptation to stable and unstable dynamics achieved by combined impedance control and inverse dynamics model. Journal of neurophysiology, 90(5):3270–3282, 2003.

26. Pruszynski, J. A. and Scott, S. H.: Optimal feedback control and the long-latency stretch response. Experimental brain research, 218(3):341–359, 2012.

27. Gawthrop, P., Loram, I., Lakie, M., and Gollee, H.: Intermittent control: a computational theory of human control. Biological cybernetics, 104(1-2):31–51, 2011.

28. Patton, J. L., Stoykov, M. E., Kovic, M., and Mussa-Ivaldi, F. A.: Evaluation of robotic training forces that either enhance or reduce error in chronic hemiparetic stroke survivors. Experimental Brain Research, 168(3):368–383, 2006.

29. Guadagnoli, M. A. and Lee, T. D.: Challenge point: a framework for conceptualizing the effects of various practice conditions in motor learning. Journal of motor behavior, 36(2):212–224, 2004.

30. Hodgson, A. J. and Hogan, N.: A model-independent definition of attractor behavior applicable to interactive tasks. Systems, Man, and Cybernetics, Part C: Applications and Reviews, IEEE Transactions on, 30(1):105–118, 2000.

31. Franklin, D. W., Liaw, G., Milner, T. E., Osu, R., Burdet, E., and Kawato, M.: Endpoint stiffness of the arm is directionally tuned to instability in the environment. The Journal of Neuroscience, 27(29):7705–7716, 2007.

32. Crevecoeur, F. and Scott, S. H.: Beyond muscles stiffness: Importance of state-estimation to account for very fast motor corrections. PLoS computational biology, 10(10):e1003869, 2014.

33. Van Eesbeek, S., De Groot, J. H., Van der Helm, F. C., and De Vlugt, E.: In vivo estimation of the short-range stiffness of cross-bridges from joint rotation. Journal of biomechanics, 43(13):2539–2547, 2010.

34. McClelland, D. C.: How motives, skills, and values determine what people do. American Psychologist, 40(7):812, 1985.

35. Rawolle, M., Schultheiss, M., and Schultheiss, O. C.: Relationships between implicit motives, self-attributed motives, and personal goal commitments. Frontiers in psychology, 4, 2013.

36. Ziebart, B. D., Bagnell, J., and Dey, A. K.: Modeling interaction via the principle of maximum causal entropy. In Proceedings of the 27th International Conference on Machine Learning (ICML-10), pages 1255–1262, 2010.

37. Strausser, K. A. and Kazerooni, H.: The development and testing of a human machine interface for a mobile medical exoskeleton. In Intelligent Robots and Systems (IROS), 2011 IEEE/RSJ International Conference on, pages 4911–4916. IEEE, 2011.

38. Hargrove, L. J., Simon, A. M., Young, A. J., Lipschutz, R. D., Finucane, S. B., Smith, D. G., and Kuiken, T. A.: Robotic leg control with emg decoding in an amputee with nerve transfers. New England Journal of Medicine, 369(13):1237–1242, 2013.

39. Englehart, K. and Hudgins, B.: A robust, real-time control scheme for multifunction myoelectric control. Biomedical Engineering, IEEE Transactions on, 50(7):848–854, 2003.

40. Young, A. J., Hargrove, L. J., and Kuiken, T. A.: Improving myoelectric pattern recognition robustness to electrode shift by changing interelectrode distance and electrode configuration. Biomedical Engineering, IEEE Transactions on, 59(3):645–652, 2012.

41. Belen'kii, V., Gurfinkel, V., and Pal'tsev, E.: On the control elements of voluntary movements. Biofizika, 1967.

42. Mirabella, G.: Should i stay or should i go? conceptual underpinnings of goal-directed actions. Frontiers in systems neuroscience, 8, 2014.

43. Miles, F. and Eighmy, B.: Long-term adaptive changes in primate vestibuloocular reflex. i. behavioral observations. J neurophysiol, 43(5):1406–1425, 1980.

44. Pine, Z. M., Krakauer, J. W., Gordon, J., and Ghez, C.: Learning of scaling factors and reference axes for reaching movements. Neuroreport, 7(14):2357–2362, 1996.

45. Takahashi, C., Scheidt, R. A., and Reinkensmeyer, D.: Impedance control and internal model formation when reaching in a randomly varying dynamical environment. Journal of neurophysiology, 86(2):1047–1051, 2001.

46. Scheidt, R. A., Dingwell, J. B., and Mussa-Ivaldi, F. A.: Learning to move amid uncertainty. Journal of neurophysiology, 86(2):971–985, 2001.

47. Niu, C. M., Corcos, D. M., and Shapiro, M. B.: Temporal shift from velocity to position proprioceptive feedback control during reaching movements. Journal of neurophysiology, 104(5):2512–2522, 2010.

48. Gough, B.: GNU scientific library reference manual. Network Theory Ltd., 2009.

49. MATLAB: version 8.3.0.532 (R2014a). Natick, Massachusetts, The MathWorks Inc., 2014.

50. Mistry, M., Theodorou, E., Schaal, S., and Kawato, M.: Optimal control of reaching includes kinematic constraints. Journal of neurophysiology, 110(1):1–11, 2013.

51. Mirabella, e. a.: Context influences on the preparation and execution of reaching movements. Cogn Neuropsychol, 25(1), 2008.

52. Mirabella, G., Iaconelli, S., Modugno, N., Giannini, G., Lena, F., and Cantore, G.: Stimulation of subthalamic nuclei restores a near normal planning strategy in parkinsons patients. PloS one, 8(5), 2013.

53. Haggard, P.: Human volition: towards a neuroscience of will. Nature Reviews Neuroscience, 9(12):934–946, 2008.

54. Ghez, C., Favilla, M., Ghilardi, M., Gordon, J., Bermejo, R., and Pullman, S.: Discrete and continuous planning of hand movements and isometric force trajectories. Experimental Brain Research, 115(2):217–233, 1997.

55. Ziebart, B., Dey, A., and Bagnell, J. A.: Probabilistic pointing target prediction via inverse optimal control. In Proceedings of the 2012 ACM international conference on Intelligent User Interfaces, pages 1–10. ACM, 2012.

56. Sakitt, B.: Visual-motor efficiency (vme) and the information transmitted in visual-motor tasks. Bulletin of the Psychonomic Society, 16(5):329–332, 1980.

57. Krebs, H. I., Aisen, M. L., Volpe, B. T., and Hogan, N.: Quantization of continuous arm movements in humans with brain injury. Proceedings of the National Academy of Sciences, 96(8):4645–4649, 1999.

58. Grafton, S. T. and Tunik, E.: Human basal ganglia and the dynamic control of force during on-line corrections. The Journal of Neuroscience, 31(5):1600–1605, 2011.

59. Rohrer, B. and Hogan, N.: Avoiding spurious submovement decompositions ii: a scattershot algorithm. Biological cybernetics, 94(5):409–414, 2006.

60. Conditt, M. A. and Mussa-Ivaldi, F. A.: Central representation of time during motor learning. Proceedings of the National Academy of Sciences, 96(20):11625–11630, 1999.

61. Rohrer, B. and Hogan, N.: Avoiding spurious submovement decompositions: a globally optimal algorithm. Biological cybernetics, 89(3):190–199, 2003.

62. Liao, J. Y. and Kirsch, R. F.: Characterizing and predicting submovements during human three-dimensional arm reaches. PloS one, 9(7):e103387, 2014.

63. Huang, F. C. and Patton, J. L.: Individual patterns of motor deficits evident in movement distribution analysis. In Rehabilitation Robotics (ICORR), 2013 IEEE International Conference on, pages 1–6. IEEE, 2013.

64. Olds, J. and Milner, P.: Positive reinforcement produced by electrical stimulation of septal area and other regions of rat brain. Journal of comparative and physiological psychology, 47(6):419, 1954.

65. Shadmehr, R., de Xivry, J. J. O., Xu-Wilson, M., and Shih, T.-Y.: Temporal discounting of reward and the cost of time in motor control. The Journal of Neuroscience, 30(31):10507–10516, 2010.

66. Horowitz, J. R., Abdel Majeed, Y., and Patton, J. L.: A fresh perspective on dissecting action into discrete submotions. 2016.

67. Zipf, G. K.: Human behavior and the principle of least effort. 1949.

68. Inman, V. T., Eberhart, H. D., et al.: The major determinants in normal and pathological gait. J Bone Joint Surg Am, 35(3):543–558, 1953.

69. Huang, H. J., Kram, R., and Ahmed, A. A.: Reduction of metabolic cost during motor learning of arm reaching dynamics. The Journal of neuroscience, 32(6):2182–2190, 2012.

70. Woledge, R. C., Curtin, N. A., and Homsher, E.: Energetic aspects of muscle contraction. Monographs of the physiological society, 41:1, 1985.

71. Hammersley, J. M. and Clifford, P.: Markov fields on finite graphs and lattices. 1971.

72. Kahneman, D. and Deaton, A.: High income improves evaluation of life but not emotional well-being. Proceedings of the national academy of sciences, 107(38):16489–16493, 2010.

73. Fechner, G. T.: Elemente der psychophysik. Breitkopf und Hrtel, 1860.

74. Kamper, D. G., McKenna-Cole, A. N., Kahn, L. E., and Reinkensmeyer, D. J.: Alterations in reaching after stroke and their relation to movement direction and impairment severity. Archives of physical medicine and rehabilitation, 83(5):702–707, 2002.

75. Majeed, Y. A., Abdollahi, F., Awadalla, S., and Patton, J.: Multivariate outcomes in a three week bimanual self-telerehabilitation with error augmentation post-stroke. In Engineering in Medicine and Biology Society (EMBC), 2015 37th Annual International Conference of the IEEE, pages 1425–1431. IEEE, 2015.

76. Liu, W., Principe, J., and Haykin, S.: Kernel Adaptive Filtering: A Comprehensive Introduction. Adaptive and Cognitive Dynamic Systems: Signal Processing, Learning, Communications and Control. Wiley, 2011.

VITA

Justin received his BS in Biology (2005) at Kalamazoo College and a BS/MS in Biomedical Engineering at Washington University (2008). He pursued his PhD in Biomedical Engineering at the University of Illinois at Chicago and has completed his research at The Rehabilitation Institute of Chicago in the Robotics Lab. His first publication (1) received the Sarah Baskin Award (2015) for excellence in research and was featured in both national and international news. His second (2) and third publications (66) were extensions of the first paper and a bridge to a final publication that will be submitted to high impact journals in the near future.