NORTHWESTERN UNIVERSITY

Motion as an Information Signal in Physical Human-Robot Interaction

A DISSERTATION

SUBMITTED TO THE GRADUATE SCHOOL

IN PARTIAL FULFILLMENT OF THE REQUIREMENTS

for the degree

DOCTOR OF PHILOSOPHY

Field of Mechanical Engineering

By

Kathleen Fitzsimons

EVANSTON, ILLINOIS

December 2020

ABSTRACT

Motion as an Information Signal in Physical Human-Robot Interaction

Kathleen Fitzsimons

Robots can be capable partners when interacting with humans, but their value is largely dependent on how information is communicated in that partnership. In physical human-robot interaction, information is communicated via motion—configurations, velocities, forces, and torques. The autonomy interprets these implicit signals using metrics, which ultimately drive the the autonomy. This thesis focus on how motion measures affect the performance of the closed loop controller and our ability to statistically characterize differences in motion due to deficit, assistance, and learning.

Perhaps the most common way to implement a control solution is to include a feedback loop around the error with respect to a referent. In human motion, a single trajectory cannot capture all the possible solution strategies or variance in a single task. Therefore, I begin by describing a hybrid shared control that avoids specifying a time series of states by using methods from model predictive control to assess user action. The resulting controller improved training outcomes compared to unassisted practice and exhibited several features that are critical to learning in physical human-robot interaction (pHRI). Analysis

of the study of the hybrid shared control showed that a measure of the information about the task encoded in the motion—ergodicity—was able to statistically capture the effect of assistance and training when error and task specific measures were not able to detect one or the other. I conclude this thesis by demonstrating how one could close the loop on this information measure, such that robot provides forceful feedback that supports the task goal rather than reduces local errors.

Acknowledgements

There are many people in my life who have supported me academically and personally throughout my PhD, but I have to start somewhere, so I will start with my advisor Todd. I would like to express my sincere appreciation for all the answers to quick questions that turned out to be 45+ minute discussions on topics ranging from minute statistical details to advising philosophy. Your support and advice has made me a better researcher and teacher. Most of all, I have appreciated the culture that you have fostered in your lab, and I hope I can give my students the same friendly, fun environment that I have enjoyed as a student. I would also like to thank my committee members, Jules Dewald without whom I would not have been able to pursue the interdisciplinary work contained in this thesis, and Brenna Argall, who has pushed me to think about the bigger picture by asking tough but thoughtful questions.

To my past and current lab mates, than you for your friendship, for listening to me complain about silly coding errors, and happily eating all the products of my stress-baking. You made work feel more like home. I would like to especially thank Ahalya, Ola, and Tommy, who were fantastic collaborators that provided valuable insight and kept me excited about the work we were doing.

I would also like to thank my siblings, Megan, Jamie, Lauren, Nick, Lucy, and Mitch for their love and support. I must especially acknowledge my brother, Matt, and my mom, Paula. They are the reason I started down this path and continue to be my inspiration

in both my personal and professional life. Finally, thank you to Emery for being there through every deadline, meltdown, and every moment of celebration along the way to this thesis.

Dedication

To my daughter, you did not help with this thesis, and you often make things more difficult. I love you anyways, and you have made me stronger, more patient, and more fulfilled than I ever could have imagined.

List of abbreviations

AAN	assist-as-needed
HRI	human-robot interaction
HSC	hybrid shared control
MIG	mode insertion gradient
OCIP	optimal controller inner product
pHRI	physical human-robot interaction
PRA	percent of rejected actions
VF	virtual fixtures

Table of Contents

List of Tables

List of Figures

The control trajectories also spend time near the goal state, but to a lesser extent. On the right, the trained trajectories also have high density near $\theta = \pm\pi$, but there are large bands of high density in the region $-1.5 \leq \theta \leq 1.5$ and $-4 \leq \dot{\theta} \leq 4$. This suggests that the trained group's motions were more consistent with the task goal—making the statistics of the trained group closer to the spatial statistics of the reference Dirac delta distribution, so the ergodic measure of the trained group is lower than that of the controls. 67

3.9 The results of the OCIP study demonstrate that subjects trained in week 1 retain high performance levels in week 2 as measured by RMS error and ergodicity. In the first 2 blocks of trials, the error and ergodicity of the control group are higher than that of the trained group. The trained group retains their initial performance level, while the control group continues to improve—eventually reaching the same level of performance as the trained group. It appears the feedback helped with retention because the learning was more structured. Note that the performance measures in week 1 (gray) were not used in the statistical analysis to avoid measuring the effects of the assistance itself. 70

3.10 The MIG filter study demonstrated that the filter successfully assisted subjects in set 2 compared to controls. Moreover, trained subjects outperformed the control group in set 3. Note: error bars indicate

CHAPTER 1

Introduction

As robots have transitioned from highly controlled settings in manufacturing environments to use alongside and in collaboration with humans, the robotics community has developed numerous machine learning approaches to help these autonomous systems navigate these uncertain and unstructured environments. Often, the goal of learning is to emulate human behavior or human learning strategies. Rather than examining robot learning, the ideas in this thesis focus on using robots as tools to enhance human learning processes—specifically human learning of physical tasks.

Robot-mediated training is not new. Some commercial robots have been available since the 2000's for rehabilitation. The interest in robot-mediated training stems from the fact that robots can support many repetitions of an exercise or task and can provide quantitative evaluation of the progress made during training. This reasoning assumes that we have some reference definition of the task, but what should that be and how should we adapt support to the individual needs of the user? The simplest reference definition is a trajectory, but a reference trajectory imposes a strict adherence to a time-series of states that does not correspond to stereotypical human motion. There are often multiple equally good ways to perform a task—suggesting a need to encode this natural variation into the reference or adapt the control online. This raises the question of what quantity should be used to evaluate the human for the purposes of closing the loop and adapting the high level support goals of the system. The most common measure used to update

the control at a low level is error with respect to some desired system states, but to adapt the support at a high level there are many performance heuristics used to evaluate skill or deficit in a particular task. Unlike error, heuristics often cannot be translated from one task to another. For instance, a heuristic for reaching is unlikely to be relevant to a balance task.

These metrics dictate whether or not the robot should be providing assistance, resisting an incorrect movement, or artificially increasing the difficulty of the task, so the chief concern of this thesis is not to ask what a robot should do to support learning. Instead we take a step back from that and ask how should human motion be measured and evaluated in the loop? The way that we measure human movement has a direct impact on how we can design algorithms with features that support principles of motor learning. This thesis describes approaches to managing physical human-robot interactions that do not require a strict time-dependent task definition by (1) focusing on the future impact of user actions and (2) spatial statistic measures of motion quality.

1.1. Main Contributions

This thesis addresses control design in physical human-robot interaction when the goal of the interaction is enhanced human motor learning. Specifically, it explores a training strategy that neither actively assists users or amplifies error, instead the robot gets out of the way of the user when action support the task goal or rejects user actions when they do not. This strategy is enabled by measuring the user action and its future impact on the dynamic system. Through user studies, I show how this control strategy enhances training and explore the impact of performance measures on our ability to analyze and predict

human motion. I introduce an information-theoretic approach to control in pHRI and empirically compare it to an error-based assist-as-needed (AAN) controller. *Thus, this thesis focuses on how one can design and validate control methods for physical human-robot interaction (pHRI) that enable one to create statistically consistent patterns of movement rather than trajectories or velocity profiles.*

1.1.1. Task-based Hybrid Shared Control

Traditional robotic control techniques have been designed to minimize error with respect to a desired trajectory or produce motions that minimize an objective function consisting of error and effort components. Early rehabilitation robots used recorded trajectories from human experts, healthy references, or optimal task completions and 'replayed' them with position controllers. However, these guidance-based approaches can provide too much assistance, so 'assist-as-needed' (AAN) shared control paradigm were introduced—usually by adjusting the gains on an impedance controller. While impedance-based AAN controllers can interfere less often based on performance heuristics, impedance control is fundamentally about desired velocity profiles rather than the task goal.

In some cases, the autonomy challenges the user through error amplification or random perturbations instead of providing feedback designed to guide the user towards the desired goal. Depending on the relative difficulty of the task, one approach or the other may be more effective. However, the results of robotic training in general are mixed. There are relatively few instances where robot-mediated training leads to more than small improvements over human-mediated training or unassisted practice. This is perhaps because the

adaptation of these training paradigms focuses on performance heuristics as opposed to online evaluation of user actions based on the task goal.

Using task-based evaluation could be especially beneficial for dynamic tasks with multiple acceptable strategies. One example of this type of task is cart-pendulum inversion, where a desired trajectory would become problematic both because slight perturbations can significantly change the desired actions and swinging the pendulum up from the right or left are equally good solutions. Rather than asking whether to assist or perturb user inputs, I ask whether or not the user input should be rejected or the robot should be transparent to the user according to whether or not the current action supports task objective. This control decision is made based on an evaluation of the impact of the current action on the system dynamics in the future using methods from model predictive control.

The contributions of this chapter are as follows:

(1) I describe a novel approach to shared control that switches between full user autonomy and full rejection of user inputs called hybrid shared control.

(2) Through user studies, I show that this approach enhances learning compared to unassisted practice

(3) I demonstrate that hybrid shared control exhibits several features important to motor learning, including improving user performance, adaptation to initial skill, and assist-as-needed.

In this work, I developed the implementation of hybrid shared control on an admittance controlled robot and designed a human subject study of training and retention over a one week period. I would like to acknowledge Aleksandra Kalinowska, who introduced an alternate acceptance criteria for hybrid shared control (the mode insertion gradient) with

which we were able to run a second human subject study, emphasizing the effectiveness of hybrid shared control in single training session. This work was published in [31].

1.1.2. Using Ergodicity for Analyzing and Predicting Human Motion

Measures are used for analysis, tracking progress, and to update intervention in robot-mediated training. Energy and error are useful metrics because they allow us to reason about the underlying principles of neuromotor control and because many well-developed engineering techniques are based on minimizing these quantities. The problem is that stereotypical motion—such as reaching, cleaning, self-feeding, and walking—have substantial variation between equally qualitatively successful trials, both within and between individuals. So how should one quantitatively assess performance and incorporate multiple demonstrations into our mathematical definition of a task when there are multiple equally reasonable approaches to solving the task?

One can develop task-specific performance heuristics that capture the qualitative description of task success and are independent of the task strategy. Yet, these are not transferable to other tasks and do not provide insight into the underlying mechanism in deficit or learning. In this work, I propose that we treat the body as an information channel and motion as information-carrying signal. I use ergodicity to measure how much information about the task is encoded in the motion. and demonstrate how this measures captures phenomena that would not otherwise be captured with task-specific performance heuristics or common engineering measures such as error.

The contributions of this chapter are as follows:

(1) We show that ergodicity, which can be interpreted as the degree to which a trajectory encodes information about a task, correctly predicts changes due to reduction of a persons existing deficit or the addition of algorithmic assistance.

(2) We also show that the measure captures changes from training with robotic assistance when other common measures for assessment failed to capture at least one of these effects.

In this chapter, I develop the rationale for using ergodicity as a measure of motion quality using reaching data from a stroke participant that was provided by Ana Maria Acosta and experimental data from the work in chapter 3. This work was published in [30].

1.1.3. Ergodic Shared Control

Interactions with complex systems require safe and reliable interfaces. However, such interfaces must be able to account for substantial uncertainty resulting from the unpredictable and highly variable nature of human behavior. In the first chapter of this thesis, I discuss instantaneous rejection or acceptance of user actions based on an objective function that includes terms for error from a desired target, control effort, and barrier functions. The impact of user actions on this objective function was evaluated at each time instant to avoid underlying problem with a time-dependent reference definition. Perhaps the most important feature of this control approach is its adaptability in real-time. The approach requires no predefined trajectory, runs on an indefinite time horizon, and automatically adapts to user skill. Yet when it comes to encoding multiple strategies and the natural variability in human movement, it suffers from the same shortcoming in defining the task that position controllers or path controllers do.

In the second chapter, I explore an alternative measure—the ergodic measure—to capture differences in task performance due to deficit and learning over time. Standard objectives such as error or task success fail to capture both of these differences statistically. The ergodic measure encodes the natural variation found in human motion using a distribution over relevant states as the reference definition. In this chapter, we present a method for hybrid shared control that uses the ergodic measure to encode the statistical definition of the task into the control framework rather than simply avoiding the problems associated with time-dependent trajectories or task definitions that can only account for one strategy. We compare this approach to an error-based path controller—virtual fixtures—in an experimental study of timed drawing. Ergodic hybrid shared control enables one to provide assistance based on global task goals as opposed to the local interactions managed by error-based approaches and provides a framework in which demonstrations can be used to define a task without loss of the variance encoded in those demonstrations.

The contributions of this chapter are as follows:

(1) I present a novel method for providing corrective feedback based on a statistical task definition that can be generalized to a broad set of tasks.
(2) Through a user study, I compare this training paradigm to a fixed assist-as-needed control paradigm called virtual fixtures.
(3) I demonstrate that ergodicity-based assistance can reduce error even compared to error-based assistance and leads to better task completion.

The work in this chapter has been submitted to ACM Transactions on Human-Robot Interaction and is currently under review.

Overall, this thesis provides a foundation for evaluating motion in terms of information measures as well as a potential solution to close the loop on this type of evaluation.

1.2. Dissertation Outline

This thesis begins with background information on human robot interaction and a discussion of related work in metrics for motion evaluation. The remaining chapters are dedicated to key works undertaken during my PhD and describe how the results motivate further lines of study.

Chapter 3 introduces the idea of a shared control paradigm that is hybrid in the sense that it discretely switches between full user autonomy and full rejection of user inputs. The decision to accept or reject is made using task-based criteria to instantaneously assess whether or not the user action will advance the system towards the task goal. This is in contrast to other control paradigms that modulate the level of assistance based on past performance and adjust the level of assistance on longer time frames. We find that using this task-based acceptance/rejection paradigm, the controller exhibits many features that are necessary for motor learning including improving task execution, adapting to user skill, and adapting to online performance. The controller enhances training within session and after one week compared to unassisted practice. The results of this human subject experiment motivated further analysis into the measures used to track the quality of task execution.

Chapter 4 uses the data from the experiments described in chapter 3 as well as data from a stroke participant to compare task-specific measures such as task-success and balance time to task-independent measures like error and ergodicity. Task-specific measures

failed to capture the effect of training. Error, on the other hand, statistically capture the effect of training, but did not capture changes in performance due to deficit or the addition of assistance. Ergodicity, which can be thought of as the amount of information about a task encoded in a motion, captured both of these effects. Therefore, it could be considered a useful metric to assess and predict human motion, particularly in the context of physical human-robot interaction.

Chapter 5 describes a way to close the loop on the ergodic metric using hybrid shared control. It is experimentally compared to training with an error-based control for HRI, virtual fixtures, in a timed drawing task. The results showed that ergodic-based assistance could increase performance in terms of reducing error and reducing the distance from ergodicity, while the error-based assistance sometime improved performance but led users to adopt a slower, more conservative strategy. Therefore, the group that trained with ergodic feedback learned to draw more complete images, whereas the group that trained with error feedback generated drawings that were precise, but less complete and less accurate.

CHAPTER 2

Background and Related Work

This thesis is principally concerned with enhancing human learning through physical human-robot interaction (pHRI). Therefore, this chapter provides background information on the broader landscape of research in human-robot interaction (HRI) and the unique issues associated with robot-mediated training and rehabilitation. Although HRI has a huge variety of applications, ultimately metrics are used to drive the autonomy, so a discussion of the relevant work on measures of human motion is also presented.

2.1. Human-Robot Interaction

Human-robot interaction is a huge research domain that covers a broad range of applications from humanoid robots for behavioral therapy to autonomous vehicles. The research problems addressed in the field of HRI can be both robot-centric, where human inputs are leveraged to improve the performance of the robot, or human-centric, where interfaces are designed enable people to easily control or receive assistance from a robot. Some work also focuses on human-robot collaboration where the is a focus on effective teaming. Regardless of how the HRI is framed, a key challenge to robots interfacing with or working in close proximity to humans in any context is communication. Whether it is a social robot processing natural language or an assistive wheelchair recording inputs from a joystick, the robot must acquire information from humans and correctly interpret it in the context of the environment and task. However, communication modalities in HRI are

diverse and complex. Some systems rely on explicit communication such as processing natural language input [123] or using binary queries [110]. Increasingly, autonomous systems acquire knowledge about attention, intent, or emotional state using implicit methods like tracking eye gaze and body posture [124]. Importantly, this communication is not one way. Robots act on incoming signals to generate situation-appropriate responses. In any situation, the autonomy uses metrics to drive the algorithmic responses.

Physical human-robot interaction is a special case of communication because watching a robot walk around you (e.g., a cleaning robot) is much different than walking around *with* the robot(e.g., an exoskeleton). There are a 3 main motivations for considering consider physical interactions. First and foremost is that as humans and robots increasingly occupy the same physical spaces, unintended collisions need to be detected and managed to ensure safety [17]. Therefore, new tactile sensors and sensing modalities continue to be developed to improve the robot's sense of touch using force/torque sensors, capacitive sensing arrays, and robot 'skins' [20]. Second, physical corrections and guidance from humans can be used to help robots learn new behaviors or skills [6, 60, 72]. Finally, continuous intentional forceful interactions are necessary in cases where the human and robot are acting as a team—comanipulating an object [83] or sharing control of a vehicle [121]—or when the robot is physically supporting training exercises. In the case of unintended collisions or physical corrections, robots receive explicit communication via their tactile sensors, but in cases where humans and robots are completing a task in cooperation, its difficult for the human or the robot to assign explicit meanings to force/torque signals except in some limited task-free applications like rendering haptic objects in a virtual environment or force amplifying exoskeletons. Furthermore, implicit communication through forceful

interactions is not as well-studied as implicit information gained from non-tactile signals such as eye gaze [2] or explicit communication using natural language [123]. Therefore there is much more uncertainty about how each side of the human-robot pair will interpret the other's physical contributions to a joint task.

This thesis focuses on interactions in which the goal is for the human to learn from tactile interactions. It follows that we should, in general, expect reliance on the robot to decrease as users become more competent—meaning that our interaction strategy needs to change and adapt based on some measure of progress. Yet, there is not a clear way to measure competency or progress and algorithmically we need some metric to decide when and how to modulate the interaction strategy. Here, I take a task-oriented approach to close the loop on assistance and suggest that one can detect learning and understanding of a task by measuring the implicit information encoded in motion.

2.2. Reference Definition and Performance Metrics for Autonomy

Characterizing the quality of motion is most often done by measuring error with respect to target location and transit time to that target [97]. These measures have been well validated in the context of quasistatic tasks—those tasks where the timing of the tasks is not critical to accomplishing the task and where there is not very much variation in strategy. Some tasks, however, are very sensitive to timing and strategy (e.g., upright balance). In figure 2.1, we examine *inversion*—a classic problem in the field of control that depends on a strategy for first exciting dynamics and then stabilizing near an unstable equilibrium. Strategies for inversion can vary substantially from subject to subject and timing can vary as a function of perturbations to a trajectory. As a result,

the divergence between two trajectories can be quite large even when both accomplish a task well. Metrics, such as those that focus on error between a desired trajectory and the subject trajectory can overprescribe the behaviors necessary to accomplish a task.

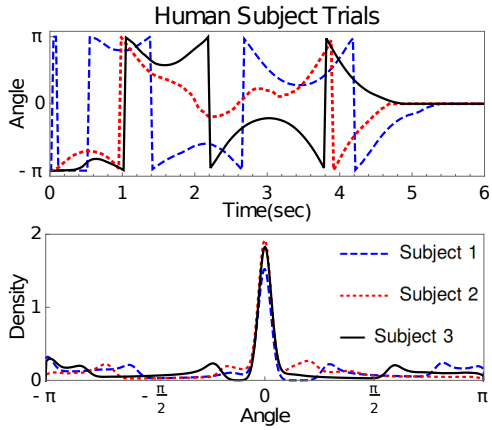

Figure 2.1. Different approaches to cart-pendulum inversion lead to very different trajectories. Nevertheless, the spatial statistics of the trajectories end up being very similar.

Why do people exhibit so much variation in their trajectories when they are nominally doing the same task? One possibility is that small differences in mechanics and interpretation of the goal lead to large differences in outcomes. I would instead argue that the measure of motion quality itself may be to blame—that error (the integrated norm of the difference between two curves) does not capture what makes two motions similar to each other. As an alternative, one can use measures of spatial statistical similarity of time evolutions (e.g. ergodicity) to quantitatively assess motion.

In humans, we observe stereotypical movement patterns for things like reaching, self-feeding, and walking despite the fact that there are infinitely many trajectories that can satisfy those high level task goals. The ability of the nervous system to plan [39], control [127], and learn [142] movements is commonly modeled using a control response to

a task-based objective [141]. Even in a simple mechanics-based model, human-like walking has been interpreted as an energy-minimizing trajectory that bears similar qualitative characteristics to human gait [14]. The objective function (energy, error) in these cases is typically explicitly dependent on the trajectory, most often the *error* with respect to a desired trajectory. Moreover, without the objective to describe the task, planning and control cannot be used to analyze or predict motion. As a result, error and energy are the *de facto* standards for creating a metric that describes how well a particular motion accomplishes a task.

In both the upper and lower limb tracking error or kinematic error is measured by comparing trajectories to some model of a normative trajectory. Error is a useful metric both because it is not task-specific and because many well-developed engineering techniques are based on minimizing error with respect to some reference trajectory [77, 103]. These optimal control techniques define the relationship between desired movement and motor control for both the study of human motor control [127] and the advancement of humanoid robots [62], but how should one rigorously define stereotypic movement patterns, particularly when there is a large amount of variation among normative motions? One possible answer is to approximate normative movements of the upper limb based on mathematical models such as a minimum jerk trajectory in reaching [58, 108, 139, 140]. Other options—used in rehabilitation robotics—are to use the movement of the "good" limb [40, 74, 131] or to scale/parameterize the time, amplitude, offset, and range of recorded movements from healthy subjects using online optimizations or machine learning occurring over many iterations of the same task [3, 5, 21, 26, 107]. For most tasks there is no model, so researchers choose one or more performance measures that are specific to the task such as

work area [24], movement speed [48], or a combination of velocity threshold, aim, and maximum reach [58] in order to evaluate human performance and/or modulate robotic assistance.

One reason these task specific measures capture the qualitative description of task success is that they are *independent of the strategy*. The result is that movement gets evaluated in highly task-specific manners, each one essentially a special case that does not generalize. For healthy subjects, variations in strategy may play less of a role in reaching tasks, but there are many strategies that subjects might use to enable balance, dynamic inversion, or locomotion. Given the amount of variation that exists in strategies used by healthy subjects, how should one choose a reference motion to measure performance? Measuring the statistical variation as a function of state (between a trajectory and the *distribution* of normative motions) might be preferable to measuring the error as a function of time (between the trajectory and an average normative motion). Comparing the statistics of two distributions is often done using measures of information, so I assert that information measures—for which there are information maximizing techniques [81, 133]—have the potential to be used to analyze movement, predict features of neuromotor control, and teach robots through observation of task executions.

CHAPTER 3

Task-Based Hybrid Shared Control for Training Through Forceful Interaction

Despite the fact that robotic platforms can provide both consistent practice and objective assessments of users over the course of their training, there are relatively few instances where physical human robot interaction has been significantly more effective than unassisted practice or human-mediated training. This chapter describes a hybrid shared control robot, which enhances task learning through kinesthetic feedback. The assistance assesses user actions using a task-specific evaluation criterion and selectively accepts or rejects them at each time instant. Through two human subject studies (total n=68), we show that this hybrid approach of switching between full transparency and full rejection of user inputs leads to increased skill acquisition and short-term retention compared to unassisted practice. Moreover, we show that the shared control paradigm exhibits features previously shown to promote successful training. It avoids user passivity by only rejecting user actions and allowing failure at the task. It improves performance during assistance, providing meaningful task-specific feedback. It is sensitive to initial skill of the user and behaves as an 'assist-as-needed' control scheme—adapting its engagement in real time based on the performance and needs of the user. Unlike other successful algorithms, it does not require explicit modulation of the level of impedance or error amplification during training and it is permissive to a range of strategies because of

its evaluation criterion. We demonstrate that the proposed hybrid shared control paradigm with a task-based minimal intervention criterion significantly enhances task-specific training.

3.1. Introduction

Approaches to designing kinesthetic feedback for robotic training platforms lie on a spectrum from antagonistic and resistive strategies that are dynamically updated based on user performance to passive assistive strategies in which users have a consistent guide during training. Training regimens at either end of the spectrum have been shown to be appropriate depending on the type and relative difficulty of the task. Passive assistance in the form of virtual fixtures [109] or record and replay strategies can provide task-relevant feedback to users by demonstrating correct movements. However, this type of guidance may not engage or challenge users because it does not dynamically adapt to different users or changes in user performance. Training in which errors are amplified rather than reduced by guidance has been effective in inducing adaptations in healthy and impaired individuals [89] during quasistatic reaching, but guidance was more effective in a timing-based motor task when individuals were less skilled [82]. Active assistance or shared control has been introduced as an alternative where the level of assistance or impedance is modulated based on performance heuristics. Though the results of robotic training are mixed, meta-analysis of studies using robotics in therapeutic settings demonstrate small but significant improvements in patient outcomes compared to usual care [57].

Here we present a hybrid shared control paradigm that lies in the middle of that spectrum—it does not resist or aid correct actions but requires user action for task completion. The autonomy evaluates user inputs based on criteria that capture how well the current input contributes to task completion. If the filtering criterion is met, the controller is transparent to the user. When the criterion is not met, the robot physically rejects the user input, providing feedback but not guidance. Rather than adjusting the relative contributions of the robot and human on a continuum based on heuristics over past performance of the user, we hypothesize that using an evaluation criterion to instantaneously switch between full user control and full rejection of user actions by the autonomy is sufficient to improve user performance, adapt to user skill, and ultimately enhance learning of a task.

The user input is evaluated at each time instant, using methods from model predictive control, which allows us to avoid prescribing a desired trajectory over time. This enables users to try different task completion strategies, to make errors, and to fail—all of which are critical to learning [55, 64, 125]. Additionally, the fact that we choose to only reject user input rather than replacing user input means that users must engage in the task actively to achieve success. The results of two user studies demonstrates that the controller-filter also adapts to the initial skill of the users, and adjusts the level of assistance based on current user performance much like an assist-as-needed controller. It does this without any pre-training assessments of the user's initial skill and without evaluating the overall performance of the subject within the current trial or any preceding trials. We find that this form of hybrid shared control is an effective training tool for both improving skill acquisition and retention of skill one week post-training.

In this chapter we show that a hybrid approach to switching between full user autonomy and full rejection of user inputs is an effective way to enhance learning through forceful interaction with a robot. Furthermore, we show, through two user studies, that the task-based switching control leads to improved subject performance while the assistance is engaged, decreased intervention for highly skilled users, and assistance that increases when subject performance is poor and becomes more transparent when subjects perform well.

The chapter is organized as follows. First, we review relevant work in robotic training in Section 3.2 and our prior work in Section 3.3. We introduce the hyrbid shared control algorithm in Section 3.4.1 and discuss the task-based criteria used to assess user inputs in Section 3.4.2. The experimental platform and protocol is discussed in Sections 3.4.3 and 3.4.6, respectively. Experimental results of two user studies are given in Section 3.5—discussing the training effect in Section 3.5.1 and the relevant features in Section 3.5.2-3.5.4. Finally, a discussion of the results and their implications for future work is given in Section 3.6.

3.2. Relevant Background

Using robotics in training provides a platform for consistent, high intensity repetitions that are not limited by the time the coach or therapist has available. In rehabilitation settings, specifically, devices can provide support and safety—reducing the physical and cognitive load of the caregiver. Patients who receive additional therapy with robotics often have improved clinical outcomes compared to patients receiving the standard of care [59, 74, 104, 119, 134]. Furthermore, robotics can quantitatively assess users [122]

and have the potential to systematically tailor the interaction to the user's skill or level of impairment. As a result, there is interest in facilitating training and rehabilitation through forceful interaction between robots and humans.

Numerous devices and control strategies have been developed to support physical human robot interaction (pHRI) and modulate it based on principles of motor learning. Despite the development of novel hardware and software to facilitate pHRI for training and therapy, there are relatively few instances where robotics have been used to significantly improve learning outcomes. Gains are often modest [80, 96] or equivalent to a similar amount of human-mediated training [18, 69, 132]. The success of robot-mediated therapy is highly dependent on the principles used to design robotic assistance and the corresponding features of training interfaces, which vary greatly from one implementation to another.

Traditional robotic control techniques have been designed to minimize error with respect to a desired trajectory or produce motions that minimize an objective function consisting of both error and effort components. Early rehabilitation robotics used a recorded trajectory from a human expert or healthy reference and 'replayed' it with position controllers [11, 15]. Alternatively, the reference was generated from an optimal task completion, such as minimum jerk reaching in the upper limb [33, 41]. Robotically assisting subjects to perform these normative movements has led to moderate improvements in training outcomes compared to unassisted practice [9, 49, 76]. This type of guidance has been especially effective when the learned task is difficult relative to the subject skill level [36] or the subject has a high level of impairment [13]. However, haptic guidance can actually interfere with learning [95, 112, 138] or lead to 'slacking' by the user [77, 105].

When learning a task, the central nervous system encodes not only a sequence of joint positions but also a feedback control loop—making motor output necessary to learning [116]. So while it is necessary for robotic trainers to be able to assist subjects in completing the task, especially when subjects have limited ability or skill, too much support—leading to user passivity—is not conducive to learning.

Rather than assisting subjects with task completion, some training paradigms act antagonistically to task goals, making aspects of the task more difficult and allowing failure. For instance, robotics have been used to introduce random noise-based disturbances into training. Supported by studies demonstrating that mistakes or errors actually enhance learning [125], training with this approach has been shown to improve training outcomes compared to progressive guidance strategies and unassisted practice [63]. Perturbation-based training could also improve the robustness of robot-mediated training—in human-robot teaming training with perturbations led to increased performance across task variants [101]. Alternatively, control strategies that explicitly amplify errors have been developed and have also been shown to improve motor learning in the upper limb [25, 27, 89], though the effects may be transient or may not generalize to other similar tasks [87]. Interestingly, error amplification is most effective when the users are not novices [82], suggesting that this antagonistic strategy is not appropriate for unskilled or highly impaired individuals. Finally, another approach is to allow users to make errors rather than enhancing them explicitly. Simply enabling kinematic variability has proved to be more effective than enforcing strict repetitive movement patterns [64]. As a result, impedance-based shared control has been widely adopted in pHRI to increase kinematic variability and allow users to make errors [55].

While shared control approaches are often implemented to augment user inputs such that task performance is optimized [19], it does not necessarily improve training outcomes [85]. The efficacy of blending control signals of a human expert [52] or robotic teacher [92, 100] with students through shared control varies depending on the task and mode of assistance [95]. Generally, shared control for training is considered most effective when the robot provides only as much assistance as is necessary based on estimates of user intent [65, 144], motor contribution [107], or other performance heuristics.

Assist-as-needed control schemes are implemented by dynamically updating the relative contributions of the robot and human. Updates to the relative contributions are made by adjusting the gains of an impedance controller based on measured outcomes [58, 91], introducing forgetting factors that adjust robot effort according to a schedule [26, 140], or implementing a repulsive potential field at the boundary of a virtual tunnel around a desired path [21].

Numerous implementations of assist-as-needed controllers have been developed for robots that support gait rehabilitation in exoskeletons [21], provide end-point guidance for upper limb tasks [28], offer support at anatomical joints in upper limb exoskeletons [140], and enhance sports training [78, 102, 135] with mixed results.

Given that approaches at either end of the assistive/resistive spectrum seem to be effective in some cases and ineffective in other training scenarios, one might ask what features of the interfaces discussed above create conditions conducive to motor learning? One idea that is consistent across training strategies is the need for user engagement and active participation [77], often accomplished by modulating the assistive or antagonistic forces based on subject performance trial to trial. However, it is still unclear how to best

implement real-time modulation. Literature suggests that it is necessary for platforms to be capable of assisting subjects in completing the desired task, especially when the user is unskilled. Yet, allowing or enhancing errors is critical to learning. In this chapter, we describe a novel shared control paradigm that, through an initial human subject study, we find to be successful in improving learning. We then explore the features of the shared control paradigm in the context of previous findings.

3.3. Prior Work

An algorithm for filtering control inputs was proposed in [129] for noise driven swing-up problems based on the hypothesis that noisy inputs can be a rich source of control authority if filtered in a meaningful task-specific way. This filter was implemented by combining a controller and a filter into a single computational unit that cancels noise samples not driving the system towards a desired control direction. In [32] and [50], we modified this algorithm to allow for filtering of *user input*. User inputs were either accepted or rejected based on the criteria described in Sections 3.4.2.1 and 3.4.2.2. When they were not accepted, they may be either rejected by the automation (as shown in Figure 1) or replaced with input prescribed by a control policy. In the experiments described in this work and our previous work, subject inputs were not replaced—allowing users to fail both allowed us to evaluate the participants' success rate during trials with and without the shared control and to evaluate the training effect of the kinesthetic feedback provided to them.

Previous experiments on a touchscreen platform in [32] represented an infinite actuation scenario for the filter, since user inputs were able to be completely rejected in

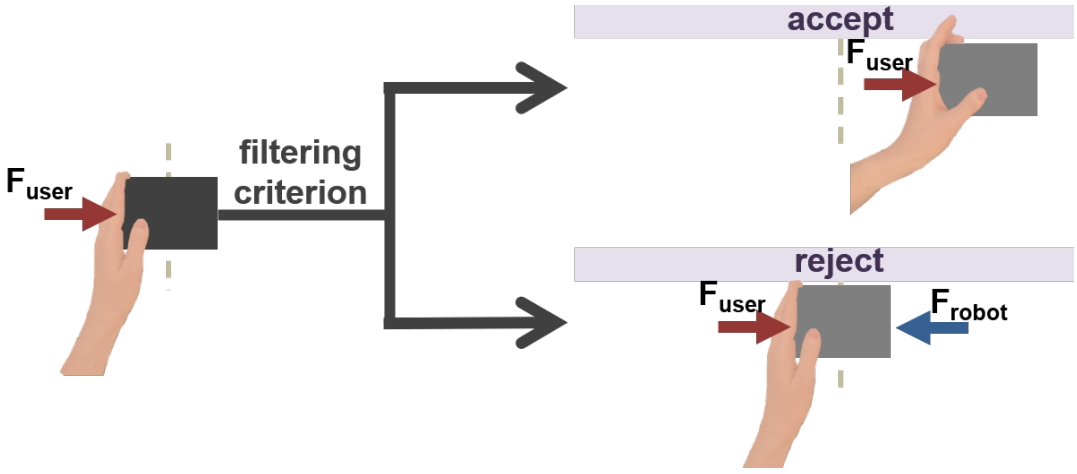

Figure 3.1. Robotic responses of hybrid share control on the example of a hand pushing a mass. The robot filters user input by physically accepting or rejecting it. When a user action is accepted, the robot admits the force. When a user action is not accepted, the robot rejects it by applying an equal and opposite force.

software. A haptic stylus (Phantom Omni by Sensable) on the other hand provided kinesthetic feedback, but did not have sufficient power to do more than weakly resist user inputs. We found that both implementations were able to effectively assist subjects in swinging up a cart-pendulum system compared to their baseline performance. The touch-screen platform indicated significantly higher success rates and lower time to success for the swing-up task. Although the assistance mode on the haptic platform did increase the success rate, there was no significant difference in time to success between the baseline and the assistance mode. This was likely due to the fact that the haptic interface did not generate enough force to strictly enforce the filter's acceptance criterion.

Therefore, we realized the mechanical filter on a higher power robotic system described in Section 3.4.3. Preliminary results of this work have been discussed in [50], where we noted a modest training effect compared to controls with unassisted practice as well

as a low, but significant correlation between the controller intervention rate and the participant's initial skill level. In this work, we extend these results by evaluating the progression of subject performance over time. We also present results using an alternative acceptance criterion and assess the skill retention of the trained group after one week.

3.4. Methods

3.4.1. Hybrid Shared Control

The hybrid shared control algorithm works as follows. Given a system and an operator, assume that a user input is measured every t_s seconds. The user input is assessed based on one of the acceptance criterion described in Section 3.4.2—roughly asking whether the user understands the task goal or an optimal control strategy for task completion. When the acceptance criterion is met, if the magnitude of the user command is within the allowed limits, the command is applied to the system. Otherwise, saturation may be applied.[1] On the contrary, if the criterion is not met, one of two alternatives can be followed: a) the system input can be set equal to zero (user command is "rejected") or b) the system input can be set equal to the nominal control value. The latter case would result in potentially never-failing interfaces, serving both training and safety purposes. Note that in our experimental setup we followed the first approach; the rationale behind this choice is that being allowed to fail in the task should provide clear indications as to whether the filtering algorithm has any effect on performance. When inputs were rejected in these experiments, a force equal and opposite to the force of the user is exerted at the end-effector. This results in the interface being transparent when user inputs are accepted

[1]Saturation limits may correspond to physical constraints e.g. angle or torque/force limits etc.

or velocity being held constant when inputs are rejected. This process is illustrated in Algorithm 1.

Algorithm 1 Hybrid shared control algorithm

Initialize current time t_0, sampling time t_s, time horizon length T, final time t_f, input saturation u_{sat} and angle tolerance γ.

while $t_0 < t_f$ **do**
    Infer user input u_{user} from sensor data
    Calculate the quantities in eq. 3.2 or 3.3 for time T.
    if Filter Criterion is True **then**
        if $|u_{user}| < u_{sat}$ **then**
            Use u_{user} as current input, $u_{curr} = u_{user}$
        else
            Apply saturated user input $u_{curr} = u_{sat}$
    else
        Completely "reject" u_{user} ($u_{curr} = 0$)
    Apply u_{curr} for $t \in [t_0, t_0 + t_s]$
    $t_0 = t_0 + t_s$
end while

3.4.2. Acceptance Criteria

In this chapter, we use two criteria. Both are reasonable interpretations of the hybrid philosophy of shared control. The Mode Insertion Gradient (MIG) assumes the user must be generating *descent directions* while the Optimal Controller Inner Product (OCIP) insists that the user *agrees* with the optimal control. Because of this difference, MIG is more relevant to assessing how well a person understands a task in the moment, whereas OCIP is more relevant to whether the person is being *taught* by the optimal control solutions we compute. Naturally these two interpretations have considerable overlap, but in different situations the choice may matter. For instance, a driver-assist wheelchair may need to interpret the quality of motion control a person is providing without having

an explicit need to train the user and potentially having reason to believe that the user needs flexibility in his/her implementation (leading to MIG being a better choice). On the other hand, technologies geared toward rehabilitation may want to steer a person's motor control towards a normative set of expected solutions (leading to OCIP). The practical consequences of these two interpretations of acceptance is that in the MIG study the acceptance criteria was met much more frequently and user actions were rejected less often than in the OCIP study.

3.4.2.1. Mode Insertion Gradient Criterion. The mode insertion gradient $\frac{dJ}{d\lambda}$ is most often used in mode scheduling problems to determine the optimal time τ to insert control modes from a predetermined set [4, 12, 22, 35, 136]. In these cases, it gives an estimate of the sensitivity of the cost function to the timing of a switch from one control mode to another. Therefore, a negative MIG at a specific time indicates that a mode switch at that time would decrease the cost compared to not switching modes. Often, the goal is to choose an application time when the MIG is most negative, to optimize the benefit of switching control modes. Here we use the mode insertion gradient as a measure of the sensitivity of the cost to a change from the nominal control, u_1, to a particular user input, u_2. Instead of using the MIG to decide *when* to switch modes, we use it to decide *whether* to switch modes and allow user input. To aid in this evaluation, we consider the MIG over the entire time horizon T and thus use the integral of it as our evaluation criterion. Our approach to calculating the MIG criterion is outlined below.

The mode insertion gradient is usually defined as

$$\frac{dJ}{d\lambda}(\tau) = \rho(\tau)^T \left[f(x(\tau), u_2(\tau)) - f(x(\tau), u_1(\tau)) \right] \tag{3.1}$$

for a system with dynamics

$$\dot{x}(t) = f(x(t), u(t), t) = g(x(t)) + B(x(t), t)u(t),$$

where $\dot{x}(t)$ is linearly dependent on the control u. In (3.1), state x is calculated using nominal control, u_1, and ρ is the adjoint variable calculated from the nominal trajectory $x(t)$,

$$\dot{\rho} = -\nabla l_1(x) - D_x f(x, u_1)^T \rho,$$

where $l_1(x, t)$ is the incremental cost and $\rho(t_0 + T) = \nabla m(x(t_0 + T))$ is the terminal cost. Moreover, in the work presented here, we define the nominal control, u_1, to be equivalent to the calculated controller action ($u_1(t) = u_{controller}$), and we define u_2 with the piece-wise function below,

$$u_2(t) = \begin{cases} u_{user} & t \leq t_0 + t_s \\ u_1 & t_0 + t_s < \ t \leq \ t_0 + T \end{cases}$$

where t_s is the sampling time, T is the time window over which we are evaluating system behavior, and u_{user} is a user input recorded at current time t_0. The control mode, u_2, is defined by a combination of user input at current time t_0 and actions from an optimal controller over time T into the future[2]. It is worth noting that u_1 is not a schedule of actions that is precomputed ahead of time, instead we calculate the best sequence u_1 every time step t_s based on the previously taken action and current state of the system.

[2]Sequential Action Control [130] was used to compute the nominal controller action for both criteria.However, any control policy that can be computed in real-time could be used.

The current user input, u_{user} is not included in the computation of u_1. In turn, the action sequence u_2 is defined by a combination of user input at current time t_0 and newly calculated actions from an optimal controller over time T into the future. This gives unique flexibility to the criterion and grants the user more control authority over the joint system, because any user action that could be corrected for by a future optimal action or sequence of optimal actions without destabilizing the system during the time window T will be admitted. Even suboptimal user actions will be allowed. The MIG quantitatively represents the benefit or disadvantage of allowing the user to push the system in the way that they are currently trying to move it.

When using MIG as an evaluation criterion, we calculate the integral of the mode insertion gradient over a time window T into the future

$$\int_{t_0}^{t_0+T} \frac{dJ}{d\lambda}(t)\delta t, \tag{3.2}$$

to evaluate the impact of user control u_2 on the system over time T[3]. When negative, the integral indicates that u_2—the user input—is a descent direction over the entire time horizon, which can be shown by evaluating the change in cost due to a control perturbation $u_2 - u_1$. Thus, the MIG integral can serve as the basis for evaluating the impact of a current user action on the evolution of a dynamic system over a time window into the future and has proven in our experiments to be a balanced assessment criterion—significantly improving performance while only minimally rejecting user actions.

3.4.2.2. Optimal Controller Inner Product Criterion. The optimal controller inner product (OCIP) criterion works in algorithm 1 by computing the value of a nominal

[3]The time window over which the MIG integral is evaluated, T was 1 s in the experiments discussed in this chapter. In general, the time window is chosen based on the dynamics of the system.

controller u_c based on the current state of the system. In this study, we use a model predictive controller described in [4], and when the system is near equilibrium, we switch to a linear quadratic regulator (LQR). Note that any controller could be used, but it should be capable of driving the system by itself according to the desired specification. Calculating the inner product between the user input and the nominal controller establishes whether or not the two vectors are in the same half plane (e.g. $\langle u_c, u_{user} \rangle > 0$). One can further specify that the user input vector must lie within a cone near the nominal control vector by specifying a maximum angle γ between u_{user} and u_c. If the user input lies in the same half plane as u_c and within γ radians of u_c, then the filter does nothing. This acceptance criterion is given by,

$$\langle u_c, u_{user} \rangle > 0 \textbf{ and } |\phi| \leq \gamma. \tag{3.3}$$

If the inner product between the control and the user command vector is positive, and the corresponding angle of the vectors is small, then the effect of user input on the system should be similar to that of the control vector. If the user input is not in the same half plane as u_c or not within γ radians of u_c, the input is rejected.

3.4.3. Experimental Platform

All subject data was collected using the New Arm Coordination Training device (NACT-3D) shown in Figure 3.2. The NACT-3D is a powerful haptic admittance-controlled robot that can be used to render virtual objects, forces, or perturbation in three degrees of freedom. This device is similar to that described in [122] and [23], to quantify upper limb motor impairments and provide a means to modulate limb weight support during

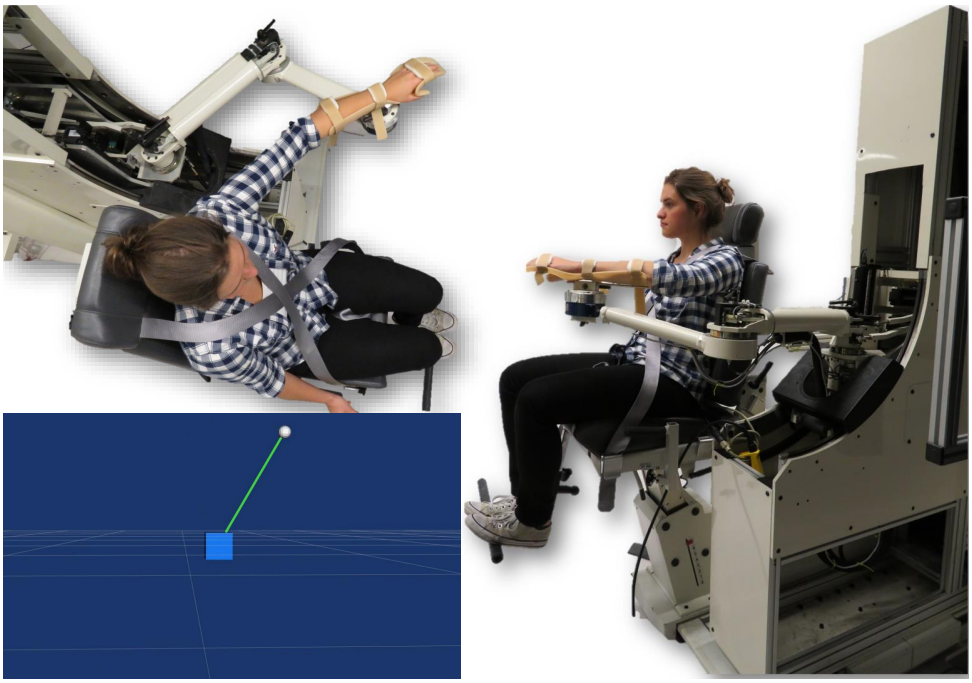

Figure 3.2. The New Arm Coordination Training 3D (NACT-3D) device provides haptic feedback in three dimensions to simulate a specified inertial model via admittance control. A force-torque sensor at the end-effector provides input to the admittance control loop. During this experiment, high stiffness virtual springs were used to restrict user motion in the z-direction while allowing them to move freely in the x-y plane. The display (bottom left) provided real-time visual state feedback of the cart-pendulum system.

reaching. While in use, the subject is seated in a Biodex chair connected to the base of the NACT-3D with their arm secured in forearm-wrist-hand orthosis. The NACT-3D is capable of exerting forces at this interaction point between the user and the robot in the x, y, and z directions only. The impedance control is updated at 1000 Hz.

The NACT-3D can move its end effector within a workspace defined both by its design limits (a radius of approximately 0.6m around the participants shoulder in the half plane in front of the participant's chest) and safety limits set by the investigators. The splint

can rotate passively but no torque can be exerted by the robot. At the point where the splint is mounted, a force-torque sensor measures the subject input which is fed back to the admittance controller. The peak push-pull force that can be exerted by the device of the device at the end effector is approximately 4.7 kN. The force measured at the end effector is sent to a host computer for use in the assistance algorithm to compare the user input to the control policy and perform the filter update at a rate of 60 Hz. In Figure 3.3, $f(s)$ is the subject control input which is used in the filtering algorithm. At start up, the haptic model is set such that the model of the end effector accounts for the mass of the subject's arm as well as an inertia parameter defined by the investigator.

During testing, a display provided real-time visual state feedback to the user about the cart-pendulum system s/he was trying to invert. High stiffness virtual springs in the haptic model were used to restrict user motion to a horizontal plane corresponding to the path of the cart in the virtual display. When user inputs met the criterion being used, they were accepted and the robot behaved according to the control scheme described in Figure 3.3. When user inputs did not meet the criterion for acceptance, the user input $f(s)$ was ignored by the admittance controller, such that the robot maintained its velocity at the time of rejection. Although the device was capable of replacing the user input with an input prescribed by an optimal controller, we chose to simply reject user actions. In this way, we provide feedback only by corrections without demonstrating or guiding the user in the correct action.

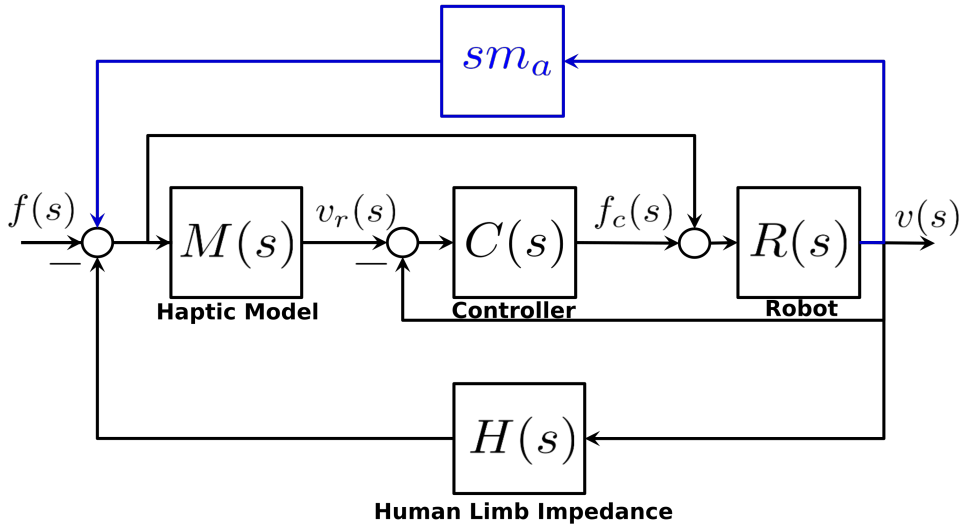

Figure 3.3. A voluntary force $f(s)$ is measured at the robot's end-effector using a six degree of freedom force-torque sensor (JR3) and passed through a model $M(s)$ that determines the velocity $v_r(s)$ the robot should move with. The reference velocity is tracked by the low level velocity controllers of each motor drive. The human also delivers involuntary impedance forces due to movement, given by dynamics transfer $H(s)$. Acceleration information is fed back as a pseudo-force for extra inertia reduction of the system.

3.4.4. Experimental Task

Users were tasked with controlling a simulated two-dimensional cart-pendulum system, which they were instructed to swing up to the unstable equilibrium (the system was initially resting at the downward stable equilibrium). The equations that describe the underactuated cart-pendulum system shown bottom left in Figure 3.2 are given by:

$$\dot{x} = f(x,u) = \begin{pmatrix} \dot{\theta} \\ \frac{g}{l}\sin\theta + u\cos\theta - \frac{b}{ml^2}\dot{\theta} \\ \dot{x}_c \\ u \end{pmatrix} \tag{3.4}$$

where the state vector x consists of the angular position and velocity of the pendulum and the position and lateral velocity of the cart, $x = [\theta, \dot{\theta}, x_c, \dot{x}_c]$, the input u is the lateral acceleration of the cart, g is the acceleration due to gravity, b is the damping coefficient, l is the pendulum length and m the mass at the tip.

Users kinematically controlled the cart acceleration (and thus position) by moving their arm from left to right in the horizontal plane subject to the constraints of the admittance controller outlined in Figure 3.3. To avoid confusion associated with conflating the task-related forces with forces generated by the assistance algorithm [95], no haptic feedback related to the system dynamics was displayed to the user during either nominal task execution or in addition to the assistance. In both the assisted and unassisted cases, users had to rely solely on visual state feedback to understand the system dynamics.

3.4.5. Sample Response

The mechanical filtering imposed by the robotic platform forces changes in the user input. Figure 3.4 shows a sample response of user inputs in assistance mode. Shortly before $t = 4$ s, we see an example of a rejected user action. Although the user input (gray) is a positive acceleration, the filtered input (red) is zero, and the velocity of the cart (green) is held constant. The optimal control signal (blue) indicated that a negative acceleration should be applied, but this was not used to replace the user input nor was it communicated to the user. At around $t = 3.5$ s, the user attempts a negative acceleration, and the prescribed optimal controller is also negative. Under the OCIP criterion, this action is allowed and the cart velocity decreases. This demonstrates how the mechanical filter can effectively yield to skilled users while assisting unskilled users.

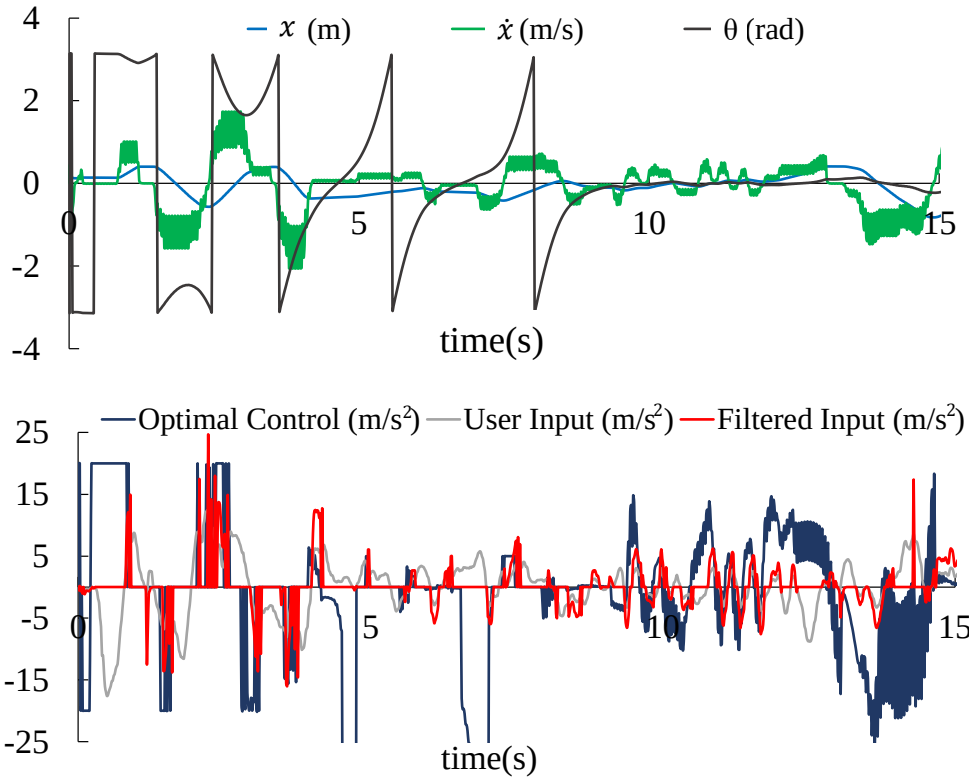

Figure 3.4. Sample response of a subject using the NACT-3D with the OCIP criterion. The NACT-3D is able to directly shape user input. We can see that even relatively large user inputs (gray) can be reduced to zero in the filtered input (red). Top: the states of the cart-pendulum system. The subject kinematically controls the cart position x_c (and $\dot{x}_c$) through the cart's lateral acceleration. We see the subject is able to stabilize the pendulum for 5 s. Bottom: The reference signal and user input used in (3.3) to generate the filtered input that drives the system.

Unlike the haptic stylus used in [32], the robotic platform used in the studies discussed in this chapter was capable of fully rejecting the physical motions of the subjects because of its underlying control architecture and sufficient actuation capabilities. While the haptic stylus, relied on users to interpret the feedback and correct their motion, the device described in Section 3.4.3 could actively correct motion while giving feedback and did not rely so heavily on the subjects interpretation of the haptic feedback. This allowed

us to update the mechanical filter at a higher rate (60 Hz-100 Hz) than in the previous implementation (10 Hz), which is part of why the improvements in performance are much greater on this device. In the trial shown in Figure 3.4, the user stabilizes the pendulum at the unstable equilibrium at $t = 9$ s and maintains that configuration for 5 s.

3.4.6. Experimental Protocol

Subjects used an upper limb robotic platform (NACT-3D) as an interface to control a simulated cart-pendulum system with state vector $x = [\theta, \dot{\theta}, x_c, \dot{x_c}]$ and horizontal acceleration of the cart as control input. During experimental trials, the user's goal was to invert the pendulum to its unstable equilibrium. User input was inferred from a force sensor at the robot's end-effector.

At the beginning of each session subjects were seated and secured in a Biodex chair and their left arm was secured in the orthosis on the NACT-3D (Figure 3.2). The system and task was demonstrated to them at the start of the testing using a video of a sample task completion. Subjects were instructed to attempt to swing up the pendulum to the upward unstable equilibrium and balance there for as long as possible. Subjects were instructed to continue to try to do this until the 30 second trial was over even if they succeeded at balancing near the equilibrium more than once. Depending on the study, subjects performed sets of 30 trials with short breaks in the same session or in sessions scheduled approximately one week apart as shown in Figure 3.5.

Subjects were recruited locally, and had to be healthy, able-bodied adults (in the age range of 18 to 50) with no prior history of upper limb or cognitive impairments. Only right-hand dominant participants were accepted into the study, and each subject

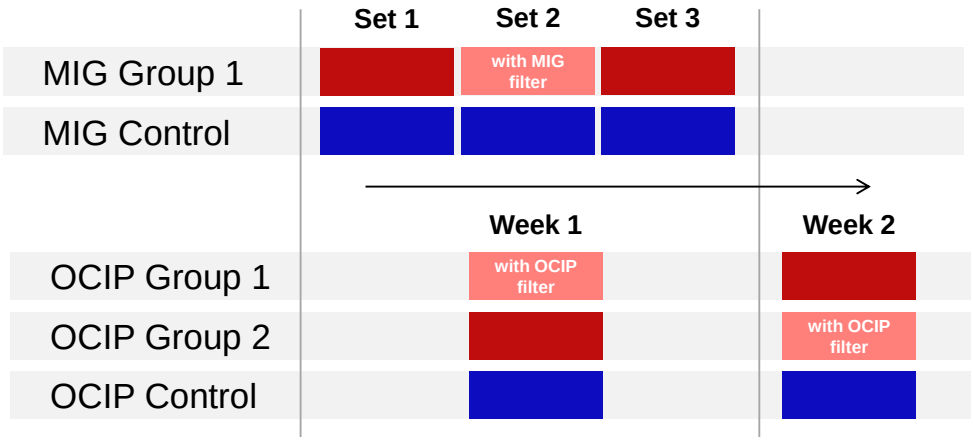

Figure 3.5. Each rectangle represents a set of 30 trials. MIG study participants completed all sets on the same day. OCIP participants completed sets one week apart.

performed the task with their left limb. All study protocols were reviewed and approved by the Northwestern University Institutional Review Board, and all subjects gave written informed consent prior to participation in the study.

3.4.6.1. MIG Study. Twenty-eight subjects (9 males and 19 females) consented to participate in the MIG study. All subjects in the MIG study completed three sets of thirty 30-second trials with short breaks between sets. Upon enrollment, subjects were randomly placed into either a control ($n = 10$) or training group ($n = 18$). During the second set, feedback in the form of a filter using the MIG criterion was engaged for the training group, while the control group completed each of the three sets without any feedback. Again, each user did three sets of thirty trials: set 1 (both groups: no feedback), set 2 (control: no feedback, training: feedback in the form of a mechanical filter using MIG), set 3 (both groups: no feedback).

3.4.6.2. OCIP Study. Fifty-three subjects (17 males, 36 females) consented to participate in this study. Each subject completed 2 sessions being approximately one week apart.

Upon enrollment in the study, each subject was placed into 1 of 3 groups. If placed in the training group ($n = 20$), the subject completed the first session with the OCIP filter and received no assistance in the second session. If a subject was placed in the non-training group ($n = 20$), they performed the task without assistance in the first session and used the OCIP filter in the second session. Finally, a control group ($n = 13$) performed the task without assistance in both the first and second session.

3.4.7. Performance Measures

The full state and user inputs were recorded in each trial and were used to calculate task-specific performance measures as well as more general measures such as error. The task-specific performance measures used to evaluate subjects in both studies is predicated on a notion of success. The definition of success that was used was based on the region of attraction for a linear quadratic regulator capable of stabilizing the system dynamics defined in the experiments. A trial was considered successful when a subject reached an angle of ± 0.15 rad and angular velocity of ± 0.6 rad/s. This definition of success was used to determine the time to success of the users in each experiment. In addition, if a subject was successful, the total time spent at the angle and angular velocity defined as success was recorded as the balance time. When users were successful multiple times in the same trial, time spent in the balance region was cumulative.

While these outcome-based measures provide clear indication about whether or not users could meet task goals, they neglect the behavior of users away from the goal state. Therefore, we use two measures—error and ergodicity—that use the full trajectory data to characterize task performance. The root mean square (RMS) error of each trajectory

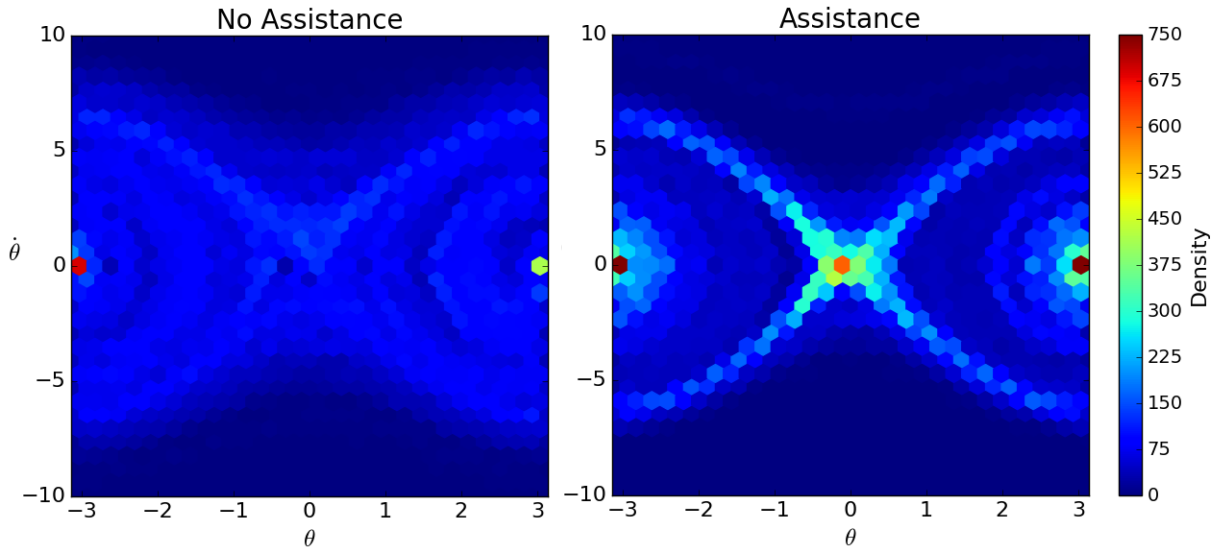

Figure 3.6. A histogram of all the trajectories recorded in the OCIP study demonstrates how the statistics of unassisted and assisted trajectories differ from one another. The histogram of unassisted trajectories (left) has its highest density at $\theta = \pm\pi$ which is the farthest point from the goal state. The rest of the distribution is diffuse over the state space. Although the histogram of the assisted trajectories (right) also has a high density at $\theta = \pm\pi$, the distribution is not as diffuse as that of the unassisted trajectories. There are bands of high density spreading outward form the goal state $(\theta, \dot{\theta}) = (0, 0)$. The spatial statistics of the assisted trajectories are more similar to the reference distribution, because there is a high density at and around the goal state. This outcome is captured by measuring the distance from ergodicity of the trajectories in each group with respect to the reference distribution.

generated by the users was calculated with respect to the desired position in an inverted unstable equilibrium (zero-vector of the states). RMS error was normalized by the RMS error of a constant trajectory at the stable equilibrium, equivalent to the error of the user not moving from the initial conditions. Finally, we also compared the experimental conditions through an analysis of the spatial distribution of trajectories that we observe under each condition. For instance, in the histogram of states recorded for all subject trajectories (Figure 3.6), one can see that trajectories in which subjects received assistance

have high density values near the goal state. To quantify the comparison of the distributions, we compute a metric on the ergodicity [79, 81] of each trajectory with respect to a Dirac delta $\delta(x - s)$ function centered at the unstable equilibrium $(\theta, \dot{\theta}) = (0, 0)$. The ergodic measure captures how well the time averaged statistics of the trajectory match the statistics of the reference distribution. The value of this metric was determined by calculating the weighted distance between the Fourier coefficients of the trajectory and those of the distribution. The ergodic metric gives us the distance from ergodicity, such that trajectories which were highly ergodic had lower ergodicity than those that were less ergodic.

The controller intervention was measured as the percent of rejected actions (PRA). PRA measured the fraction of user inputs that were rejected, where we defined an action to be a non-zero user input.

3.4.8. Statistical Analysis

The MIG experiment consisted of 30 baseline trials, 30 trials with or without the MIG filter, and 30 trials post-training for a total of 90 trials. These were grouped into blocks of 5 trials to evaluate subject performance over time. The analysis consisted of two-factor (block and group) repeated measures ANOVA tests, using the baseline and post-training data only. The ANOVA's were used to compare the effect of the MIG filter and unassisted practice on each of the performance measures. Trials from set 2 are removed from the analysis to avoid including the assistance itself as a factor in the experiment. In the OCIP study, subjects trained with the filter received no prior exposure to the task without

assistance. Student's t-tests were used to evaluate the difference between the week 2 performance of the trained group and the week 2 performance of the control group.

The relevant features of the hybrid shared controller were evaluated statistically. First, the ability of the shared controller to assist subjects in completing the task was tested in each study. In the MIG study, this was done by comparing the experimental group to controls with an equivalent amount of practice using a two-sample t-test. The effect of the OCIP criterion as an assistive controller was tested in a counter-balanced fashion using paired two-sample t-tests on all performance metrics. Second, the sensitivity of the shared controller to the initial skill of the users was evaluated by performing Peason's R correlation tests between the level of controller intervention and the performance of users in their first set of unassisted trials. Finally, the assist-as-needed feature of the shared controller was shown by testing the correlation between the level of controller intervention and the current performance of subjects.

3.5. Results

The results were reported as follows. First, the training effect of each study was statistically tested in Section 3.5.1. The results demonstrated that training with the hyrbid shared controller increased subject performance in later trials within the same session (Section 3.5.1.1) and in a session one week after training (Section 3.5.1.2). An analysis of the hyrbid shared controller was performed to test for three characteristics of effective pHRI. In Section 3.5.2, the performance improvement made while the criterion was engaged was evaluated in both the MIG study and the OCIP study. In Section 3.5.3 and 3.5.4, the correlation of the percent of rejected actions with the initial skill and current

performance are reported to evaluate the sensitivity of the shared controller to user skill and its ability to assist-as-needed, respectively. In each section the relevant statistics are reported first, followed by a summary and interpretation of the results.

3.5.1. Training Effect

The effectiveness of the filter as a training tool was assessed in both experiments. In the MIG study, we consider only skill acquisition within a single session. We assess the retention of skill over the course of one week in the OCIP study.

3.5.1.1. MIG Study: Skill Acquisition. Two-factor repeated measures ANOVAs were used to assess the effects of the group (between-subjects) and set (within-subjects) on all performance measures listed in Section 3.4.7. The training group and control group were evaluated based on the baseline trials (set 1) and the post-training trials (set 3) only. Set 2 was left out of the ANOVA, so that effects of the assistance itself would not be measured in the analysis. In order to assess how subject performance evolved over time, the baseline and post-training sets were analyzed using blocks containing five individual trials. Therefore, there were 6 blocks in each set as shown in Figure 3.7.

The factorial ANOVA of the balance time revealed that block was the only significant factor ($p < 2 \times 10^{-16},\ F(11, 286) = 10.775$). The main effect of group and interaction effect of group and block were not significant for balance time ($p > 0.05$). When an analysis of variance was performed on the time to success, again, the main effect of block was significant ($p = 3.81 \times 10^{-15},\ F(11, 286) = 9.848$) and the main effect of group was not significant ($p = 0.533,\ F(1, 25) = 0.399$). However, the interaction effect of group and block was significant ($p = 0.0135,\ F(11, 286) = 2.222$). The control and trained

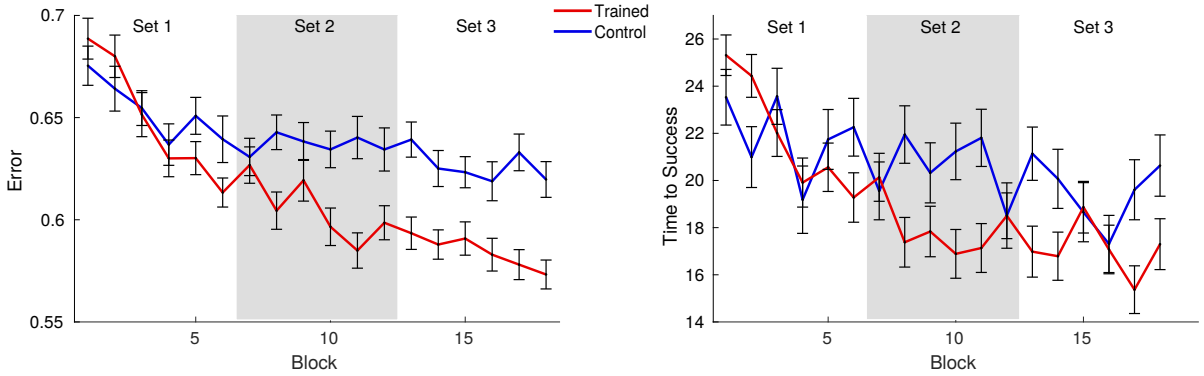

Figure 3.7. The MIG study showed that subjects improved with practice in all sets regardless of training group, however, there was a significant interaction effect between training group and block when ANOVAs were applied to three of the four performance metrics. This suggests that although subjects in each group started around the same performance level, the trained group attained a higher level of performance than the the control group during the post-training trials. Note that the set 2 performance (gray) was not included in the ANOVA to avoid measuring effects of the assistance itself.

group performed similarly in the baseline trials. The time to success decreased even before the training set (Figure 3.7). However, the control group essentially plateaued during the training set and saw large fluctuations in the time to success during the post-training trials. The time to success of the trained group decreased during training and was maintained in the post-training trials.

The group also was not a significant factor affecting the RMS error ($p = 0.223,\ F(1, 25) = 1.560$), but main effect of subset ($p < 2 \times 10^{-16},\ F(11, 286) = 20.620$) and the interaction of group and subset ($p = 0.004,\ F(11, 286) = 2.575$) were significant. When the error of the control group and trained group was plotted over time (Figure 3.7), the control group error decreased initally but leveled off. The error of the trained group continued to decrease during training and in the post-training trials.

When the distributions of the trajectories were compared using the ergodic metric, the significant factors were the subset ($p < 2 \times 10^{-16},\ F(11, 286) = 18.311$) and the interaction between group and subset ($p = 0.030,\ F(11, 286) = 1.983$). The main effect of group was not significant ($p = 0.294,\ F(1, 25) = 1.151$). The progress of the ergodic metric over time was similar to that of the RMS error.

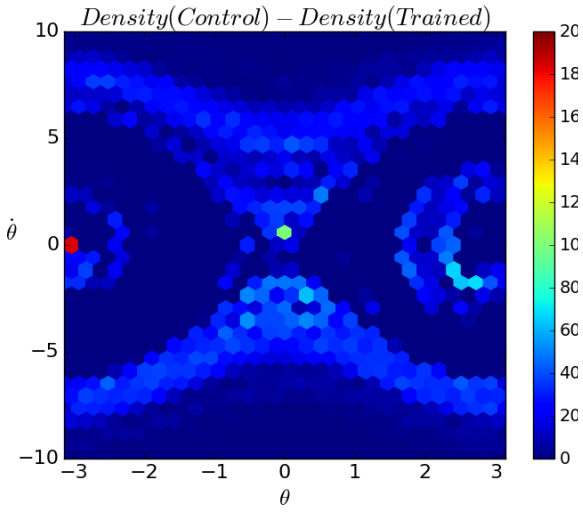

(a) The density function of trained group trajectories subtracted from the control trajectories density.

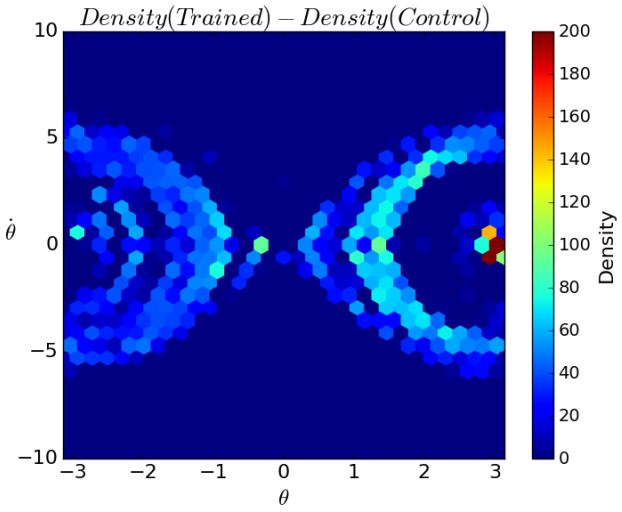

(b) Control trajectories density function subtracted from the post-training trajectories density.

Figure 3.8. Trajectories from Week 2 of the OCIP study showed that subjects who trained with the hyrbid shared controller spent more time near the goals state $(\theta, \dot{\theta}) = (0, 0)$ than subjects who practiced unassisted. On the left, the week 2 control trajectories have higher densities than the post-training trajectories at higher angular velocities as well as in bands near $\theta = \pm\pi$ which is the farthest angle from the goal state. The control trajectories also spend time near the goal state, but to a lesser extent. On the right, the trained trajectories also have high density near $\theta = \pm\pi$, but there are large bands of high density in the region $-1.5 \leq \theta \leq 1.5$ and $-4 \leq \dot{\theta} \leq 4$. This suggests that the trained group's motions were more consistent with the task goal—making the statistics of the trained group closer to the spatial statistics of the reference Dirac delta distribution, so the ergodic measure of the trained group is lower than that of the controls.

The results of the ANOVA of each of the performance measures showed that subset was a significant factor—implying that regardless of the training in set 2, all subjects performed better in later sets than in their initial sets. The significant interaction effect observed in three out of the four metrics demonstrates that while the subjects started at the same performance level, subjects in the trained group attained a higher performance level than the control group.

3.5.1.2. OCIP Study: Short-term Retention. The effect of training was assessed by comparing the week 2 session of the trained group to the week 2 session of the control group. The two groups were not significantly different in terms of the task-specific measures of success. However, the trained group had significantly lower RMS error, and the distributions of the trained group's trajectories were more similar to the reference distribution, resulting in a much lower ergodic measure than the control group. A two-sample t-test was performed on the task specific performance measures, finding no difference between trained group and untrained group in terms of their time spent balanced ($p = 0.1687, t(988) = 1.378$) and time to success ($p = 0.1935, t(988) = 1.301$). The two-sample t-test of the RMS error showed a significant difference between the trained ($mean = 0.621,\ SD = 0.058$) and control ($mean = 0.629,\ SD = 0.061$) groups ($p = 0.0499, t(988) = -1.963$). The t-test of the ergodic metric also showed a significance difference ($p = 2.266 \times 10^{-4}, t(988) = -3.701$) between the trained group ($mean = 0.705,\ SD = 0.177$) and the control group ($mean = 0.751,\ SD = 0.207$). Although subjects who trained with the OCIP criterion were not successful more often than the control group, they did spend a higher proportion of their time near the goal state as can be seen by the histogram of their trajectories shown in Figure 3.8. *These results*

suggest that subjects learned more and retained that skill one week after training when they trained with assistance rather than simply practicing the task unassisted.

The progress of the two groups over the second session (Figure 3.9) was analyzed further by performing mixed design ANOVAs on the training group (between participants) and block (within participants) using all four measures.

The balance time of the control group and the trained group in the second session was analyzed with a 2 (training groups) x 6 (blocks) mixed design ANOVA, which showed no significant main effects or interactions effects. The main effect of training group was not significant $F(1, 31) = 1.202,\ MSE = 1.25,\ p = 0.28,\ Cohen'sf = 0.08$. The main effect of block also was not significant $F(5, 155) = 2.018,\ MSE = 0.44,\ p = 0.079,\ Cohen'sf = 0.11$, nor was the interaction of training and block significant $F(5, 155) = 1.05,\ MSE = 0.23,\ p = 0.39,\ Cohen'sf = 0.08$.

The mixed design 2 x 6 ANOVA design was also applied to the time to success, and the main effect of training group was not significant $F(1, 31) = 0.334,\ MSE = 103.4,\ p = 0.567,\ Cohen'sf = 0.05$. The main effect of block was not significant either $F(5, 155) = 1.34,\ MSE = 66.32,\ p = 0.25,\ Cohen'sf = 0.09$. The interaction effect of block and training group also was not significant $F(5, 155) = 1.34,\ MSE = 66.50,\ p = 0.25,\ Cohen'sf = 0.09$.

The same mixed design ANOVA was used to analyze the RMS error in each trial. The main effect of block was significant $F(5, 155) = 4.336,\ MSE = 0.011,\ p = 0.001,\ Cohen'sf = 0.19$, but the main effect of training was not significant $F(1, 31) = 0.76,\ MSE = 0.035,\ p = 0.39,\ Cohen'sf = 0.15$. The interaction effect of training group and block also was not significant $F(5, 155) = 1.61,\ MSE = 0.004,\ p = 0.16,\ Cohen'sf = 0.12$.

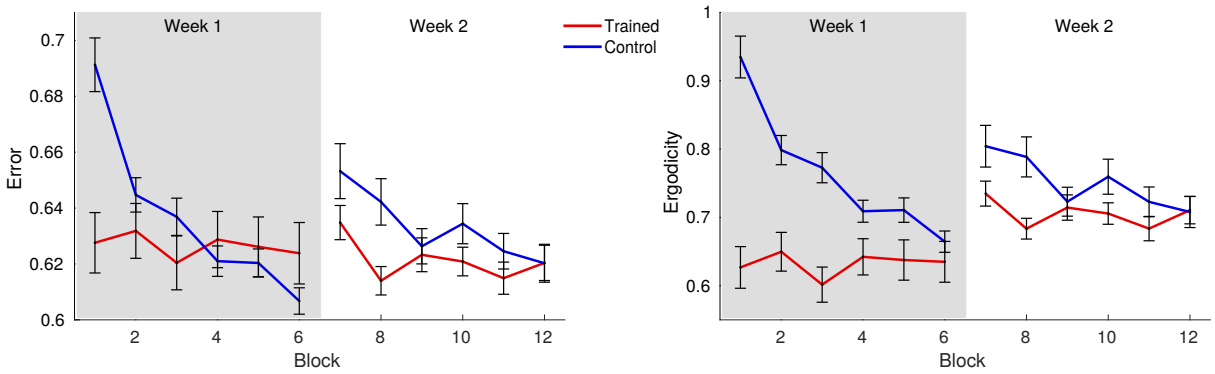

Figure 3.9. The results of the OCIP study demonstrate that subjects trained in week 1 retain high performance levels in week 2 as measured by RMS error and ergodicity. In the first 2 blocks of trials, the error and ergodicity of the control group are higher than that of the trained group. The trained group retains their initial performance level, while the control group continues to improve—eventually reaching the same level of performance as the trained group. It appears the feedback helped with retention because the learning was more structured. Note that the performance measures in week 1 (gray) were not used in the statistical analysis to avoid measuring the effects of the assistance itself.

The analysis of the ergodic metric using the mixed design ANOVA revealed a significant main effect of block $F(5, 155) = 2.88,\ MSE = 0.08,\ p = 0.0163,\ Cohen'sf = 0.15$, and a significant interaction effect of block and training group $F(5, 155) = 2.33,\ MSE = 0.06,\ p = 0.045,\ Cohen'sf = 0.14$. The main effect of training was not significant $F(1, 31) = 1.056,\ MSE = 0.49,\ p = 0.312,\ Cohen'sf = 0.17$.

In Figure 3.9, the control group performed worse at the beginning of the second session that it did at the end of the first session, and their performance increased in terms of error over the course of the session. The trained group also improves moderately during the second session. The ANOVA of the ergodic metric is also able to detect the significant improvement during the second session by the control group as well as the interaction effect of group and training. This interaction is a result of the trained group performing better

under the ergodic metric at the beginning of the second session and maintaining that performance, while the control group eventually reached the same level of performance. *Training with the OCIP criterion in week 1 speeds learning, and skill is retained after one week though the improvements due to unassisted practice are not retained.*

3.5.2. Task-based Assistance

We evaluate the ability of the hybrid shared controller to assist subjects in completing the task while it is engaged. In the MIG study, we compare the the control group to the group recieving assistance during their second set of trials. In the OCIP study, the order in which subjects received assistance was counterbalanced, such that subject performance in the assisted session was compared to the same subject's performance in the unassisted session.

3.5.2.1. MIG Study. Comparisons between the control and experimental groups are shown in Figure 3.10. Two-sample t-tests showed that there was no significant difference between the control group ($n = 10$) and experimental group ($n = 18$) baseline performance in terms of their balance time ($p = 0.178,\ t(793.22) = 1.35$), time to success ($p = 0.9497,\ t(644.23) = 0.063$), error ($p = 0.411,\ t = 749.28 = -0.822$), or ergodicity ($p = 0.507,\ t(711.17) = -0.6631$). During the training set (set 2), the experimental group ($mean = 2.36,\ SD = 3.47$) maintained the pendulum in the balanced position for significantly longer ($p = 7.674 \times 10^{-8},\ t(832.55) = 5.42$) than the control group ($mean = 1.37,\ SD = 1.78$). The group receiving assistance also reached the balance position more quickly than the group practicing the task without assistance ($p =$

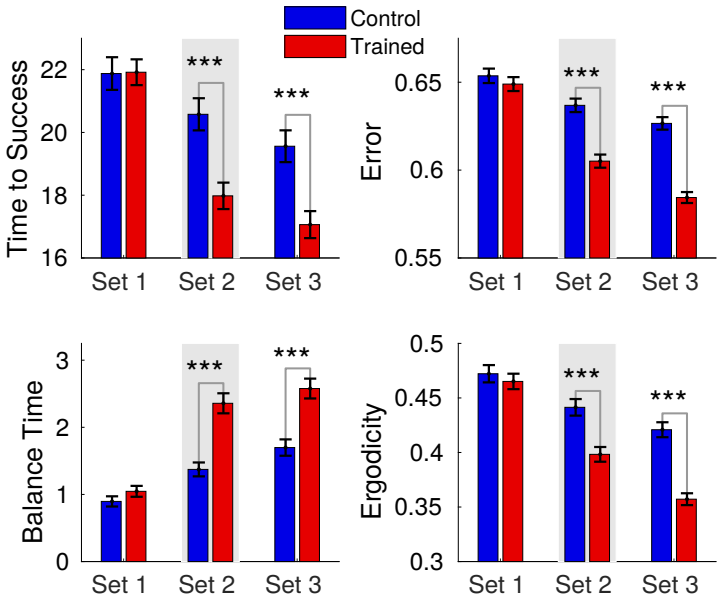

Figure 3.10. The MIG filter study demonstrated that the filter successfully assisted subjects in set 2 compared to controls. Moreover, trained subjects outperformed the control group in set 3. Note: error bars indicate standard error; significance is indicated by $^{*}p < 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$.

9.87×10^{-5}, $t(666.93) = -3.9174$), so the experimental group ($mean = 17.98$, $SD = 9.79$) had a lower time to success than the control group ($mean = 20.58$, $SD = 8.84$).

The RMS error of the experimental group ($mean = 0.605$, $SD = 0.087$) was also significantly lower ($t(753.59) = -5.925$, $p = 4.738 \times 10^{-9}$) than that of the control group ($mean = 0.636$, $SD = 0.066$). Finally a comparison of the trajectory distributions of the experimental group in terms of ergodicity ($mean = 0.398$, $SD = 0.157$) to the distributions of the control group ($mean = 0.441$, $SD = 0.131$) showed that the filter was effective able to effectively assist subjects in the task ($t(707.63) = -4.2435$ $p = 2.494 \times 10^{-5}$).

The two experimental groups performed similarly in their baseline trials, but in set 2, the group using the filter outperformed the control group in terms of balance time, time

to success, RMS error, and the ergodic metric. *This demonstrates that the hyrbid shared controller using the MIG criterion meets the basic requirement of assisting subjects with the task while in use.*

3.5.2.2. OCIP Study. In the study of the OCIP criterion, subjects were randomly placed into either a group who used the shared controller in the 1st session ($n = 20$) or a group who used the shared controller in the second session ($n = 20$). Therefore, the ability of the hybrid shared controller to provide assistance was tested in a counterbalanced fashion. Pairwise student's t-test were used to compare performance with and without the assistance of the filter on a subject by subject basis. Subjects did not have significantly lower error ($t(1199) = -1.674,\ p = 0.0949$) when using the OCIP filter ($mean = 0.626,\ SD = 0.102$) compared to unassisted trials ($mean = 0.632,\ SD = 0.062$). Under all other metrics, subjects performed better on the day that they used the OCIP filter compared to their performance on the day they performed the task without assistance ($p < 10^{-14}$) as shown in Table 3.1. *These results showed that the shared controller with the OCIP criterion was able to help subjects complete the task more frequently.*

3.5.3. Hybrid Shared Control Adapts to Initial Skill

We previously reported that there was a relationship between participant skill level—estimated based on performance in unassisted trials—and the frequency of controller intervention in the MIG filter mode in [50]. In that case, we calculated the success rate of the 30 trials from set 1 to approximate user skill level. We then used *Percent of Rejected Actions* (PRA) values from individual trials in set 2 from the same users to identify the correlation.

	No Assistance		Assistance					
			$n = 40$					
Measure	μ	SD	μ	SD	$\mathbf{\Delta}\mu$	t	df	p
Success Rate	0.348	0.163	0.792	0.182	$\mathbf{0.444}^{***}$	12.314	39	5.162×10^{-15}
Balance Time	0.191	0.411	1.661	1.913	$\mathbf{1.47}^{***}$	26.519	1199	2.541×10^{-122}
Time to Success	25.333	7.648	20.068	7.824	$\mathbf{-5.265}^{***}$	-17.202	1199	1.926×10^{-59}
Error	0.632	0.062	0.626	0.102	$\mathbf{-0.006}$	-1.674	1199	9.477×10^{-2}
Ergodicity	0.739	0.191	0.631	0.283	$\mathbf{-0.108}^{***}$	-11.261	1199	4.954×10^{-28}

Table 3.1. The OCIP filter assisted subjects in completing the task more frequently and at a higher level of performance in four out of five measures when subjects were randomly assigned to use the filter in either the first of second session. Paired two-sample t-tests were performed in *R* [99] comparing the unassisted and assisted trials of the 20 subjects receiving subjects in the first session and the 20 subjects receiving assistance in the second session. Significant differences in means are indicated by $^{*}p < 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$. Note that the degree of freedom (df) for success rate is 39 since there is only one rate per subject.

The PRA of OCIP filter had a moderate correlation to the initial skill of the subjects under all of our performance measures. We evaluated the correlation between initial skill of the untrained group ($n = 20$) who received no assistance in week 1 and the PRA in individual trials of the group when they did not receive assistance from the filter in week 2. We found that there is again a significant correlation between the initial skill of the users as measured by the success rate and mean performance measures in week 1 and the PRA of those subjects in week 2. In this case, the correlation coefficients, shown in Table 3.2, were slightly higher, indicating a moderate correlation between the subject's initial performance and the filter's response to their inputs. The correlations of each performance metric matched the expected sign corresponding to a decrease in PRA in response to an increase in the user's initial skill. *Although the hybrid shared controller is*

Measure	Test Sign	**r**	p
Success Rate	$-$	$\mathbf{-0.235}$	5.898×10^{-9}
Balance Time	$-$	$\mathbf{-0.427}$	$< 2.2 \times 10^{-16}$
Time to Success	$+$	$\mathbf{0.2444}$	1.2898×10^{-9}
Error	$+$	$\mathbf{0.302}$	4.308×10^{-14}
Ergodicity	$+$	$\mathbf{0.282}$	2.078×10^{-12}

Table 3.2. There were moderate correlations between the initial skill of the user and PRA of the OCIP filter in all measures ($p < 0.05$). Pearson's correlation tests were performed in R [99] by applying a linear model to the mean of performance metrics first session and percent of action rejected by the OCIP filter. The expected sign of the correlation coefficient (r) for a shared control scheme that is sensitive to the initial skill of the user is indicated in the column on the right.

not tailored to either high skill or low skill users, it adapts to user skill level and could be appropriate for both novices and expert users.

3.5.4. Hybrid shared control Assists-As-Needed

In addition to testing the relationship between the initial skill of the user and the level of controller intervention, the responsiveness of the controller to user performance in the current trial was tested using Pearson's product-moment correlation. There were high significant correlations between user performance within a single trial and the PRA in that trial. These correlations and significance values are reported in Table 3.3. The test sign indicated in the table indicates the expected sign of the correlation coefficient when the controller accepts more user inputs in response to high user performance. Under each metric, the correlation meets this expectation. *This demonstrates that the robotic assistance adapts in real-time to the needs of the users without including high-level performance heuristics to tune the relative contributions of the human and the robot.*

Measure	Test Sign	**r**	p
Balance Time	$-$	$\mathbf{-0.616}$	$< 2.2 \times 10^{-16}$
Time to Success	$+$	$\mathbf{0.602}$	$< 2.2 \times 10^{-16}$
Error	$+$	$\mathbf{0.677}$	$< 2.2 \times 10^{-16}$
Ergodicity	$+$	$\mathbf{0.706}$	$< 2.2 \times 10^{-16}$

Table 3.3. The PRA of the OCIP filter was highly correlated with the current performance of the users under all measures ($p < 0.05$). Pearson's correlation tests were performed by applying a linear model to the performance measures in each trial in the OCIP study and the PRA in the same trials. The expected sign of the correlation coefficient (r) for a shared control scheme that is sensitive to the performance of the user is indicated in the column on the right.

3.6. Discussion

Despite the breadth of research, there are relatively few instances where physical human robot interaction has been significantly more effective than unassisted practice or human-mediated training. In the work presented here, experimental results demonstrate that our implementation of a task-based hybrid shared control paradigm enhances the effect of training compared to unassisted practice. On average, subjects who trained with our robotic feedback improved significantly more than subjects who trained with an equivalent amount of unassisted practice. Based on analysis of the spatial statistics of the post-training trajectories, the training group was capable of more controlled movement with significantly more time spent near the goal state. Moreover, subjects who trained with the proposed MIG shared control scheme continued to improve even after the assistance was removed, while members of the control group plateaued in their performance. Finally, through our studies, we observed that subjects both experienced immediate improvement from training with feedback and exhibited short-term retention

of the acquired skill. These results demonstrate that the proposed hybrid shared control paradigm enhances task learning through forceful interaction.

In order to understand why the algorithm was effective, we examine the unique characteristics of the hybrid shared control paradigm as well as qualities that coincide with existing best practices in robotic training. Reviewing the motor learning literature, several features of pHRI can be identified to lead to effective training. For one, a necessary condition for effective training through forceful interaction is that the automation should be able to assist subjects in completing a task while assistance is engaged. In our experimental results, we show that the hybrid shared control paradigm is capable of improving success in accomplishing a dynamic task during the trials in which it was engaged. In the MIG study in set 2, subjects performed better across all metrics when assistance was engaged, even though on average they started off at the same skill level in set 1. Similarly, the subjects in the OCIP study performed better with assistance compared to their own unassisted trials.

Secondly, interfaces should avoid user passivity and require substantial user effort. This is inherent to our algorithm because the hybrid controller never actively assists with task completion by only rejecting, but not replacing, incorrect actions. As a result, users are allowed to fail at the task and when they succeed, they succeed through their own actions. While impedance-based assist-as-needed controllers can interfere less based on performance heuristics, impedance control is based on desired velocity profiles rather than the task goal. The hybrid shared control paradigm discussed in this chapter uses a task-based criterion in order to measure whether or not it is needed. This allows the controller

to effectively get out of the way when users are progressing towards the task goal on their own—maximizing their effort.

Building on the principle of requiring effort from the patient, shared control paradigms have been shown to be more effective when they adapt the level of assistance over time, assisting only as much as is necessary. The need for modulating the level of assistance can be due to two factors: (1) differing initial user skill level and (2) varying user performance over time. Users are expected to progress in their training over time. However, it is not enough for the level of assistance to decrease over time or after a certain performance target had been reached—there are cases, where subjects fatigue or become less engaged if the task is too difficult, so interfaces must be able to adjust both up and down in response to the automation's current assessment of the user. In our results, we show that the proposed shared control paradigm adapts to user initial skill and exhibits properties of an 'assist-as-needed' controller, reducing or increasing its intervention according to user performance in real-time. In future studies, it would be interesting to explicitly assess fatigue in between or during trials. In this way, we could adjust assistance based on current levels of fatigue and/or control for the effects of fatigue in study outcomes.

All in all, we present here a hybrid shared control paradigm that significantly improves task learning. We use a task-based criterion to discretely switch between full user control and full rejection of user control, which allows us to synthesize an interface with characteristics important for motor learning. Experimental data confirms that the shared control scheme exhibits these characteristics.

We also found that within a single session, trained subjects attained a higher level of performance than their counterparts who practiced unassisted. Yet at the end of

the second session in week 2, control subjects reached the same level of performance as the trained group. This is likely due to the difference in when the hybrid shared control was introduced, and indicates an opportunity to explore the scheduling of assisted and unassisted practice over the course of a training regimen. In future work, we plan to test subjects in higher-dimensional tasks and make comparisons to other assist-as-needed controllers, such as path controllers, active constraints, and other impedance-based approaches. In addition, we are exploring ways to define more complex tasks where it may be difficult to define a desired trajectory or goal state.

3.7. Conclusion

Numerous devices and control strategies have been developed to facilitate forceful interaction between humans and robots for the purposes of training specific skills or tasks. However, it is difficult to show the efficacy of these robots in promoting skill learning. Some types of robot-mediated training may be detrimental to learning, and others might be no more effective than an equivalent amount of unassisted practice. Interfaces for pHRI that have been shown to successfully enhance training have several features explicitly included in their design to enhance motor learning. Specifically, the automation must be able to assist users in completing the task and adapt the assistance to the needs of the individual user in terms of both initial skill and current performance in order to promote user engagement.

In this work, we investigate the use of a hybrid shared control method for assistance and training. The interface allowed subjects to make errors and even fail at the task. While the application of the filter improved subject success rates, it did not make subjects

successful all of the time. It also avoided enforcing a specific trajectory by evaluating the effect of user inputs on a continuous basis. Results from two user studies with different task-based acceptance criteria demonstrate the method's effectiveness in both assistance and training. Analysis of the correlations between the level of controller engagement and the initial skill of the users showed that the filter is sensitive to users' skill level. While the filter inherently adapts with every measurement of the user inputs, the strong correlation between performance measures and the level of controller intervention shows that this instantaneous adaptation results in a controller that also assists as needed according to the performance of the user in an individual trial.

CHAPTER 4

Ergodicity Reveals Assistance and Learning from Physical Human-Robot Interaction

This chapter applies information theoretic principles to the investigation of physical human-robot interaction. Drawing from the study of human perception and neural encoding, information theoretic approaches offer a perspective that enables quantitatively interpreting the body as an information channel, and bodily motion as an information-carrying signal. We show that ergodicity, which can be interpreted as the degree to which a trajectory encodes information about a task, correctly predicts changes due to reduction of a person's existing deficit or the addition of algorithmic assistance. The measure also captures changes from training with robotic assistance. Other common measures for assessment failed to capture at least one of these effects. This information-based interpretation of motion can be applied broadly, in the evaluation and design of human-machine interactions, in learning by demonstration paradigms, or in human motion analysis.

4.1. Introduction

Hundreds of devices have been designed and built to facilitate forceful interactions between humans and an autonomous system for the purposes of training, safe collaboration, and physical task assistance [17, 38, 83]. The goal of these systems is to augment the capabilities of the human either by providing feedback that enhances the training of a person in a certain task or by eliminating an existing deficit, such as weakness or

discoordination due to a neuromotor pathology. As a result, these robotic systems have unique requirements for sensing, actuation, and algorithmic design. In particular, the algorithmic component must be able to infer the quality of the measured behaviors or tasks performed by the joint human-robot system, implying the use of an appropriate metric as the basis for evaluating and modulating the human-robot interaction. However, it is unclear what metrics are appropriate for automating human-machine collaboration. Should one use measures from traditional robotic control such as trajectory error, biologically relevant measures (e.g., energy), or task-specific measures of motion quality? The choice of metric has implications beyond modulation of the interaction, including evaluation of the effectiveness of the physical interaction between the human and robot. One of the purposes of the interaction should be discernible improvements in performance from the assistance and learning from the assistance. In this chapter, we show that ergodicity—a measure of the task information encoded by a movement—can predict the presence of robotic assistance and detect the training effect of assistance.

In the human body, sensory information and motor commands are transmitted by nerve fibers conducting action potentials from one synapse to the next. Information theory provides a means to measure the information contained in such signals and to characterize the communication channel [117]. Part of the difficulty of analyzing neural signals are the idiosyncratic sources of variability, but applying information theoretic principles to the nervous system has allowed us to understand and analyze neural coding and organization [7, 86, 98] as well as cognitive perception [118, 143, 145]. Here we provide evidence that the motions resulting from neuromotor signals can be understood

as information-carrying signals themselves, and that information measures can also be used to analyze the movement and predict features of neuromotor control.

While signals analysis has been broadly applied to study the transfer of information from stimuli to cortex, principles from information theory are rarely used to model the ouput of motor commands. Instead theories of motor coordination are often developed on the basis of constrained optimization, often substituting the behavioral goal with the minimization of a measured quantity such as error or energy [14, 126]. These are useful metrics both because they allow us to reason about the underlying principles of neuromotor control and because many well-developed engineering techniques are based on minimizing these quantities [29, 77, 89, 103]. Therefore one can characterize human-like walking as an energy-minimizing trajectory [14] or reaching in the upper limb as a minimum jerk movement [33, 41]. When generic measures fail to capture important features of motion, they are supplemented with qualitative analysis (e.g., similarity to normal patterns of motion) [14] and task-specific measures of success such as work area [24], movement speed [48], or a combination of velocity threshold, aim, and maximum reach [58].

One of the reasons task-specific or outcome-based measures capture the qualitative description of the behavioral goal is that they are independent of the motion strategy, whereas energy or error is typically explicitly dependent on a specific desired trajectory. For instance, one might travel forward and then to the right to a reach a target, or one could follow a diagonal path, resulting in the same level of task success using two disparate strategies. Even if the average of the two paths were used as a reference, the resulting desired trajectory may not convey the task goal, and the variance and other statistics of the set of motion strategies are not part of a typical control architecture. Stereotypical

motion—such as reaching, self-feeding, and walking—have substantial variation between equally qualitatively successful trials, both within and between individuals. Because of the inherent stochasticity in neuromotor commands and the resulting task executions, we use a distribution $\phi(x) : \mathbb{R}^n \mapsto \mathbb{R}$ over a state space $\mathcal{X}$ to define a task goal. We assess a motion by asking *how much information about* $\phi(x) \in \mathbb{R}, x \in \mathcal{X}$ *is encoded in the movement* $x(t) : \mathbb{R} \mapsto \mathcal{X}$ *where* $x(0) = x_0$ *and* $x(t) \in \mathcal{X}$ *for some time* t.

There are a few natural ways to describe a task statistically rather than by specifying a goal state or a goal trajectory. If there is a particular goal state s, one can represent this as a Dirac delta function $\delta(x - s)$. Or, if the task definition is a consequence of measuring many instances of task execution, the collection of observations will form a distribution $\phi(x)$ in the domain $\mathcal{X}$. As more demonstrations of the target reaching task are added to the set of observations, the collective time spent at the goal state generates a higher peak at the state s, asymptotically approaching a delta function at the goal state. To quantify information content in a motion, one needs to measure a trajectory $x(t) \in \mathcal{X}$ that describes movement of the body. If $x(t)$ were itself a distribution across all of $\mathcal{X}$, one could use the Kullback-Leibler divergence D_{KL} [61] to measure how well $x(t)$ communicates information about $\phi(x)$. However, $x(t)$ is a trajectory, taking on only one state at each time t, and as a consequence D_{KL} between $x(t)$ and $\phi(x)$ will be generally infinite.

To compare a trajectory $x(t)$ to a distribution $\phi(x)$ while avoiding the underlying problems with using D_{KL}, we use ergodicity—which relates the temporal behavior of a signal to a distribution. A trajectory $x(t)$ is *ergodic* with respect to a distribution $\phi(x)$ if, for every neighborhood $\mathcal{N} \subset \mathcal{X}$, the amount of time $x(t)$ spends in $\mathcal{N}$ is proportional

to the measure of $\mathcal{N}$ provided by $\phi(x)$. On a long enough time horizon, measuring a perfectly ergodic $x(t)$ gives a complete description of $\phi(x)$. However, since a trajectory can only visit every neighborhood of $\mathcal{X}$ on an infinite time horizon, a finite time horizon $x(t)$ cannot be perfectly ergodic. Instead, we can ask that $x(t)$ be *maximally ergodic*, by introducing a metric on ergodicity, so that the *time-averaged* statistics of $x(t)$ best capture the statistics of $\phi(x)$ in a specified time horizon T, subject to system dynamics and constraints. Ergodicity can be measured by several metrics [114, 115]; here we use the spectral approach [79], which characterizes ergodicity by comparing spatial Fourier coefficients of $x(t)$ to coefficients of $\phi(x)$—giving us the distance from ergodicity. In Fig. 4.1, we show two hypothetical cases of the trajectory of the center of mass during walking compared to an idealized reference trajectory based on typical gait patterns. The high quality execution is not temporally aligned with the reference and may represent faster or slower walking than the reference. Nonetheless, the time-averaged statistics of the trajectory match that of the reference distribution. The low quality execution provides an example of what one might obtain from an impaired individual with poor balance or motor coordination. The ergodic metric used here gives us the distance from ergodicity, such that trajectories which are highly ergodic, like the high quality execution in Fig. 4.1, have a lower ergodicity than those that are less ergodic.

We use this information measure to analyze two cases of assisted motion—where the lack of assistance may be interpreted as a deficit relative to the assisted condition. First, we looked at data in supported reaching for a participant with an abnormal tendency to flex the elbow when lifting the arm at the shoulder, where we can see in Fig. 4.2 that

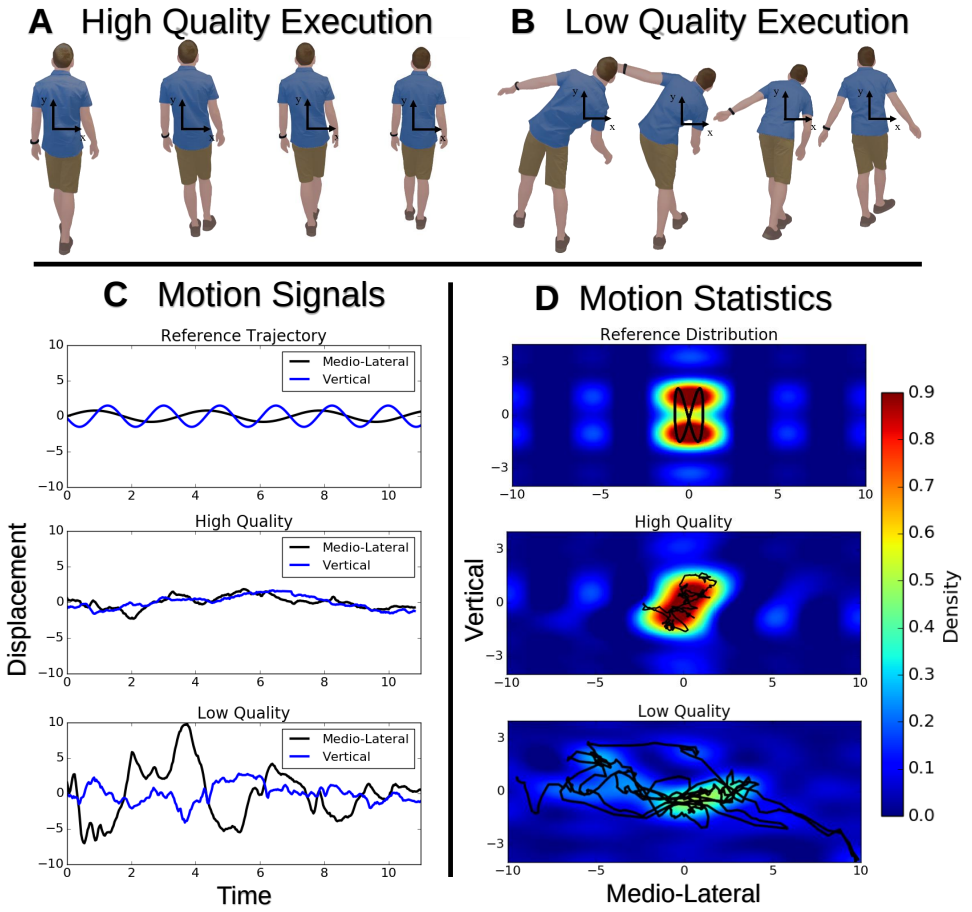

Figure 4.1. **Illustration of Motion Signals and Statistics using the Center of Mass in Walking.** For the task of walking on a line, we can distinguish between two hypothetical cases—a high quality execution (A) and a low quality execution (B) by tracking the vertical and medio-lateral displacement of the person's center of mass. These displacements can be characterized as motion signals (C) with a reference or desired trajectory that is based on typical gait patterns. As a trajectory, the high quality execution does not exactly track the reference trajectory in time, but when we look at the Fourier reconstruction of the trajectory statistics (D), we can see that the high quality execution is very similar to the reference distribution. In contrast, the low quality execution has spatial statistics that are very different from the reference distribution.

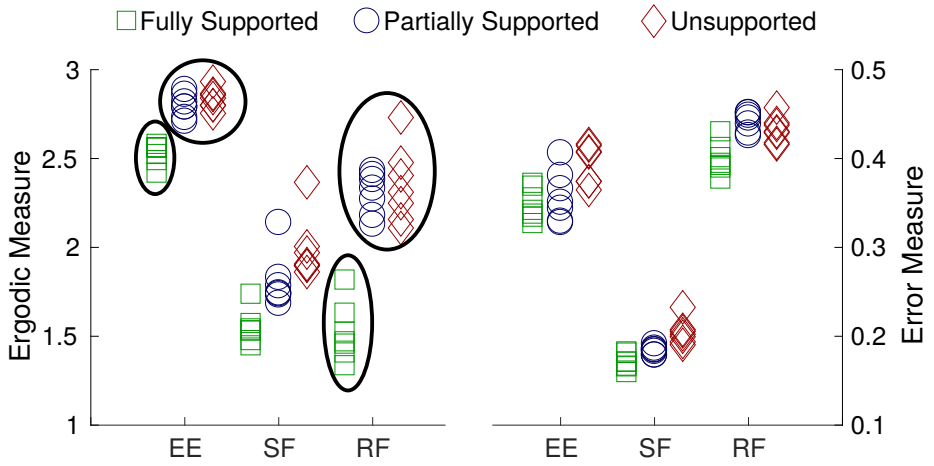

Figure 4.2. **Target-Reaching Trials of a Stroke Participant.** A stroke patient was asked to reach to one of three targets (EE=Elbow Extension, SF = Shoulder Flexion, RF=Reach Forward) in different areas of their workspace. The ergodic Measure (left) provides clear distinction between the level of full arm support and partial or no arm support in the case of both EE and RF (as indicated by the circled data). The error measure (right) provides little distinction between the fully supported case and the partially supported. Each marker represents a trial from the same individual.

the ergodic metric distinguishes between different levels of arm weight support—although error does not—even for a single participant.

Motivated by this individual result, we collected data from (healthy) human participants robotically assisted with a dynamic inversion and balance task, to see whether the presence of assistance and learning on the part of the participants can be detected. In both cases, the outcome is affirmative. Based on our task-specific measures, there is a clear distinction between the assisted and unassisted conditions. Analysis of the error measure does not show such a distinction based on assistance, but does detect a training effect based on the significantly lower RMS error of the training group in their second session compared to the control. The ergodic measure detects both the presence

of assistance and the training effect. These results suggest that measures of the information encoded by a movement can be used to predict the presence of assistance and that such measures capture outcomes that would not otherwise be captured in task-specific performance measures.

4.2. Materials and Methods

4.2.1. Quantifying Ergodicity.

One metric for quantifying ergodicity of the trajectory is the sum ε of the weighted square distance between Fourier coefficients of the distribution ϕ_k and the coefficients of the spatial Fourier transform of the trajectory c_k,

$$\varepsilon = \sum_{k_1=0}^{K} ... \sum_{k_n=0}^{K} \Lambda_k |c_k - \phi_k|^2, \tag{4.1}$$

where at each time the state is n-dimensional and there are $K+1$ coefficients along each dimension [79]. The subscript k in eq. 4.1 is a multi-index over the coefficients of the multi-dimensional Fourier transform. The coefficient $\Lambda_k = (1 + ||k||^2)^{-s}$ where $s = \frac{n+1}{2}$ places larger weights on lower frequency information, creating a Sobolev norm [79]. Using Fourier basis functions of the form,

$$F_k(x) = \frac{1}{h_k} \prod_{i=1}^{n} cos\left(\frac{k_i \pi}{L_i} x_i\right), \tag{4.2}$$

where h_k is a normalizing factor [79] and L_i is a measure of the length of the dimension, we can compute the coefficients of a spatial distribution or time-averaged trajectory using

Eq. 4.3 and Eq. 4.4 respectively[1].

$$\phi_k = \int_X \phi(x) F_k(x) dx \tag{4.3}$$

$$c_k = \frac{1}{T} \int_0^T F_k(x(t)) dt \tag{4.4}$$

The $x(t)$ defined here represents the set of ergodic states $x_\varepsilon(t)$, or states which are relevant to the task, which may be a subset of the full set of dynamic states $x(t) \in \mathbb{R}^n$ (i.e., $n_\varepsilon \leq n$). For example, for the cart-pendulum inversion task, the full dynamic state vector is $x(t) = [\theta(t), \dot{\theta}(t), x_{cart}(t), \dot{x}_{cart}(t)]$, but the relevant ergodic states for the inversion task are $x_\varepsilon(t) = [\theta(t), \dot{\theta}(t)]$. For the comparisons made in this chapter, RMS error was also calculated based on the ergodic states only.

4.2.2. Experimental Protocol and Analysis.

Fifty-three participants (17 males, 36 females) consented to participate in this study[2]. At the beginning of each session, the system and task was demonstrated to the participant using a video of a sample task completion. Participants were instructed to attempt to swing up the pendulum to the upward unstable equilibrium and balance there for as long as possible. Participants were instructed to continue to try to do this until the 30 seconds were over even if they succeeded at balancing near the equilibrium more than once. The full 30 seconds was used in calculation of all metrics, and the trajectory statistics were averaged over this time horizon. In each session, 30 trials were completed with short breaks upon request of the participant. In order to assess the effect of assistance, 40

[1]The choice of T

[2]This study protocol was approved by the Northwestern University Institutional Review Board and all the participants signed an informed consent form.

participants were tested in two sessions (one week apart), one session with assistance and one session without assistance. The order of the sessions was randomized to account for any learning effects. Paired two-sample t-test were performed to evaluate differences between the session with assistance and sessions without assistance. Another group of 13 participants performed two unassisted sessions one week apart to establish a baseline for learning from unassisted practice. A two-sample t-test was evaluated to test the difference in means between the second session of the control group and the second session of the group of 20 participants who used assistance in their first session.

4.3. Statistical Results.

The results of the comparison of the unassisted and assisted trials using two sample t-tests are summarized in table 3.1. These tests pair samples from each participant according to the order in which they were performed in each session—accounting for the variance between participants. The hypothesis testing was performed in R [99] by subtracting the unassisted condition from the assisted condition, showing improvement due to assistance in all five measures. Both the task-specific measures and ergodicity—measured as the distance from ergodicity—capture the effect of the assistance.

Additional analysis of the effect of assistance as participants progress through trials was performed by grouping individual trials into blocks of five trials, such that the effect of both the assistance and the trials could be assessed from independent and interaction effects.

The time spent at the balance position during each trial was analyzed with a 2 (assistance/no assistance) x 12 (blocks of 5 trials) mixed design ANOVA, which showed a

significant main effect of the assistance mode $F(1, 418) = 388.87,\ MSE = 1296.5,\ p < 2\times10^{-16},\ Cohen's\ f = 0.76$. The main effect of the block was also significant $F(11, 418) = 2.196,\ MSE = 7.3,\ p = 0.0139,\ Cohen's\ f = 0.19$. The assistance and block interaction effect was not significant $F(10, 418) = 1.266,\ MSE = 4.2,\ p = 0.25,\ Cohen's\ f = 0.14$.

The time to success was also computed for each trial. This measure was analyzed with the same 2 x 12 mixed design ANOVA, and showed only a significant main effect of assistance mode $F(1, 418) = 224.922,\ MSE = 16629,\ p < 2 \times 10^{-16},\ Cohen's\ f = 0.44$. The main effect of block was not significant $F(11, 418) = 0.809,\ MSE = 60,\ p = 0.63,\ Cohen's\ f = 0.09$. The assistance and block interaction effect also was not significant $F(10, 418) = 0.709,\ MSE = 52,\ p = 0.72,\ Cohen's\ f = 0.08$.

The mixed design ANOVA was also used to analyze the RMS error of the relevant states $(\theta, \dot{\theta})$ over each 30 second trial, and found no significant effects from any factor. The main effect of assistance was not significant $F(1, 418) = 1.367,\ MSE = 0.018,\ p = 0.24,\ Cohen's\ f = 0.05$. The main effect of block was also not significant $F(11, 418) = 1.399,\ MSE = 0.019,\ p = 0.17,\ Cohen's\ f = 0.18$. The interaction of block and assistance was not significant either $F(10, 418) = 0.609,\ MSE = 0.008,\ p = 0.806,\ Cohen's\ f = 0.11$.

The ergodic metric was computed over each 30 second trials using the relevant states $(\theta, \dot{\theta})$ and was analyzed using the 2 x 12 mixed design ANOVA. The only significant main effect was assistance mode $F(1, 418) = 62.51,\ MSE = 6.90,\ p = 2.38\times10^{-14},\ Cohen'sf = 0.35$. Block was not a significant main effect $F(11, 418) = 1.31,\ MSE = 0.144,\ p = 0.218,\ Cohen'sf = 0.17$, and the interaction of assistance and block was not significant $F(10, 418) = 0.691,\ MSE = 0.076,\ p = 0.73,\ Cohen's\ f = 0.12$.

These ANOVAs show that the task-specific measures and the ergodic metric detected the effect of assistance with a moderate effect size, while error did not distinguish between the assisted and unassisted conditions over the course of each session.

The comparison of the trained and control group using two sample t-tests is summarized in table 4.1. Unlike the assisted and unassisted trials, the groups in these tests

	Control $n = 13$		Post-Training $n = 20$				
Measure	μ	SD	μ	SD	t	df	p
Success Rate	0.438	0.211	0.380	0.200	-0.8032	31	0.4280
Balance Time	0.295	0.570	0.243	0.502	1.378	988	0.1687
Time to Success	24.319	7.844	24.981	7.794	1.301	988	0.1935
Error	0.629	0.061	0.621	0.058	-1.963	988	0.0499
Ergodicity	0.751	0.207	0.705	0.177	-3.701	988	2.266×10^{-4}

Table 4.1. **Two-sample t-tests of week 2 control trials and week 2 trained trials.** Hypothesis testing was performed in *R* [99] by comparing the means of the control group to the means of the trained group. Error and ergodicity—measured as the distance from ergodicity—were the only measures that revealed a significant improvement in the mean between the trained group and the control group. Note that the degree of freedom (df) for success rate is 31 since there is only one rate per participant.

are independent and therefore, the samples are not paired. The progress of the two groups over the second session (Fig. 6) was analyzed further by performing mixed design ANOVAs on the training group (between participants) and block (within participants).

The balance time of the control group and the trained group in the second session was analyzed with a 2 (training groups) x 6 (blocks) mixed design ANOVA, which showed no significant main effects or interactions effects. The main effect of training group was not significant $F(1, 31) = 1.202,\ MSE = 1.25,\ p = 0.28,\ Cohen's\ f = 0.08$. The main effect of block also was not significant $F(5, 155) = 2.018,\ MSE = 0.44,\ p = 0.079,\ Cohen's\ f =$

0.11, nor was the interaction of training and block significant $F(5, 155) = 1.05,\ MSE = 0.23,\ p = 0.39,\ Cohen's\ f = 0.08$.

The mixed design 2 x 6 ANOVA design was also applied to the time to success, and the main effect of training group was not significant $F(1, 31) = 0.334,\ MSE = 103.4,\ p = 0.567,\ Cohen's\ f = 0.05$. The main effect of block was not significant either $F(5, 155) = 1.34,\ MSE = 66.32,\ p = 0.25,\ Cohen's\ f = 0.09$. The interaction effect of block and training group also was not significant $F(5, 155) = 1.34,\ MSE = 66.50,\ p = 0.25,\ Cohen's\ f = 0.09$.

The same mixed design ANOVA was used to analyze the RMS error in each trial. The main effect of block was significant $F(5, 155) = 4.336,\ MSE = 0.011,\ p = 0.001,\ Cohen's\ f = 0.19$, but the main effect of training was not significant $F(1, 31) = 0.76,\ MSE = 0.035,\ p = 0.39,\ Cohen's\ f = 0.15$. The interaction effect of training group and block also was not significant $F(5, 155) = 1.61,\ MSE = 0.004,\ p = 0.16,\ Cohen's\ f = 0.12$.

The analysis of the ergodic metric using the mixed design ANOVA revealed a significant main effect of block $F(5, 155) = 2.88,\ MSE = 0.08,\ p = 0.0163,\ Cohen's\ f = 0.15$, and a significant interaction effect of block and training group $F(5, 155) = 2.33,\ MSE = 0.06,\ p = 0.045,\ Cohen's\ f = 0.14$. The main effect of training was not significant $F(1, 31) = 1.056,\ MSE = 0.49,\ p = 0.312,\ Cohen's\ f = 0.17$.

These ANOVAs demonstrate that the task-specific measures were not sensitive to either the improvement made by the participants throughout the second session or the benefit of the feedback provided to the trained group in the first session. The error measure indicates that users performed better over the course of the second session. In Fig. 6, the control group performed worse at the beginning of the second session that

it did at the end of the first session, and their performance increased in terms of error over the course of the session. The trained group also improves moderately during the second session. The ANOVA of the ergodic metric is also able to detect the significant improvement during the second session by the control group as well as the interaction effect of group and training. This interaction is a result of the trained group performing better at the beginning of the second session and maintaining that performance, while the control group eventually reached the same level of performance.

4.4. Results

In order to evaluate training effect and presence of assistance, participants were tested in both assistance and no assistance modes. Each participant completed 2 sessions, approximately one week apart. Upon enrollment in the study, each participant was placed into 1 of 3 groups. If placed in the training group (n=20), the participant completed the first session with assistance and received no assistance in the second session. If a participant was placed in the non-training group (n=20), they performed the task without assistance in the first session and used the assistive interface in the second session. Finally, a control group (n=13) performed the task without assistance in both the first and second session. Participants were tasked with inverting and balancing a virtual cart-pendulum system as shown in Fig. 4.3 (often studied in nonlinear control [45, 67]). For the purposes of calculating ergodicity, a delta function $\delta(x - s)$ at the goal state s was used.

In this experiment, we implemented a form of assistance that can convert pure noise input into a successful task execution by comparing the noise input to that of an optimal controller [129]. The assistance acts as a filter similar to that described in [129] and [32],

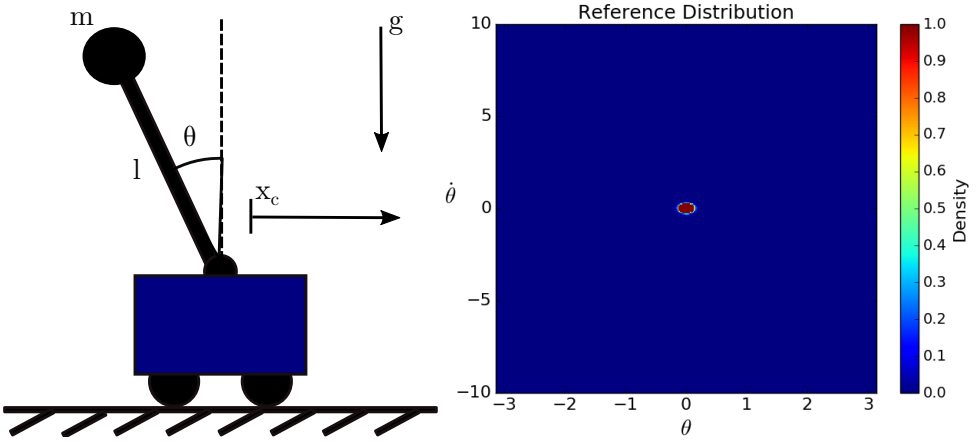

Figure 4.3. **Experimental System.** Participants controlled the cart position x_c directly and indirectly controlled the angle θ and angular velocity $\dot{\theta}$ of the cart-pendulum system (left). The goal state used to calculate the RMS error was $(\theta, \dot{\theta}) = (0, 0)$ and the distribution used as the task definition for the information measure was a Dirac delta function at $(\theta, \dot{\theta}) = (0, 0)$ (right).

such that if user inputs agree with the optimal controller, user input is not modified by the interface. When user inputs do not agree, the robot physically rejects the input, providing feedback but not guidance. The user input—acceleration of the cart as measured at the robot end-effector—is either accepted or rejected at each instant based on whether or not the input vector is in the same direction as the input prescribed by an optimal controller. Note that the objective of the optimal controller is to minimize the error between the system trajectory and the goal state at the unstable equilibrium, $s = (\theta, \dot{\theta}) = (0, 0)$. The input vector calculated by the optimal controller is never implemented. It is only used as a filtering criterion. The input is completely rejected—replacing the user input with a zero-vector—when user inputs do not agree with the optimal controller. If the participant is a perfect actor, the assistance is completely transparent. Details of the assistance algorithm can be found in chapter 3.

4.4.1. Assistance adds task information

Several task-specific performance measures were recorded, including success rate, balance time, and time to success. The training and non-training group data was aggregated to evaluate the effect of assistance on the 40 participants in a counterbalanced fashion. Paired two-sample t-tests on these task-specific measures of the participants with and without assistance showed that the participants improved with the addition of assistance. The assisted trials had a higher success rate ($p < 0.001, t(39) = 12.314$), more cumulative time spent at the goal state ($p < 0.001, t(1199) = 26.519$), and reached the goal state more quickly ($p < 0.001, t(1199) = 17.202$). The trajectories generated in the assisted condition were also more ergodic with respect to the task distribution than those without assistance when the paired two sample t-test was performed ($p < 0.001, t(1199) = 11.261$). However, the root mean square (RMS) error of the trajectory from the goal state did not show a significant difference ($p = 0.094, t(1199) = 1.674$) between the assisted and unassisted conditions. This suggests that the assistance improved task-specific performance metrics and increased the task information encoded by the movement. A two-factor (assistance and block) analysis of variance also showed that the assistance had a significant effect in terms of the task-specific measures and the ergodic metric. However, the analysis of the RMS error revealed no significant effects. When we look at the spatial statistics (see Fig. 3.6) of the assisted trials versus the unassisted trials, we see that the assisted trials spend a larger proportion of time near the origin where the target distribution is centered. The histogram of unassisted trajectories (left in Fig. 3.6) has its highest density at $\theta = \pm\pi$ which is the farthest point from the goal state. The rest of the distribution is diffuse over the state space. Although the histogram of the assisted trajectories (right

in Fig. 3.6) also has a high density at $\theta = \pm\pi$, the distribution is not as diffuse as that of the unassisted trajectories. There are bands of high density spreading outward form the goal state $(\theta, \dot{\theta}) = (0, 0)$. The spatial statistics of the assisted trajectories are more similar to the reference distribution in Fig. 4.3, because there is a high density at and around the goal state. This suggests that assistance increased the task information encoded in the movement. This outcome is captured by measuring the ergodicity of the trajectories in each group with respect to the reference distribution. The mean ergodicity of the unassisted trajectories is 0.739, and the mean ergodicity with assistance is 0.631. This lower number indicates that less information is lost in the assisted motion than the unassisted motion.

4.4.2. Learning involves increasing task information

The effect of training was assessed by comparing the week 2 session of the trained group to the week 2 session of the control group. Although performance in the task-specific measures did not improve with training, the ergodicity and error of the trained group was significantly better than the control group. A two-sample t-test was performed on the task specific performance measures, finding no difference between trained group and untrained group in terms of their success rate ($p = 0.4280, t(31) = -0.8032$), time spent balanced ($p = 0.1687, t(988) = 1.378$), and time to success ($p = 0.1935, t(988) = 1.301$). The two-sample t-test of the RMS error showed a significant difference between the trained and control group ($p = 0.0499, t(988) = -1.963$), but the effect size was small ($d = 0.127$). The t-test of ergodicity ($p = 2.266 \times 10^{-4}, t(988) = -3.701$) also detected the difference, but with a larger effect size ($d = 0.237$). A two-factor (training group and block) analysis

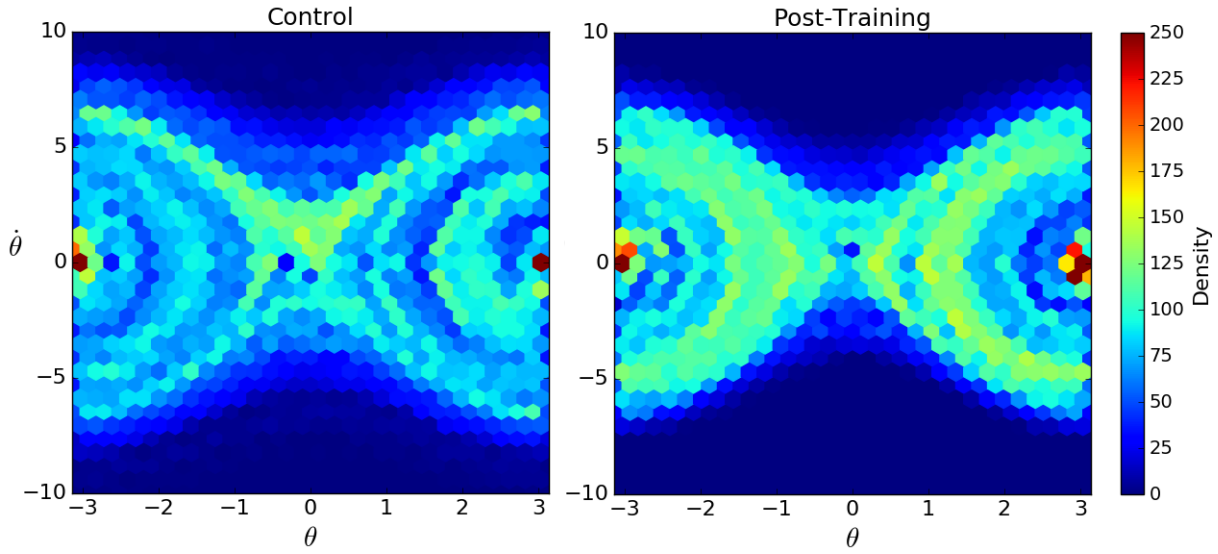

Figure 4.4. **Learning increases information.** The histogram of week 2 control trajectories (left) has its highest density at $\theta = \pm\pi$ which is the farthest angle from the goal state at $(\theta, \dot{\theta}) = (0, 0)$. The control trajectories also spend time near the goal state, but to a lesser extent. The histogram of trained trajectories (right) also has high density near $\theta = \pm\pi$, but there are large bands of high density in the region $-1.5 \leq \theta \leq 1.5$ and $-4 \leq \dot{\theta} \leq 4$. These bands make the statistics of the trained group closer to the spatial statistics of the reference distribution in Fig. 4.3. We quantify how well these statistics match that of the reference by calculating the ergodicity. The trained trajectories are on average more ergodic ($\mu = 0.705$) than the controls ($\mu = 0.751$). In other words, the trained motions communicate information about the task goal more effectively than the control motions.

of variance also showed that block and the interaction between training group and block had a significant effect in terms of the RMS error and the ergodic metric. However, the task-specific measures revealed no significant effects. This indicates that the training effect is not captured by task-specific measures but is captured by error and ergodicity. While the training effect can be detected with error, the information measure detects it with a larger effect size.

4.5. Discussion

Our results suggest that the information encoded in bodily motion provides a language for describing changes due to assistance and learning in physical human-robot interaction. Using this framework, we found that an information measure was a better predictor of changes in deficit compared to error, even when the robotic assistance was based on that error metric. Additionally we showed that when the error-based assistance was used to supplement training, task-specific measures failed to detect the effect of training, and the ergodic measure demonstrated a stronger statistical power to detect the performance changes due to training. When we consider the body as an information channel and bodily motion as an information-carrying signal, we can interpret these results in a way that captures phenomena that would not otherwise be captured by task-specific measures or standard measures such as error. This analysis could provide valuable insight into changes in performance over the course of training or therapy.

When we examined the effect of robotic assistance in reducing the deficit of a stroke participant, we found that simply reducing the deficit by arm weight support increased the information encoded in their reaching motions to multiple targets. Specifically, their motions were more ergodic with respect to a Dirac delta distribution at the target position. We found that task-oriented assistance using an error-based optimal control also reduced the information lost in task executions. These decreases in information loss indicate that the information channel (the human-robot pair) itself was improved by the addition of both forms of assistance.

Task-specific measures also reflected differences in task executions due to deficit. Clinically, task-specific measures are frequently used to assess deficit [43, 44, 53, 68], often

motivated by the fact that more standard and generalizable measures, like error and energy, fail to predict the presence of deficit. As a consequence, deficit often must be assessed in very narrow experimental conditions where the measures are applicable. These task-specific measures have several negative consequences. First, they do not translate to other motions (e.g., an assessment strategy for reaching cannot be applied to interpretation of walking or even self-feeding). Moreover, and more importantly, task-specific measures do not admit the same level of principled interpretation as measures such as state error (which captures motor control accuracy) and energy (which captures metabolic efficiency). Measuring the task information encoded by a movement provides an information-theoretic framework for interpreting motion—in a principled manner like error or energy—by capturing the qualitative description of task while not implicitly prescribing a specific strategy for task completion.

Task specific measures are able to capture qualitative task success because they measure events at the goal state or task outcomes—making them independent of the strategy. In contrast, measurement of error is typically explicitly dependent on a reference trajectory, so error measures and error-based assistance prescribe a specific strategy for task completion. Measures on information are still independent of the strategy chosen, but can be expected to detect that a strategy is encoded in the movement even if the movement does not result in task success. A movement with a focused strategy should have higher task information than a movement without a focused strategy. In the case of the swing-up and inversion problem, the assistive algorithm forces participants to employ an error-reducing strategy, increasing the task information in the assisted movement as seen in Fig. 3.6.

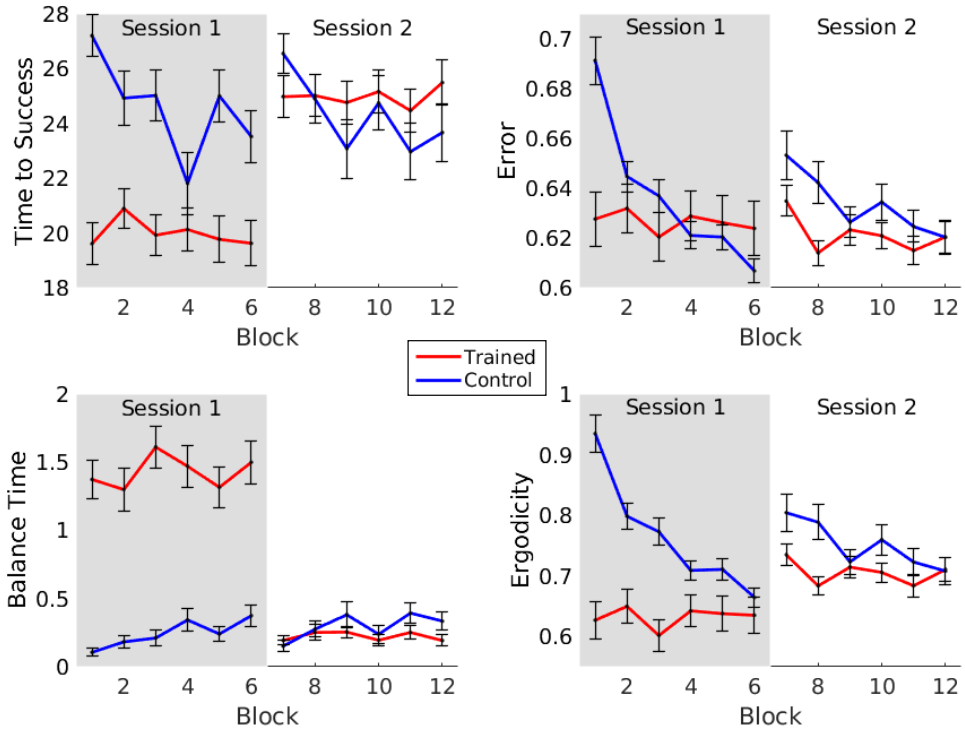

Figure 4.5. **Comparison of the control and trained group performance progress over training.** The statistical comparisons of the trained and control groups excluded the data from the first session (gray) to avoid including the effects of the assistance algorithm itself. For the task specific-measures (top row), there was no difference between the two groups, and block had not significant effect on performance. For the error and ergodic metric, block has a significant effect, especially in the control group. Under both measures, the control group performance was worse at the beginning of the second session (the first two blocks in white), but by the end of the session performed as well as the trained group. Ergodicity enables one to see the difference between the treated and untreated group, and both error and ergodicity allow one to see learning as a function of block.

Participants attempting to maintain the error reducing strategy after training with the assistance explains the training effect that we see in the error and ergodic metrics (see Fig. 4.5). However, the error-based assistance used in this experiment may not be the most effective strategy to improve task outcomes like balance time and success rate. Assistive

algorithms that provide too much guidance—thereby reducing errors during training fail to result in improved training because both error [125] and kinematic variability[55] are critical to learning. Additionally, the feedback may lead users to learn the wrong task[66]. In this case, participants may have learned to reduce the cumulative error without learning the true task goal which is to reach the unstable equilibrium regardless of the RMS error incurred before or after reaching that state.

While this work demonstrates that the ergodic measure is useful for assessing human motion post-hoc, it could also be used as a powerful tool for modulating haptic and kinesthetic feedback during training. There is already some evidence that providing feedback based on relevant measures of movement quality is more effective that trajectory error-based assistance [47]. Using the ergodic metric, one can generate distribution-based representations of tasks from sets of imperfect demonstrations—providing an alternative means for robots to 'learn' from humans by applying information maximizing techniques [81, 133]. Rather than learning a policy or task objective, one could generate a reference distribution from recordings of human task executions and use the distribution as the task objective for a deterministic information-based model-predictive controller, allowing the human-robot pair to accomplish the task from different initial conditions and under various system constraints. Furthermore, providing assistance based on the ergodic measure allows us to build joint-human robot control policies that directly encode the natural variability of human motion, such that we do not need to restrict assistance to enforce a particular goal trajectory or make inferences about which movement parameters are relevant to the task.

These results suggest that the ergodic measure can augment error or energy measures in the study of biomechanical motion—providing fine-grained insight on the progression of robot-aided training and therapy. Specifically this study supports the idea that ergodicity is a principle of motion for interpreting and predicting animal motion with potential implications for the design of effective feedback and training strategies.

CHAPTER 5

Ergodic Shared Control: Closing the Loop on pHRI Based on Information Encoded in Motion

Advances in exoskeletons and robot arms designed for physical human-robot interaction have given us increasing opportunities for providing physical support and meaningful feedback in training and rehabilitation settings. However, roboticists and researchers must choose control strategies that support motor learning and provide mathematical task definitions that are actionable for the actuation. Typical robot control architectures rely on measuring error from a reference trajectory. In physical human-robot interaction (pHRI), this leads to low user engagement, invariant practice, and few errors, which are all negative features in the context of motor learning. Furthermore, a reliance on reference trajectories means that the task definition is both over-specified—requiring specific timings not critical to task success—and lacking information about normal variability in task execution. In this chapter, we examine a new way to define tasks and close the loop on physical human robot interaction using an ergodic measure that quantifies how much information about a task is encoded in the human-robot motion. This measure can capture the natural variability that exists in typical human motion—enabling therapy based on scientific principles of motor learning. We implement an ergodic hybrid shared controller (HSC) on a robotic arm as well as an error-based controller—virtual fixtures—in a timed

drawing task. In a study of 24 participants, we compare ergodic HSC with virtual fixtures and find that ergodic HSC leads to improved training outcomes.

5.1. Introduction

Designing robotic assistance frameworks presents a set of challenges that are unique to HRI and highly dependent on the task, environment, capabilities of the user, and the ultimate goal of the interaction. Robotic interfaces can act as guides in remote teleoperation or robotic surgery—providing haptic feedback and constraints to users—or shared control may enable robots to act as partners providing assistance that compensates for a user's deficit due to disability or high cognitive load[46, 90]. While the goal of teleoperation and assistive robotics is successful task completion, rehabilitation robots must target learning—sometimes at the expense of achieving task success. Robotics have great potential for enhancing training and rehabilitation because of their ability to support many task repetitions and to quantify the performance of the trainee. However, the automation cannot simply track the same series of joint states over and over—measuring the user's error. Instead, the automation must take actions based on a quantitative task definition and the specific needs of the user.

Motor learning literature and past studies of control strategies for training provides us with several requirements that robots must meet in order to promote learning. Most important in these is a need for active involvement by the user. When a patient is suffering from severe deficit, such as immediately after a stroke, there is some benefit to passive movement. However, in the long term, simply moving a person through the motions of a task leads to slacking [42]. Training and therapy is most effective when training is

intense and patients are actively involved in the exercises [73]. User engagement can be promoted using game-like graphic interfaces for task-oriented training [113, 137] and by matching the control strategy to the relative difficulty of the task [71, 82]. This suggests that typical robot control frameworks tracking a particular joint trajectory are ill-suited for robot-mediated training.

Instead, impedance/admittance control, virtual fixtures, and potential fields are used to 'assist-as-needed' or to challenge high skill users through error augmentation, kinematic variability, and random perturbations [10, 77, 97]. Although hardware advances have made robots more versatile and made it possible for us to autonomously support an almost unlimited number of meaningful tasks, most prior work in the upper extremity has focused on path-following tasks, where the goal is to minimize tracking error or follow a normative velocity profile. This is because a fundamental problem in physical human-robot interaction is how we should define the desired behavior [77, 97]. There are two main aspects to consider. One must specify the low level task that the actuators and encoders manage—e.g., a desired trajectory, velocity profile, or torque specification—and at a higher level one must select from one of many possible strategies to achieve the desired behavior. While the simplest solution may be to minimize the error between the robot and a recorded trajectory based on expert input or average behaviors, this brings up issues of time-dependence and enforcing a particular solution to a task that may have infinitely many 'good enough' solutions. This approach also neglects a fundamental principle of motor learning—that errors and variability actually enhance learning [125]

We assert that interfaces must enable flexibility in task solutions and provide task-based feedback rather than error-based. Stereotypical movements—such as reaching,

cleaning, self-feeding, and walking—have substantial variation between equally qualitatively successful trials, both within and between individuals. This variability is necessary for motor learning, so our control objectives need to maintain a task-level view of performance. This perspective is achieved by monitoring the spatial statistics of movement rather than relying on error or task-specific performance heuristics—providing assistance based on a global measure as opposed to local interactions.

In this chapter, we present a novel method for providing corrective feedback based on a statistical task definition that can be generalized to a broad set of tasks. Rather than asking the robot to help the user minimize their tracking error, the robot intervenes to *increase the information about the task that is encoded in the motion.* This is implemented by using an ergodic measure to close the loop on our hybrid approach that switches between full user autonomy and full rejection of user inputs. In a user study, we implement ergodic hybrid shared control (HSC) on an impedance control robot (Figure 5.1) and empirically compare training with the ergodic HSC to an assist-as-needed path controller.

The chapter is organized as follows: First, in Section 5.2, we provide background information on control strategies used in robotic training applications and highlight two control design features that greatly impact the performance of these training strategies. This is followed by a discussion of our prior work and the ergodic HSC algorithm in Section 5.3. Results of applying the algorithm to simulated noise input are given in Section 5.4. The experimental protocol and design can be found in Section 5.5 followed by the results in Section 5.6. Finally, we provide a discussion of the results in Section 5.7.

Figure 5.1. Participants were seated in front of a display showing the reference image and corresponding square in which they were to copy the drawing. They held onto a handle rigidly fixed to the end-effector of the Sawyer robot. The display provided visual feedback on what they had drawn so far and the amount of time remaining. A physical card with the given reference image was displayed throughout each trial. Study participants were asked to reproduce the images on the left as quickly as possible. They were given a maximum of 10 seconds per drawing.

5.2. Related Work

5.2.1. Control Strategies for Training

Control strategies that are designed to support learning can be classified as either passive, corrective, resistive, or some combination thereof. Many early rehabilitation robots provided passive assistance to users by replaying trajectories from healthy humans or experts [11, 15, 84, 120]. This type of guidance is beneficial for patients in an acute stage of recovery or when the relative difficulty of the task is exceedingly high. As patients recover and relearn motor skills, the task should be performed under shared control [42, 88]. In an effort to avoid user passivity, Assist-As-Needed (AAN) control was introduced and is now widely used. This can be done by adjusting the relative contributions of the robot and the human based on user performance, or it can be corrective—using virtual fixtures to reject

user commands when they deviated from the planned trajectory or enter restricted regions [10]. AAN is usually achieved through impedance/admittance control [97]. Resistive strategies have also been implemented by adjusting the impedance gains on the control or specifying a novel potential field where the forces push users away from the desired trajectory or path, but results are mixed. Users benefit from resistive strategies when the relative difficulty is low, whereas users with significant deficits need assistance [71, 82]. Our approach is not to provide assistance or resistance based on some desired velocity profile, but rather to accept or reject user actions based on global knowledge of the task goal. This results in an interface that inherently adapts to user performance and skill [31], while other strategies require two separate strategies to select the appropriate level of controller intervention and the lower level modulation of controller responses to individual user actions.

5.2.2. Adaptation

To promote user engagement and participation during training, controllers must modulate assistance based on user performance or the relative difficulty of the task. The most common solution is trial-by-trial adaptation based on tracking error [8, 140]. In some cases, mean velocity [37, 54], anticipated deviation based on model of participant impairment [93], or other performance heuristics are used. This adaptation has led to moderately better therapuetic outcomes compared to conventional therapy [106]. Yet, the rate at which control parameters are adapted may affect the skill retention [71]. Regardless of the adaptive scheme or the assistive/resistive strategy, some definition of this desired behavior is required.

5.2.3. Reference Definiton

Defining the reference for pHRI involves choosing between many possible solutions to achieve the task goal and translating that solution into something that the system can feasibly execute and monitor throughout that execution. When autonomous systems are acting alone, one can simply define a trajectory over the robot joint states and use state error feedback to execute the trajectory. However, in activities of daily living, there are many possible redundant solutions to seemingly simple tasks like reaching and stepping. Often trajectories are recorded from a human expert [11, 56, 120], a healthy person [15, 94], or the unaffected limb in hemiparetic stroke [11, 74], but the time-dependence in these definitions is problematic. A task execution is not less successful because it is slower or faster than the reference. This can be avoided by defining a reference path rather than a reference trajectory. Potential fields can limit the distance from the path [21, 75] or, in the case of error augmentation, push users away from the desired path.

Still, this does not answer the higher level question of which trajectory, path, or task strategy the robot assistance should be supporting. In cases where the task goal is ambiguous, we could plan reference definitions for a short time into the future using intent detection based on EMG [70], EEG [16, 34, 111], force torque sensors [65, 107], or kinematic data [144]. Proietti et. al suggest that another potential solution is to generate trajectories from statistically consistent patterns from a sufficient number of healthy subjects [97]. There are almost no instances of this in the literature, possibly, because it is not clear how one would close the loop on statistical patterns—something that we specifically achieve in the present work. Even if the average of the two paths were used as a reference, the resulting desired trajectory may not convey the task goal, and

the variance and other statistics of the set of motion strategies are not part of a typical control architecture. In this chapter, we present an approach to reference definition based on statistical task definitions and modulate robotic assistance based on concepts from information theory.

5.2.4. Current Study

Our approach is to treat the body as a communication channel and motion as an information carrying signal by relating the time-averaged behavior of trajectories to a spatial distribution that describes the task. To quantify information content in a motion in a domain $\mathcal{X}$, one needs to measure a trajectory $x(t)$(e.g., position of an arm over time) that describes movement of the body and describe the task as a distribution $p(x)$ (e.g., the volume of positions that the arm has been in) over states in the domain. Describing a task by a distribution requires that there is a natural way to describe a task statistically rather than by specifying a goal state or a goal trajectory. For instance, if there is a particular goal state s, one can represent this as a singular function with infinite value at the goal state. Or, if the task definition is a consequence of measuring many instances of task execution as in Learning from Demonstration (LfD) [6], the collection of observations will form a distribution. As more demonstrations of the target reaching task are added to the set of observations, the collective time spent at the goal state would generate a higher peak at the state s, asymptotically approaching a singular function at the goal state. In either case, we assess a motion by asking how much information about the task is encoded in the movement. To compare a trajectory to a distribution, we use ergodicity—which relates the temporal behavior of a motion signal to a distribution. We implement a hybrid

shared controller using an ergodic measure to update the control of a robot arm online. An error-based AAN path controller is also implemented on the robot arm—a form of virtual fixtures [10]. In a study of 24 participants performing a timed drawing task, we compare our approach to virtual fixtures that provide resistance to users when they are too far from the desired path.

5.3. Ergodic Hybrid Shared Control

Our previous work has compared error-based hybrid shared control (HSC) to unassisted practice in cart-pendulum inversion task. We found that it adapted to user skill, improved performance while in use, improved training within a single session, and led to greater skill retention over a one week period [31]. An analysis of the data from those experiments led us to take a closer look at the measures that we were using to evaluate performance. Task-specific measures like time to success and balance time captured the effect of HSC as form of assistance, but did not capture any training effects over time in either the control or experimental group. The RMS error of the system states did not capture the clear improvement in performance when assistance was provided by HSC, but showed that both time and training conditions were statistically significant factors. An ergodic measure was able to capture both of these effects [30]. This unintuitive outcome indicates that error-based assistance, by far the most common way of creating an assistive device and measuring its efficacy, may not be an effective measure for many dynamic tasks.

5.3.1. Ergodic Metric

Ergodicity can be measured by several metrics [114, 115]. When we analyzed ergodic motion in our prior work, we used the spectral approach [79], which characterizes ergodicity by comparing spatial Fourier coefficients of a trajectory $x(t)$ to coefficients of a reference distribution $p(x)$—giving us the distance from ergodicity. While one can compute controls based on this formulation as in [81], there are two major drawbacks. First, this approach scales as $\mathcal{O}(|k|^n)$, where k is the maximum integer-valued Fourier term and n is the number of relevant states. Second, the use of periodic basis functions leads to artifacts in the reconstructed distribution. Both of these factors contribute to limitations in computing the measure online for complex, high-dimensional tasks. In [1], Abraham et al proposed an alternative measure of ergodicity using the Kullback-Leibler Divergence [61], where the time-average statistics of the trajectory are defined as a mixture distribution:

$$q(s|x(t)) = \frac{\eta}{t_f - t_0} \int_{t_0}^{t_f} exp\left[-\frac{1}{2}(s - x(t))^T \Sigma (s - x(t))\right] dt \tag{5.1}$$

where η is a normalization constant and $\Sigma \in \mathbb{R}_{n \times n}$ is a parameter representing the covariance of the Gaussian approximation. This is an approximation because the time-averaged statistics of the trajectory is actually a collection of delta functions parameterized by time.

Under this definition, we can define the ergodicity of the trajectory relative to the distribution $p(x)$ using the Kullback-Leibler divergence as:

$$\begin{aligned} D_{KL}(p(s)||q(s)) &= \int_{\mathcal{X}} p(s) ln \frac{p(s)}{q(s)} ds \\ &= \int_{\mathcal{X}} p(s) ln\ p(s) - p(s) ln\ q(s) ds \\ &= -\int_{\mathcal{X}} p(s) ln\ q(s) ds \end{aligned}$$

Note that we drop the first term because it does not depend on the trajectory $x(t)$. Rather than computing the integral over the entire domain $\mathcal{X}$, we approximate it by sampling such that given a set of N points $\mathcal{S} = s_1, ..., s_N$ randomly sampled over the domain $\mathcal{X}$, the KL-divergence ergodic measure [1] is computed by,

$$\varepsilon_{KL} = \sum_{i=1}^{N} p(s_i) ln \int_{t_0}^{t_f} exp\left[-\frac{1}{2}(s_i - x(t))^T \Sigma (s_i - x(t))\right] dt. \tag{5.2}$$

5.3.2. Mode Insertion Gradient

The mode insertion gradient gives an estimate of the sensitivity of the cost to switching from one mode to another at a particular time. Therefore it is used in mode scheduling problems to find the optimal time to insert a mode from predetermined set [4, 12, 22, 35, 136]. In hyrbid shared control [31], we instead use the mode insertion gradient to determine whether or not to accept a switch from the nominal control u_1 to the user action, u_{user}. We define the hybrid control, u_2 with the piece-wise function below,

$$u_2(t) = \begin{cases} u_{user} & t \leq t_0 + t_s \\ u_1 & t_0 + t_s < \ t \leq \ t_0 + T \end{cases} \tag{5.3}$$

We assume a system with dynamics

$$\dot{x}(t) = f(x(t), u(t), t) = g(x(t)) + B(x(t), t)u(t), \tag{5.4}$$

where $\dot{x}(t)$ is linearly dependent on the control u. The cost describing the task objectives is

$$J = \varepsilon_{KL} + \int_{t_0}^{t_f} l_2(x(t), u(t))dt, \tag{5.5}$$

where $l_2(x, u)$ represents the running cost associated with the control effort, safety parameters, or other secondary objectives. The mode insertion gradient is then,

$$\frac{dJ}{d\lambda} = \rho(\tau)^T \left[f(x(\tau), u_2(\tau)) - f(x(\tau), u_1(\tau)) \right]. \tag{5.6}$$

In (5.6), state x is calculated using nominal control, u_1, and ρ is the adjoint variable calculated according to Equation 5.7, below.

$$\dot{\rho} = -l(x, t) - D_x f(x, u_1)^T \rho, \tag{5.7}$$

$$subject\ to\ \rho(t_0 + T) = \mathbf{0}$$

where

$$l(x, t) = -\sum_{i=1}^{N} \frac{p_{s_i}}{q_{s_i}} exp\left[-\frac{1}{2}(s_i - x)^T \Sigma^{-1} (s_i - x) \right] (s_i - x)^T \Sigma^{-1} + \nabla l_2(x, u).$$

In the work presented here, we define the nominal control, u_1, to be equivalent to the free dynamics of the system. Nevertheless, a calculated controller action could be used to

define u_1 as in [50] and [31]. User inputs are then accepted when the following integral computed over a time window into the future is negative:

$$\int_{t_0}^{t_0+T} \frac{dJ}{d\lambda}(t)\delta t. \tag{5.8}$$

When the integral of the mode insertion gradient is negative, u_2 is a descent direction over the time horizon, T. Thus the mode insertion gradient as an acceptance criterion represents quantitatively the advantage or disadvantage of allowing the user to push the system in the way that they are currently trying to move it.

5.3.3. Algorithm Implementation

The ergodic hybrid shared control algorithm works as follows. Given a system and an operator, assume that a user input is measured every t_s seconds. The user input is used to define the control input u_2, and the system is simulated forward for T seconds into the future. The user input is assessed based on the integral of the mode insertion gradient—roughly asking whether the user understands the task goal. When the integral is negative, if the magnitude of the user command is within the allowed limits, the command is applied to the system. Otherwise, saturation may be applied. On the contrary, if the criterion is not met, the user input is rejected. When inputs were rejected in these experiments, the impedance at the end-effector was increased proportional to the difference between the current velocity and the velocity of the system at the time of the last accepted input. This results in the interface being transparent when user inputs are accepted or velocity being held constant when inputs are rejected. This process is illustrated in Algorithm 2.

Algorithm 2 Ergodic hybrid shared control

Initialize current time t_0, sampling time t_s, time horizon length T, final time t_f, input saturation u_{sat} and target distribution $p(x)$.

while $t_0 < t_f$ **do**
  Generate N samples of s_i within the domain $\mathcal{X}$
  Infer user input u_{user} from sensor data
  Simulate the system (5.4) for $t \in [t_0, t_s + T]$
  Simulate (5.7) for $t \in [t_0, t_s + T]$
  Compute $\int_{t_0}^{t_s+T} \frac{dJ}{d\lambda}(t)\delta t$ using Equation 5.6.
  if $\int_{t_0}^{t_s+T} \frac{dJ}{d\lambda}(t)\delta t < 0$ **then**
    if $|u_{user}| < u_{sat}$ **then**
      Use u_{user} as current input, $u_{curr} = u_{user}$
    else
      Apply saturated user input $u_{curr} = u_{sat}$
  else
    Completely "reject" u_{user} ($u_{curr} = 0$)
  Apply u_{curr} for $t \in [t_0, t_0 + t_s]$
  $t_0 = t_0 + t_s$
end while

5.4. Simulation Results

Ergodic HSC was first implemented on a simulated double integrator system with a sampling rate of 60 Hz, where the user input was randomly sampled from a uniform distribution $u_{user} \in [-10m/s^2, 10m/s^2]$. We can see that even using noise as input, the hybrid shared control paradigm can produce drawings that resemble the original images (Figure 5.2). Unlike the robotic platform used in the experiment, where shared control was implemented by updating impedance control parameter, the simulation is able to perfectly reject the noise input. In prior work, we have noted the differences in performance when the robotic interface has relatively low power compared to the user [32] with learning outcomes being more significant when the system has power sufficient to mechanically alter the movements of the user [31].

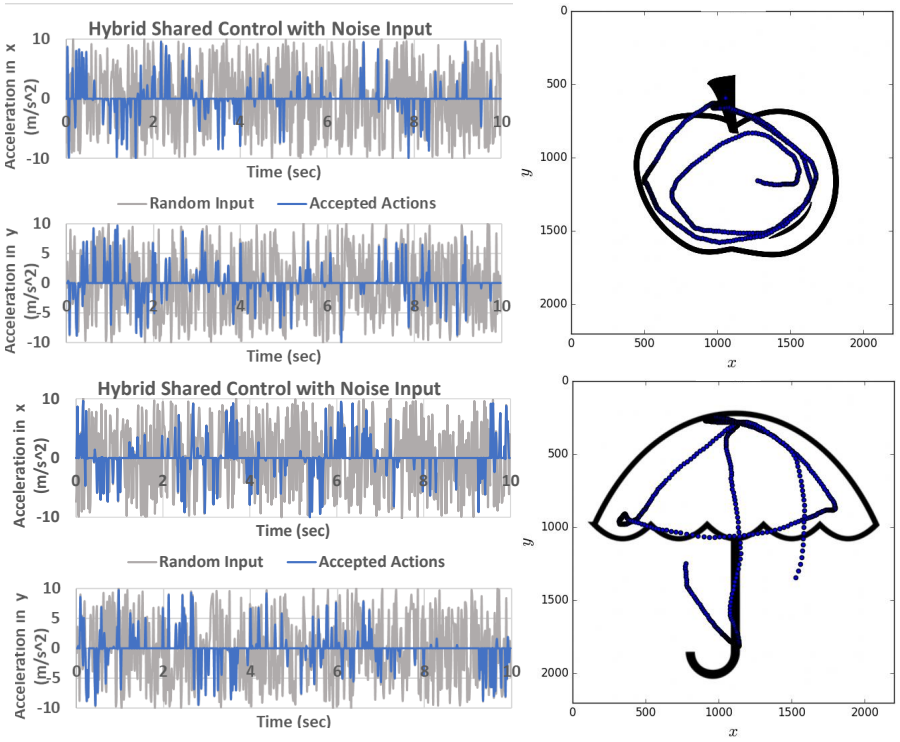

Figure 5.2. When noise is used as input, we can use hybrid shared control to reject any noise that does not represent a descent direction for the ergodic cost. This enables us to transform a noisy input (gray) that would generate a random walk to the filtered input (blue). This generates a trajectory that roughly follows the lines of the original image.

5.5. Methodology

In this study, we assess ergodic hybrid shared control in a set of timed drawing tasks, and compare it to virtual fixtures. There are many possible strategies one could use to complete a drawing, and the order of locations visited does not affect the end result. Yet, one can still make a comparison to an error-based controller by using the distance from the lines in the drawing to generate a set of virtual fixtures that limit the distance between the user's path and the desired path. The time limit was added to make the task more dynamic. Twenty-four participants were asked to perform a set of baseline tests without

assistance, followed by a training set and post-training set for each image tested. The type of assistance provided during training was randomly assigned. The experimental design allowed us to establish performance changes due to each of the control strategies and compare their effects on training.

5.5.1. Study Design/Apparatus

Participants were asked to draw the four grayscale line drawings shown in Figure 5.1. Each were sized to $2200px \times 2200px$ They were given a maximum of 10 seconds to copy each image into a box on the screen. The area on the screen corresponded to a $1m \times 1m$ horizontal plane in front of the user, much like the mapping of a mouse to a computer screen. For our experiment, we utilized a Sawyer Robot (Rethink Robotics) in the interaction control mode provided in the Intera SDK. Interaction control mode enabled us to set the type of control—force vs. impedance—of each dimension in Cartesian space. Sawyer's integrated sensors were used track the state of the end-effector as well as estimate the user inputs. The sensor information was sent to the host computer which simulated a double integrator system and executed algorithm 2—sending updates on the interaction control parameters according to Equation 5.9 and Equation 5.10, defined below. The host computer also updated the visualization provided to the users.

5.5.1.1. Virtual Fixture Implementation. Virtual fixtures were implemented by setting the planar components of the interaction control commands to force mode. The other dimensions were set to a high level of impedance, restricting the motion of the user to the horizontal plane. When users were within 100 pixels (approximately 4.5cm) of a dark pixel in the given drawing, impedance and force parameters were set to $\mathbf{0}$. When this

condition was violated, a force proportional to the distance to the nearest pixel and the current velocity was produced by the robot as in Equation 5.9,

$$F_{VF} = \frac{K_P(d_p - r)}{d_p} \begin{pmatrix} x_p - x_{user} \\ y_p - y_{user} \end{pmatrix} + K_D \mathbf{v_{user}}, \tag{5.9}$$

where d_p is the distance to the nearest dark pixel, r is the radius of the channel, and (x_p, y_p) are the coordinates of the nearest pixel. K_P and K_D are gains on the proportional and derivative terms of the feedback law, respectively.

5.5.1.2. HSC Rejection Equations. There are several ways to implement the rejection of user actions described in Algorithm 2. When using a low power haptic device, one can generate a transient virtual wall or ignore the user inputs when using an admittance controller with sufficient mechanical power to generate forces equal to that of the user. In this study, we modify the impedance parameters in the end-effector space of the robot. The task-irrelavant dimensions were set to a high impedance—again restricting the motion of the user to the horizontal plane. When user actions are accepted, impedance in the plane is set to zero. When user inputs are rejected, impedance parameters are set to track the velocity at the time of the last accepted action according to the following equation:

$$\begin{pmatrix} D_x \\ D_y \end{pmatrix} = \begin{pmatrix} sgn(v_x) & 0 \\ 0 & sgn(v_y) \end{pmatrix} \begin{pmatrix} \Delta v_x \\ \Delta v_y \end{pmatrix}. \tag{5.10}$$

5.5.2. Procedure

At the beginning of the session, participants were seated in a chair facing the robot and a display screen, and were asked to grasp a handle on the robot end-effector. Sawyer is

capable of exerting forces at this interaction point between the user and the robot in the x, y, and z directions, and can exert torques about all 3 axes. However, we maximized the impedance on the torques about these axes as well as the force in the z-direction, restricting the end effector to a horizontal plane. The position of the end-effector was measured from the joint angles using forward kinematics, and the acceleration of the end effector was measured using an inertial measurement unit installed in the end-effector. The acceleration was used as input to the simulated double integrator system. At start-up, force/torque limits were place on each degree of freedom.

A host computer was used to communicate with Sawyer during setup and operation. Using the core software architecture of the Robot Operating System (ROS), the Host received position and acceleration information from Sawyer. The host also sent messages setting the parameters of the Sawyer impedance model and controller. Information from the Sawyer was used to visualize the interaction point as a 3D cursor and the drawing history as a series of dots in the ROS visualization package, rviz. The position information also kinematically controlled the double integrator system being simulated by the host computer. The host set the parameters to either increase of decrease the impedance at the end effector or modify the forces at the end effector according to either Equation 5.9 or Equation 5.10.

At the beginning of the session, the drawing task was demonstrated to the participants by the authors, and participants were able to practice drawing on the screen using the robot as a cursor. Participants performed a baseline set of trials in which they drew each image 10 times for a total of 40 trials. The order in which they completed these drawings was randomized to minimize learning during the baseline trials. After the baseline set

of trials, participants trained with their assigned control strategy completing both the training and post-training trials for one image before moving on to the next image.

Subjects were recruited locally (n=24), and had to be healthy, able-bodied adults (in the age range of 18 to 50) with no prior history of upper limb or cognitive impairments. Only right-hand dominant participants were accepted into the study, and each subject performed the task with their right limb. All study protocols were reviewed and approved by the Northwestern University Institutional Review Board, and all subjects gave written informed consent prior to participation in the study.

5.5.3. Measurements & Statistical Analysis

We assess user performance using the metrics that close the loop in the two control strategies that were tested—error and ergodicity. The data for each image consisted of 10 baseline trials, 10 trials with either ergodic HSC or virtual fixtures, and 10 trials post-training for a total of 30 trials for each of the four images. These were grouped into sets of 10 trials to evaluate subject performance over time. The analysis consisted of two-factor (set and group) repeated measures ANOVA tests. The ANOVA's were used to compare the effect of the ergodic HSC and virtual fixture training on each of the performance measures. When significant main effects or interaction effects were detected, student's t-tests were used to evaluate the difference between the performance of the ergodic HSC group and the control group.

5.5.3.1. Error. Every t_s seconds, we measured the position of the robot end-effector and translate it to image coordinates on the domain $[0, 2200] \times [0, 2200]$. We then perform a search for the nearest dark pixel (saturation$<$ 130). The distance between the end-effector

position and the nearest dark pixel is recorded. The error measure that we report here is the mean distance from the nearest dark pixel computed for each trial.

5.5.3.2. Ergodicity. We treat each image as a discrete histogram over the domain and generate 100 random samples from that distribution as can be seen in Figure 5.3. We use these to calculate the ergodic metric according to Equation 5.2–giving us the trajectory's distance from ergodicity for each image.

5.5.3.3. Completeness. Scorers were asked to provide a rating evaluating the completeness of each drawing on a scale from 1 to 100. Each participant drawing was assigned a random code and randomly assigned to one of eight scorers via an online survey. Scorers were instructed not to judge the quality of the drawing on the basis of scale or accuracy. Each image was recieved at least one rating.

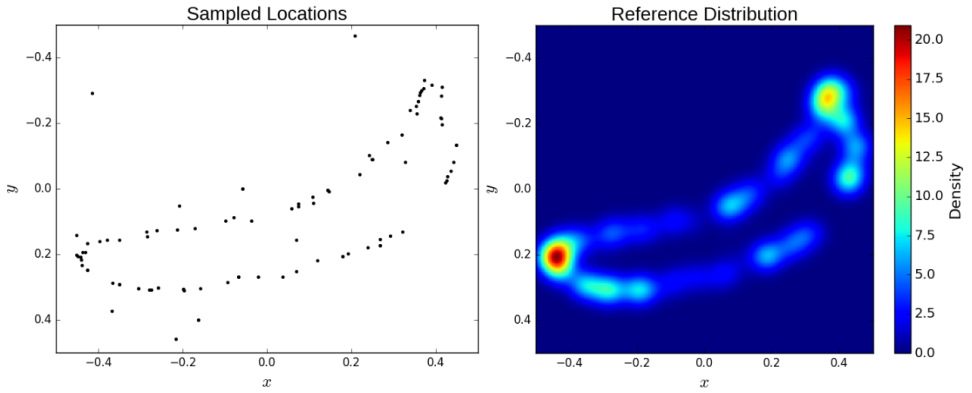

Figure 5.3. Using importance sampling, 100 sample locations are randomly chosen from the original image. The color intensity at each point is used to weight a series of Gaussian distributions that generate the reference distribution.

5.6. Experimental Results

The results were reported as follows. First, the error of each group was statistically tested in Section 5.6.1. An analysis of the ergodicity was performed to test for differences in the relative information communicated in the drawings of the HSC group and the VF group in Section 5.6.2. Finally, an analysis of the completion scores of each group was performed in Section 5.6.3. The results demonstrated that training with the ergodic HSC increased subject performance in later trials within the same session. In each section the relevant statistics are reported first, followed by a summary and interpretation of the results.

5.6.1. Error Measure

The mean error of each group in each set can be seen in Figure 5.4. The progress of the two groups over the training session was analyzed by performing mixed design ANOVAs on the HSC group (between participants) and set (within participants) using the error on all four images. Only the baseline trials (set 1) and post-training trials (set 3) were used to avoid measuring the effects of the assistance itself in the analysis.

The mean error of the apple drawings had two significant factors. The main effect of training group was not significant ($p = 0.636,\ F(1, 20) = 0.231$). However, the main effect of set was significant ($p = 5.58 \times 10^{-7},\ F(1, 454) = 25.784$), as was the interaction of training and set ($p = 0.0272,\ F(1, 454) = 4.913$). Interestingly, study participants in the VF group increased their average distance from a dark pixel both during and after training, whereas participants using HSC had similar levels of error in set 1 and set 3.

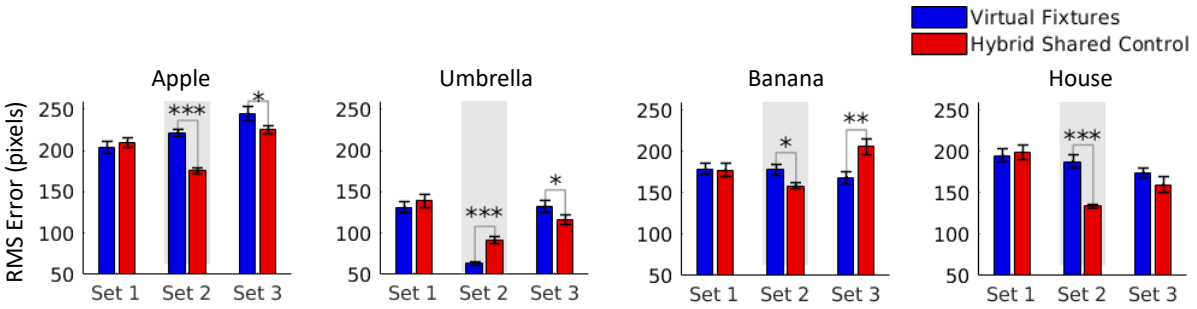

Figure 5.4. Although both groups performed similarly in terms of error in the baseline set, the group that trained with HSC had lower error when drawing the apple, umbrella, and house in the post-training trials. Interestingly, the virtual fixtures were activated based on the error measure, but the umbrella was the only image in which virtual fixtures actually reduced errors in the drawings. Note: error bars indicate standard error; significance is indicated by $^{*}p < 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$.

The mixed design ANOVA design was also applied to the error in the banana drawings, and the main effect of training group was not significant ($p = 0.4611,\ F(1, 20) = 0.565$). The main effect of set was not significant either ($p = 0.1132,\ F(1, 454) = 2.514$). The interaction effect of block and training group was significant ($p = 0.00202,\ F(1, 454) = 9.643$). This reflects the fact that the two groups performed similarly in the first set, but the VF group had lower error in the post-training set.

When the same mixed design ANOVA was applied to the error in the umbrella drawings, the main effects of set ($p = 0.071,\ F(1, 454) = 3.27$) and group ($p = 0.811,\ F(1, 20) = 0.059$) were not significant. The interaction effect of training group and set was significant ($p = 0.0495,\ F(1, 454) = 3.88$). In the case of this drawing, the VF group had higher error in the post-training trial compared to the HSC group, though the two groups had similar baseline error.

The analysis of the error in the drawings of the house revealed a significant main effect of set ($p = 1.23 \times 10^{-5},\ F(1, 454) = 19.545$), but the interaction effect of set and training

group ($p = 0.46,\ F(1, 454) = 0.546$) was not significant. The main effect of group also was not significant ($p = 0.728,\ F(1, 20) = 0.124$). As in the drawings of the apple and umbrella, the HSC group had lower error than the VF group in the post-training set, though they started at the same baseline error.

In 3 of the 4 drawings, the group that trained with virtual fixtures, which were designed to reduce error, actually performed worse during set 3 in terms of error compared to the group that trained with HSC. When we look at Figure 5.4, we can see that even when the virtual fixtures were engaged in set 2, the HSC group had lower error than the VF group when drawing each image except the umbrella. *These results demonstrate that even when feedback is based on spatial statistics, other standard measures like error can be improved though they are not directly targeted by the algorithm.*

One reason that error increases when participants train with virtual fixtures is that participants exploit these guides when they are present. For the participant shown in Figure 5.5, it is clear when they were drawing the apple, banana, and house, that they found the virtual wall and followed it such that they maintain a consistent distance from the desired lines. When the virtual fixtures are removed, this bias remains. The offsets in the post-training drawings are similar to those we see in the drawings with virtual fixtures.

5.6.2. Ergodic Measure

Two-factor repeated measures ANOVAs were used to assess the effects of the group (between-subjects) and set (within-subjects) on the ergodic measure defined Section 5.5.3.2 for each image used in the study. The HSC group and VF group were evaluated based on

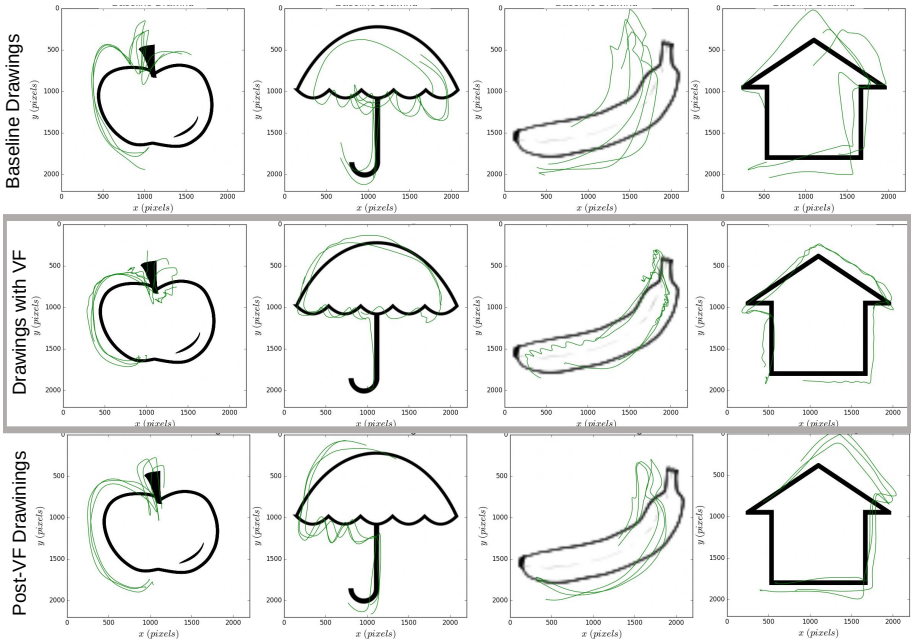

Figure 5.5. Example drawings from Participant 10 training with virtual fixtures. Each plot contains the desired image as well as three drawings by the participants. In drawings with virtual fixtures, we see two notable responses to the guidance. Wall following is evident in the drawings of the apple, banana, and house. There are also oscillations when bumping into and out of the wall particularly in the drawings of the banana and the stem of the apple.

the baseline trials (set 1) and the post-training trials (set 3) only. Set 2 was left out of the ANOVA, so that effects of the assistance itself would not be measured in the analysis.

The factorial ANOVA of the ergodic measure on the apple image revealed that the interaction effect of group and set was the only significant factor ($p = 6.17\times10^{-4},\ F(1, 454) = 11.889$). The main effects of group and set were not significant for the apple drawings ($p > 0.05$). The HSC and VF group performed similarly in the baseline trials, but the HSC group performed slightly better after the training set. When an analysis of variance was performed on the ergodicity of the banana drawings, again, there was no significant

effect of group, set, or the interaction of those two factors ($p > 0.05$). When the ergodicity of the trajectories drawing the umbrella were compared, the significant factors were the set ($p = 6.21 \times 10^{-5},\ F(1, 454) = 16.34$) and the interaction between group and set ($p = 1.95 \times 10^{-4},\ F(1, 454) = 14.11$). The main effect of group was not significant ($p = 0.613,\ F(1, 20) = 0.264$). The group was not a significant factor affecting the error in the house drawings ($p = 0.238,\ F(1, 20) = 1.477$). The main effect of set ($p = 0.73,\ F(1, 454) = 0.119$) also was not significant, but the interaction of group and set ($p = 7.90 \times 10^{-8},\ F(1, 454) = 29.795$) were significant.

The results of the ANOVA of ergodicity for 3 of the 4 drawings showed that the interaction effect of set and group was a significant factor—implying that while the participants started at the same performance level in their baseline set, participants in the HSC group attained a higher performance level in the post-training set than the VF group. The differences in the ergodic measure implies that participants training with HSC generated trajectories that encoded more information about the original image. This is likely due

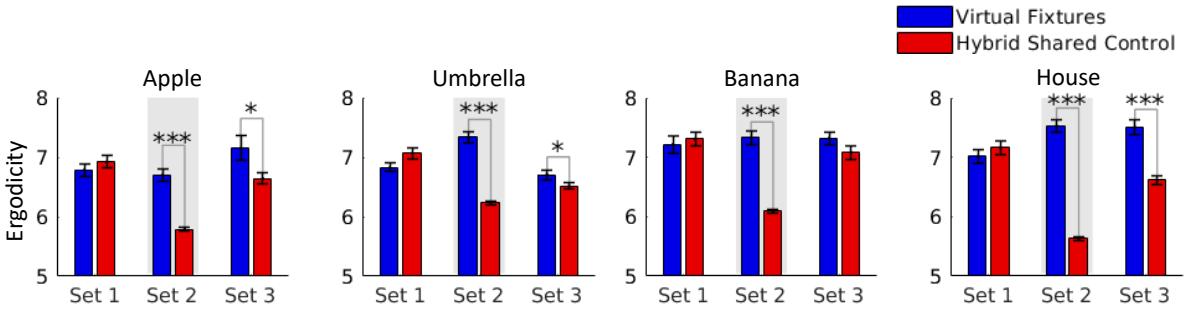

Figure 5.6. Participants using Hybrid Shared Control in set 2 generated drawing trajectories that were more ergodic with respect to the desired image than those using virtual fixtures. This advantage was maintained in set 3 when the assistance was removed in the apple, umbrella, and house drawing. The difference between the two groups for set 3 of the banana drawings was not significant. Note: error bars indicate standard error; significance is indicated by $^{*}p < 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$.

to the assistance intervening based on a measure of overall performance as opposed to the local distance measure employed by the virtual fixtures. The virtual fixtures generally led participants to draw more slowly—not completing the image. Participants receiving feedback from HSC drew images that were smaller, but more complete as can be seen in the examples in Figure 5.7.

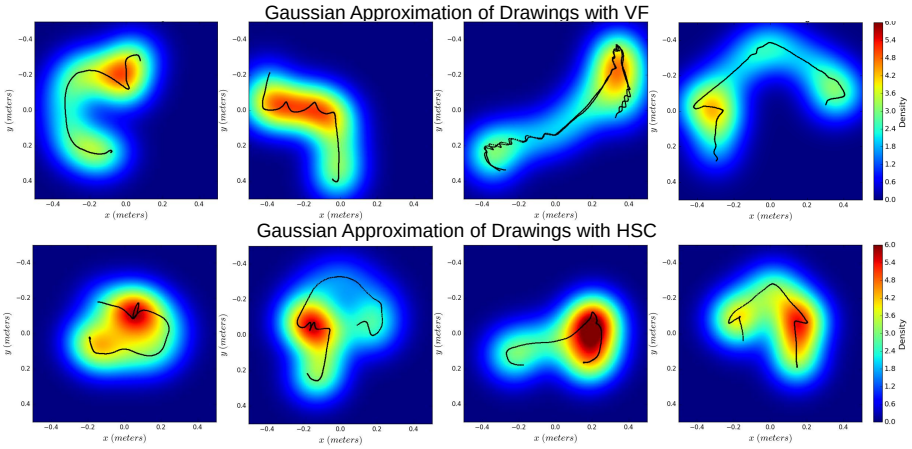

Figure 5.7. Representative examples of trajectories from Participant 10 using virtual fixture (top) and Participant 6 using hyrbid shared control (bottom). The time-average statistics of the trajectories are plotted using the Gaussian approximation from Equation 5.1. The trajectories in the top row are taken from the middle row of Figure 5.5. However, here we plot them over a representation of their spatial statistics. There is high density around the trajectory, particularly points where participants moved more slowly, such as corners on the house or cusps on the apple and umbrella. For Participant 6 using HSC (bottom), we see that they were generally moving more quickly and covered more of the original image than the drawings with VF. This example demonstrates how HSC participants drew smaller, but more complete drawings than the participants that trained with virtual fixtures. The difference between the two distributions demonstrates how the drawings with HSC contained more information about the original image than the drawings with VF.

5.6.3. Completion Score

The mean completion score of each group in each set can be seen in Figure 5.8. The change in completion percentage of the two groups over the course of training was analyzed by performing mixed design ANOVAs on the HSC group (between participants) and set (within participants) using the ratings on all four images. Only the baseline trials (set 1) and post-training trials (set 3) were used to avoid measuring the effects of the assistance itself in the analysis.

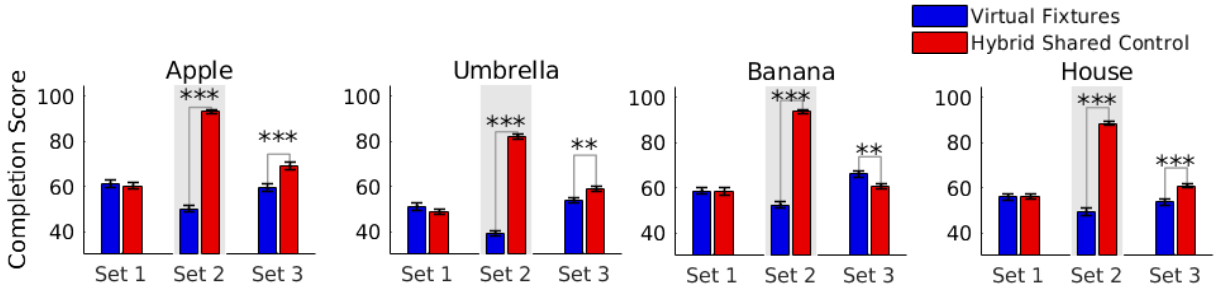

Figure 5.8. The virtual fixtures actually reduced completeness score in set 2 for every image except the banana, while the HSC greatly increased the completion scores. Therefore, the group that trained with HSC had much higher completion scores when drawing the apple, umbrella, and house in the post-training trials, even though both groups performed similarly in terms of completion score in the baseline set. Note: error bars indicate standard error; significance is indicated by $^{*}p < 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$.

The completion percentage of the apple drawings had two significant factors. The main effect of training group was not significant ($p = 0.400,\ F(1, 20) = 0.739$). However, the main effect of set was significant ($p = 0.00245,\ F(1, 460) = 9.276$), as was the interaction of training and set ($p = 5.56 \times 10^{-}5,\ F(1, 460) = 16.556$). Study participants in the VF group completed around the same amount of this drawing both before and after training, whereas participants using HSC completed 7% more on average.

The mixed design ANOVA was also applied to the completion in the banana drawings, and the main effect of training group was not significant ($p = 0.393,\ F(1, 20) = 0.761$). However, the main effect of set was significant ($p = 2.200 \times 10^{-5},\ F(1, 470) = 18.373$). The interaction effect of block and training group was significant ($p = 0.0497,\ F(1, 470) = 3.871$). This reflects the fact that the two groups performed similarly in the first set, but the VF group had completed more of the drawings in the post-training set. Both groups improved their completion scores post-training.

When the same mixed design ANOVA was applied to the completion scores in the umbrella drawings, the main effect of set ($p = 1.82 \times 10^{-8},\ F(1, 477) = 32.79$) and the interaction effect of training group and set were significant ($p = 0.005,\ F(1, 477) = 7.67$). The main effect of group ($p = 0.662,\ F(1, 20) = 0.197$) was not significant. In the case of this drawing, the VF group and the HSC group had higher completion scores in the post-training trials compared to their baseline. Although the two groups had similar baseline completion scores, the HSC group achieved significantly higher scores that the VF group post-training.

The analysis of the error in the drawings of the house revealed that the main effects of set ($p = 0.162,\ F(1, 455) = 1.963$) and group($p = 0.284,\ F(1, 20) = 1.210$) were not significant, but the interaction effect of set and training group ($p = 0.003,\ F(1, 455) = 9.190$) was significant. As in the drawings of the apple and umbrella, the HSC group had higher completion scores than the VF group in the post-training set, though they started at the same baseline error. In this drawing in particular, the completion scores of the VF group actually went down in the post-training trials.

In 3 of the 4 drawings, the group that trained with virtual fixtures, had lower completion scores on their drawings compared to the group that trained with HSC. When we look at Figure 5.8, we can see that the HSC group had much higher completion scores in set 2 and retained a modest advantage over the VF group in the post-training trials. *These results demonstrates that the HSC encouraged participants to take actions that improved the overall quality of the drawing rather than the accuracy of an individual pose.*

5.7. Discussion & Conclusions

There are numerous articles stating the potential of robotics to support training and rehabilitation because of their ability to assist users in completing many repetitions and their ability to provide quantitative feedback. The questions of how robots should assist/resist users, how to define the task, and what metrics they should use quantify success have a profound impact on the efficacy of training. Prior studies show that error, variability, and active user participation—achieved by adapting the robot support as needed—are crucial to motor learning. Furthermore, we know that there are many equally good solutions to a particular task that a person might use.

We have developed a hybrid shared controller that selectively rejects or accepts user actions based on how that action will affect the time-averaged statistics of the trajectory for some time into the future—that is, how the amount of task information present in the trajectory will be affected. Using this controller, a robot can provide physical corrective feedback during training while avoiding issues of time-dependence and selection of a particular strategy to complete a task. When user inputs increase the ergodicity of

the trajectory with respect to a distribution defining the task, the controller is transparent to the user. Otherwise, user inputs are rejected by providing an equal and opposite force—maintaining a constant velocity. In our study, we experimentally compare this novel assistance paradigm, ergodic hybrid shared control, to a standard form of assistance based on error.

Our results demonstrated that although ergodic HSC is based on a global measure of the distance from ergodicity rather than a local measure of error, it improved the error measure. Participants who trained with the error-based assistance actually performed worse in terms of error than those who trained with ergodic HSC. Because the virtual fixtures provide feedback based on distance from a local point on the desired path, there is a tendency for participants to follow the virtual fixtures. This leads to drawings that are precise but not accurate, following the same incorrect path over multiple trials. This emphasizes the fact that error is a limited measure that cannot capture broader task goals. Although error was higher in the VF trained participant drawings, one would not say that their skill in reproducing the drawings was necessarily poor. If the reference had been built off many example drawings, they likely would have fallen within range of one or more of the examples. Yet, the need to define a specific path means that some participants are penalized for being arbitrarily far from the mean.

When evaluating the impact of training in terms of the time-averaged statistics of the participants' trajectories, we found that there was a significant advantage to training with ergodic HSC. The group that trained with virtual fixtures produced consistently incomplete drawings, whereas the group receiving feedback based on the ergodic measure produced drawings that were smaller than the original image, but more complete. This

difference emphasizes the need for assistance based on global measures as opposed to local interactions. Combining the ergodic measure with the dynamics of the system in the MIG, allows ergodic HSC to be sensitive to time without being dependent on time like some other assistance strategies.

Timed drawing is not a daily task that people need assistance with or training for. Nevertheless, it shares characterstics with tasks such as cleaning, personal care, cooking, and reaching. Moreover, it can be used in a Learning from Demonstration (LfD) framework [51], so that erdogic HSC can be easily implemented for other tasks. When the task is unknown, one could potentially use intent detection to define a representation that encodes the uncertainty about the intent prediction.

CHAPTER 6

Conclusions and Future Directions

This thesis began with a discussion of the impact that metrics have on how we assess motion and drive our autonomous systems. The choice of metric affects our ability to distinguish between normal and impaired movement, and more importantly from a robotics standpoint, it affects our capacity to design systems that interact closely with people. One of the most frequently used measures in robotics is error with respect to a desired trajectory. This requires us to specify a specific strategy for movement to complete a task even though there may be many good-enough solutions. In chapter 3, I avoid specifying a trajectory by using methods from model predictive control to evaluate the user input rather than the trajectory itself. Evaluating inputs instantaneously to decide whether to accept or reject lead to improved training outcomes within a session and improved retention after one week. This is likely because task-based hybrid shared control had many of the features of pHRI that we know are critical to motor learning. However, it was still fundamentally based on error, with a state-based task definition.

When the goal of pHRI is to support stereotypical human motions, the ideal objective or reference is statistical patterns of movement. Trajectories or paths can easily over specify a task, and when used to classify expert/novice or normal/impaired motion, they may not provide reflect the significant differences between groups. Metrics based on concepts from information theory can provide us with a way to quantify the differences or similarities in statistical patterns of motion such that task performance is accurately classified.

The ergodic metric described in chapter 4 is not the only possible way to quantify the information encoded in motion. There are many representations of a distribution that could be used, including wavelets or Gaussian mixtures, as long as a comparison metric of two distributions can be defined. However, pHRI requires that the robot to be physically synchronous with the person, so metrics we choose must be actionable for the autonomy (i.e. computed online and differentiable) and should enable the robot to maintain a task-level view of performance as opposed to managing local interactions. The metric used to assist participants with drawing in the experiments described in chapter 5 was able to accomplish this such that participants learned to draw more of each image in a shorter time compared to participants who received assistance proportional to their distance from the drawing path.When we design algorithm that can maintain a nonmyopic view of performance at each instant in time as in chapter 3, we get desirable features for enhancing motor leaning and users receive feedback that represents global task information rather than local errors.

6.1. Future Directions

In the experimental work in chapter 5, we assume that we have a known distribution that is fixed in time. The tasks presented in chapter 4 and chapter 5 can be reasonably described by a single distribution, but we cannot expect this to hold for all tasks nor can we assume that they are easily described by parametric density functions like a delta function or Gaussian distribution. However, we can generate reference distributions using Learning from Demonstration (LfD). In [51], we show that there is a straight forward way to combine the Fourier coefficients of multiple trajectories to form a reference set

of coefficients. Furthermore, one can use ergodic optimal control [1, 81] to generate task executions that outperform individual demonstrations. These results suggest that the combined demonstrations capture the essential components of the task. A similar approach could be adopted to define the reference distribution in Algorithm 2.

Algorithm 2 assumes that the reference distribution ϕ does not depend on time. The static reference distribution that ergodic LfD provides can be applied to many tasks of interest, particularly task of daily living such as cooking and cleaning that may be relevant in therapeutic settings. Part of the advantage of statistical definitions of tasks is that there is not a direct dependence on time, however, there are situations in which timing is significant. When a task is performed in a way that changes the dynamic model of the task, we should expect that a new reference is necessary for each change. That is when an object is picked up, adding mass, or contacts are made that induce new constraints, the new dynamics require a new reference distribution. For these types of more complex tasks, I anticipate that it would be necessary to segment sequential tasks like peg-in-hole insertion or block stacking into sub tasks using machine learning. There would also be value in generating or updating distributions based on intent. This would involve applying intent detection methods such as those in [38, 48, 49] and [65, 144] to distributions as opposed to joint torques or trajectories, respectively.

Throughout this thesis, I use examples of 2-dimensional tasks. Certain aspects of extending the ergodic measure and control to higher dimensional tasks are trivial. The sample-based metric described in [1] resolves the computational issues that arise with the spectral approach, however, the choice of the parameters, Σ and N can significantly affect the behavior of the ergodic control. These two parameters are strongly related in that

to get an accurate representation of the original distribution without random artifacts related to sampling, there must either be a large number of sample locations, N, with small variance, Σ or a smaller N with large variance Σ. The latter representation would be a coarser representation of the task, such that less precision is required to generate a trajectory that is highly ergodic with respect to the representation. In the case of the drawing task in chapter 5, this actually enables the shared control to generate drawings that are complete but not to scale in the time allowed. If the task required more precision in terms of the state, we would choose a lower value of Σ and would need to increase the number of samples to avoid having erroneously low densities in areas of the state space that were not sampled as often as others. For planar tasks, where the desired distribution and its Gaussian mixture representation can be visualized, one can manually select these two parameters to manage the trade-off between solution flexibility and task accuracy. In higher dimensional settings, these parameters may be very difficult to choose especially when the reference is a result of many demonstrations of a task where the sensitivity of the task to the accuracy of the movements is unknown. Therefore, in future work, the parameters of the ergodic metric should be auto-tuned such that the information lost by using the randomly sampled Gaussian mixture representations instead of the true distribution is minimized subject to constraints on the maximum number of samples that can be used in the online controller.

In the work described in chapter 4, I compare the statistical outcomes of measures with regard to different levels of assistance in stroke and healthy participants. Increasing levels of assistance can be thought of as decreased deficit in this context. However in a clinical setting, we would be interested in predicting deficit across healthy, mild, and

moderate impairments. In future work, a rigorous comparison of clinical measures and task measures should be conducted across multiple tasks and deficit level using healthy age-matched controls. This comparison would characterize the extent to which one can use information measures to predict deficit.

Up until this point, I have discussed immediate extensions of the work in this thesis, but the information theoretic view of motion that is taken in this work opens up more exciting questions about the tasks that we choose and the mode of communication between the autonomy and the user. The commercial availability of wearable sensors, actuation, and virtual reality systems makes it possible to envision autonomous systems capable of assessing the motion of aging and impaired populations as well as prescribing therapeutic exercises to mitigate disability. As of now, we lack the algorithmic foundations to map measures of the user's state to available actuation. Future work should focus on creating methods for quantitatively assessing users based on measures that are actionable for the automation. Utilizing methods for system identification and active exploration as opposed to hand-crafted solutions, one could automatically generate exercises to resolve uncertainties about a user's impairment. Therapy regimens can also be designed to optimally challenge users by using quantitative measures in a framework informed by knowledge and expertise from the areas of human factors and psychology. Augmented or virtual reality systems can be developed to place these automatically generated tasks into a motivating real-world context. I have already shown that measurement of the information encoded by a motion about a task can discriminate well between different levels of impairment using a predefined task. It would be interesting to continue to use this information-theoretic view of motion to develop and apply measures of motion quality by

framing the assessment and design of challenging tasks as information-maximizing problems. In assessment, we will ask how we can maximize our confidence in our assessment by directing users to tasks that allow us to discriminate between different types and levels of impairments. Additionally, we can challenge users by developing tasks that exploit this assessment—ruling out tasks where we know users already have high task-embodiment.

Perhaps the largest hurdle in making such systems effective, especially for highly-impaired individuals, is the sensorimotor feedback that could be provided through haptic and robotic systems. For at-home semi-autonomous therapy to be widely adopted, devices must be low cost and low power. For instance, an ideal device might be a haptic arm band that can provide feedback through skin stretch, squeeze, and vibration. This necessarily limits bandwidth, so we must learn how we can compress task-specific motion information to generate low-dimensional feedback that enhances task performance and training. In a one-dimensional task, we have seen that users can improve their task performance when vibrotactile feedback is introduced into a balance task[128], because the feedback chosen was simple and intuitive. For higher dimensional tasks and multimodal haptic or kinesthetic feedback, the online synthesis of feedback is no longer trivial. If we connect principles from communication engineering to motor control, we can distill the task-relevant signals such that they are easily interpreted and incorporated with the native sensory signals of users to improve their task performance for more complex tasks.

References

[1] Ian Abraham, Ahalya Prabhakar, and Todd D Murphey. An ergodic measure for active learning from equilibrium. *Transactions on Automation Science and Engineering*, 2020.

[2] Henny Admoni and Brian Scassellati. Social eye gaze in human-robot interaction: a review. *Journal of Human-Robot Interaction*, 6(1):25–63, 2017.

[3] Erwin Aertbeliën and Joris De Schutter. Learning a predictive model of human gait for the control of a lower-limb exoskeleton. In *IEEE Int. Conf. on Biomedical Robotics and Biomechatronics*, pages 520–525, 2014.

[4] Alexander R Ansari and Todd D Murphey. Sequential action control: closed-form optimal control for nonlinear and nonsmooth systems. 32(5):1196–1214, 2016.

[5] Daisuke Aoyagi, Wade E Ichinose, Susan J Harkema, David J Reinkensmeyer, and James E Bobrow. A robot and control algorithm that can synchronously assist in naturalistic motion during body-weight-supported gait training following neurologic injury. 15(3):387–400, 2007.

[6] Brenna D Argall, Sonia Chernova, Manuela Veloso, and Brett Browning. A survey of robot learning from demonstration. *Robotics and Autonomous Systems*, 57(5):469–483, 2009.

[7] Bruno B Averbeck, Peter E Latham, and Alexandre Pouget. Neural correlations, population coding and computation. *Nature Reviews Neuroscience*, 7(5):358, 2006.

[8] Sivakumar Balasubramanian, Ruihua Wei, Mike Perez, Ben Shepard, Edward Koeneman, James Koeneman, and Jiping He. Rupert: An exoskeleton robot for assisting rehabilitation of arm functions. In *2008 virtual rehabilitation*, pages 163–167. IEEE, 2008.

[9] Jérémy Bluteau, Sabine Coquillart, Yohan Payan, and Edouard Gentaz. Haptic guidance improves the visuo-manual tracking of trajectories. *PLoS One*, 3(3):e1775, 2008.

[10] Stuart A Bowyer, Brian L Davies, and Ferdinando Rodriguez y Baena. Active constraints/virtual fixtures: A survey. *IEEE Transactions on Robotics*, 30(1):138–157, 2013.

[11] Charles G Burgar, Peter S Lum, Peggy C Shor, and HF Machiel Van der Loos. Development of robots for rehabilitation therapy: The Palo Alto VA/Stanford experience. *J. Rehabilitation Research and Development*, 37(6):663–674, 2000.

[12] Timothy M Caldwell and Todd D Murphey. Projection-based optimal mode scheduling. *Nonlinear Analysis: Hybrid Systems*, 21:59–83, 2016.

[13] B Cesqui, S Aliboni, S Mazzoleni, MC Carrozza, F Posteraro, and S Micera. On the use of divergent force fields in robot-mediated neurorehabilitation. In *IEEE Int. Conf. on Biomedical Robotics and Biomechatronics (BioRob)*, pages 854–861. IEEE, 2008.

[14] Steve Collins, Andy Ruina, Russ Tedrake, and Martijn Wisse. Efficient bipedal robots based on passive-dynamic walkers. *Science*, 307(5712):1082–1085, 2005.

[15] Gery Colombo, Matthias Joerg, Reinhard Schreier, and Volker Dietz. Treadmill training of paraplegic patients using a robotic orthosis. *J. Rehabilitation Research and Development*, 37(6):693–700, 2000.

[16] Janis J Daly and Jonathan R Wolpaw. Brain–computer interfaces in neurological rehabilitation. *The Lancet Neurology*, 7(11):1032–1043, 2008.

[17] Agostino De Santis, Bruno Siciliano, Alessandro De Luca, and Antonio Bicchi. An atlas of physical humanrobot interaction. *Mechanism and Machine Theory*, 43(3):253–270, 2008.

[18] Bruce H. Dobkin and Pamela W. Duncan. Should body weightsupported treadmill training and robotic-assistive steppers for locomotor training trot back to the starting gate? *Neurorehabilitation and Neural Repair*, 26(4):308–317, 2012.

[19] Anca D Dragan and Siddhartha S Srinivasa. A policy-blending formalism for shared control. *Int. J. Robotics Research*, 32(7):790–805, 2013.

[20] Vincent Duchaine, Nicolas Lauzier, Mathieu Baril, Marc-Antoine Lacasse, and Clément Gosselin. A flexible robot skin for safe physical human robot interaction. In *2009 IEEE International Conference on Robotics and Automation*, pages 3676–3681. IEEE, 2009.

[21] Alexander Duschau-Wicke, Joachim von Zitzewitz, Andrea Caprez, Lars Lunenburger, and Robert Riener. Path control: a method for patient-cooperative robot-aided gait rehabilitation. 18(1):38–48, 2010.

[22] Magnus Egerstedt, Yorai Wardi, and Henrik Axelsson. Transition-time optimization for switched-mode dynamical systems. *IEEE Trans. Automatic Control*, 51(1):110–115, 2006.

[23] Michael D Ellis, Yiyun Lan, Jun Yao, and Julius PA Dewald. Robotic quantification of upper extremity loss of independent joint control or flexion synergy in individuals with hemiparetic stroke: a review of paradigms addressing the effects of shoulder abduction loading. *J. NeuroEngineering and Rehabilitation*, 13(1):95, 2016.

[24] Michael D. Ellis, Theresa Sukal-Moulton, and Julius P. A. Dewald. Progressive shoulder abduction loading is a crucial element of arm rehabilitation in chronic stroke. *Neurorehabilitation and Neural Repair*, 23(8):862–869, 2009. PMID: 19454622.

[25] Jeremy L Emken, Raul Benitez, Athanasios Sideris, James E Bobrow, and David J Reinkensmeyer. Motor adaptation as a greedy optimization of error and effort. *J. Neurophysiology*, 97(6):3997–4006, 2007.

[26] Jeremy L Emken, Susan J Harkema, Janell A Beres-Jones, Christie K Ferreira, and David J Reinkensmeyer. Feasibility of manual teach-and-replay and continuous impedance shaping for robotic locomotor training following spinal cord injury. 55(1):322–334, 2008.

[27] Jeremy L Emken and David J Reinkensmeyer. Robot-enhanced motor learning: accelerating internal model formation during locomotion by transient dynamic amplification. *IEEE Trans. Neural Syst. Rehabil. Eng.*, 13(1):33–39, 2005.

[28] M Ferraro, JJ Palazzolo, J Krol, HI Krebs, N Hogan, and BT Volpe. Robot-aided sensorimotor arm training improves outcome in patients with chronic stroke. *Neurology*, 61(11):1604–1607, 2003.

[29] Moria E Fisher, Felix C Huang, Zachary A Wright, and James L Patton. Distributions in the error space: Goal-directed movements described in time and state-space representations. In *IEEE Int. Conf. on Engineering in Medicine and Biology*, pages 6953–6956, 2014.

[30] Kathleen Fitzsimons, Ana Maria Acosta, Julius P Dewald, and Todd D Murphey. Ergodicity reveals assistance and learning from physical human-robot interaction.

Science: Robotics, 4, 2019.

[31] Kathleen Fitzsimons, Aleksandra Kalinowska, Julius Dewald, and Todd D Murphey. Task-based hybrid shared control for training through forceful interaction. *Int. J. Robotics Reserach*, 39(9):1138–1154, 2020.

[32] Kathleen Fitzsimons, Emmanouil Tzorakoleftherakis, and Todd D Murphey. Optimal human-in-the-loop interfaces based on maxwell's demon. In *American Control Conference*, pages 4397–4402, July 2016.

[33] Tamar Flash and Neville Hogan. The coordination of arm movements: an experimentally confirmed mathematical model. *J. Neuroscience*, 5(7):1688–1703, 1985.

[34] Antonio Frisoli, Claudio Loconsole, Daniele Leonardis, Filippo Banno, Michele Barsotti, Carmelo Chisari, and Massimo Bergamasco. A new gaze-bci-driven control of an upper limb exoskeleton for rehabilitation in real-world tasks. *IEEE Transactions on Systems, Man, and Cybernetics, Part C (Applications and Reviews)*, 42(6):1169–1179, 2012.

[35] Humberto Gonzalez, Ram Vasudevan, Maryam Kamgarpour, Shankar S. Sastry, Ruzena Bajcsy, and Claire Tomlin. A numerical method for the optimal control of switched systems. In *IEEE Conf. on Decision and Control*, pages 7519–7526, 2010.

[36] Mark Guadagnoli and Kristina Lindquist. Challenge point framework and efficient learning of golf. *Int. J. Sports Science & Coaching*, 2(1_suppl):185–197, 2007.

[37] Marco Guidali, Alexander Duschau-Wicke, Simon Broggi, Verena Klamroth-Marganska, Tobias Nef, and Robert Riener. A robotic system to train activities of daily living in a virtual environment. *Medical & biological engineering & computing*, 49(10):1213, 2011.

[38] Abhishek Gupta, Marcia K O'Malley, Volkan Patoglu, and Charles Burgar. Design, control and performance of ricewrist: a force feedback wrist exoskeleton for rehabilitation and training. *Int. J. Robotics Research*, 27(2):233–251, 2008.

[39] Christopher M Harris and Daniel M Wolpert. Signal-dependent noise determines motor planning. *Nature*, 394(6695):780–784, 1998.

[40] Stefan Hesse, Gotthard Schulte-Tigges, Matthias Konrad, Anita Bardeleben, and Cordula Werner. Robot-assisted arm trainer for the passive and active practice of bilateral forearm and wrist movements in hemiparetic subjects. *Archives of physical medicine and rehabilitation*, 84(6):915–920, 2003.

[41] Neville Hogan. An organizing principle for a class of voluntary movements. *J. Neuroscience*, 4(11):2745–2754, 1984.

[42] Neville Hogan, Hermano I Krebs, Brandon Rohrer, Jerome J Palazzolo, et al. Motions or muscles? some behavioral factors underlying robotic assistance of motor recovery. *Journal of Rehabilitation Research and Development*, 43(5):605, 2006.

[43] Yu Wei Hsieh, Ching Yi Wu, Keh Chung Lin, Ya Fen Chang, Chia Ling Chen, and Jung Sen Liu. Responsiveness and validity of three outcome measures of motor function after stroke rehabilitation. *Stroke*, 40(4):1386–1391, 2009.

[44] Vincent S Huang and John W Krakauer. Robotic neurorehabilitation: a computational motor learning perspective. *J. NeuroEngineering and Rehabilitation*, 6(1):5, 2009.

[45] Alberto Isidori. A tool for semi-global stabilization of uncertain non-minimum-phase nonlinear systems via output feedback. *IEEE Trans. on Automatic Control*, 45(10):1817–1827, 2000.

[46] Siddarth Jain and Brenna Argall. Probabilistic human intent recognition for shared autonomy in assistive robotics. *ACM Transactions on Human-Robot Interaction (THRI)*, 9(1):1–23, 2019.

[47] William H Jantscher, Shivam Pandey, Priyanshu Agarwal, Sadie H Richardson, Bowie R Lin, Michael D Byrne, and Marcia K O'Malley. Toward improved surgical training: Delivering smoothness feedback using haptic cues. In *IEEE Haptics Symposium*, pages 241–246. IEEE, 2018.

[48] LE Kahn, WZ Rymer, and DJ Reinkensmeyer. Adaptive assistance for guided force training in chronic stroke. In *IEEE Int. Conf. Engineering in Medicine and Biology*, volume 1, pages 2722–2725, 2004.

[49] Leonard E Kahn, Michele L Zygman, W Zev Rymer, and David J Reinkensmeyer. Robot-assisted reaching exercise promotes arm movement recovery in chronic hemiparetic stroke: a randomized controlled pilot study. *J. Neuroengineering and Rehabilitation*, 3(1):12, 2006.

[50] Aleksandra Kalinowska, Kathleen Fitzsimons, Julius P Dewald, and Todd D Murphey. Online user assessment for minimal intervention during task-based robotic assistance. In *Robotics: Science and Systems*, 2018.

[51] Aleksandra Kalinowska, Ahalya Prabhakar, Kathleen Fitzsimons, and Todd D Murphey. Ergodic lfd: Learning from what to do and what not to do. In *Robotics and*

Automation Letters, Submitted. 2021.

[52] Behzad Khademian and Keyvan Hashtrudi-Zaad. Shared control architectures for haptic training: Performance and coupled stability analysis. *Int. J. Robotics Research*, 30(13):1627–1642, 2011.

[53] Tomoko Kitago, Johnny Liang, Vincent S Huang, Sheila Hayes, Phyllis Simon, Laura Tenteromano, Ronald M Lazar, Randolph S Marshall, Pietro Mazzoni, Laura Lennihan, and John W Krakauer. Improvement after constraint-induced movement therapy: recovery of normal motor control or task-specific compensation? *Neurorehabilitation and Neural Repair*, 27(2):99–109, 2013.

[54] Verena Klamroth-Marganska, Javier Blanco, Katrin Campen, Armin Curt, Volker Dietz, Thierry Ettlin, Morena Felder, Bernd Fellinghauer, Marco Guidali, Anja Kollmar, et al. Three-dimensional, task-specific robot therapy of the arm after stroke: a multicentre, parallel-group randomised trial. *The Lancet Neurology*, 13(2):159–166, 2014.

[55] Alexander C Koenig and Robert Riener. The human in the loop. In *Neurorehabilitation Technology*, pages 161–181. Springer, 2016.

[56] S Kousidou, NG Tsagarakis, C Smith, and DG Caldwell. Task-orientated biofeedback system for the rehabilitation of the upper limb. In *2007 IEEE 10th International Conference on Rehabilitation Robotics*, pages 376–384. IEEE, 2007.

[57] Hermano Igo Krebs. Twenty+ years of robotics for upper-extremity rehabilitation following a stroke. In *Rehabilitation Robotics*, pages 175–192. Elsevier, 2018.

[58] Hermano Igo Krebs, Jerome Joseph Palazzolo, Laura Dipietro, Mark Ferraro, Jennifer Krol, Keren Rannekleiv, Bruce T Volpe, and Neville Hogan. Rehabilitation robotics: Performance-based progressive robot-assisted therapy. *Autonomous Robots*, 15(1):7–20, 2003.

[59] Hermano Igo Krebs, Bruce T Volpe, Dustin Williams, James Celestino, Steven K Charles, Daniel Lynch, and Neville Hogan. Robot-aided neurorehabilitation: a robot for wrist rehabilitation. *IEEE Trans. Neural Syst. Rehabil. Eng.*, 15(3):327–335, 2007.

[60] Klas Kronander and Aude Billard. Learning compliant manipulation through kinesthetic and tactile human-robot interaction. *IEEE transactions on haptics*, 7(3):367–380, 2013.

[61] Solomon Kullback and Richard A Leibler. On information and sufficiency. *The Annals of Mathematical Statistics*, 22(1):79–86, 1951.

[62] Jean-Paul Laumond. Anthropomorphic action in robotics. In *Brain-inspired intelligent robotics: The intersection of robotics and neuroscience*, pages 39–41. Science/AAAS, 2016.

[63] Jaebong Lee and Seungmoon Choi. Effects of haptic guidance and disturbance on motor learning: Potential advantage of haptic disturbance. In *IEEE Haptics Symposium*, pages 335–342, 2010.

[64] Michael D Lewek, Theresa H Cruz, Jennifer L Moore, Heidi R Roth, Yasin Y Dhaher, and T George Hornby. Allowing intralimb kinematic variability during locomotor training poststroke improves kinematic consistency: a subgroup analysis from a randomized clinical trial. *Physical Therapy*, 89(8):829–839, 2009.

[65] Ming Li and Allison M Okamura. Recognition of operator motions for real-time assistance using virtual fixtures. In *Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems*, pages 125–131. IEEE, 2003.

[66] Yanfang Li, Volkan Patoglu, and Marcia K O'Malley. Negative efficacy of fixed gain error reducing shared control for training in virtual environments. *Transactions on Applied Perception*, 6(1):3, 2009.

[67] Daniel Liberzon. *Switching in Systems and Control*. Springer, 2012.

[68] Keh Chung Lin, Yi An Chen, Chia Ling Chen, Ching Yi Wu, and Ya Fen Chang. The effects of bilateral arm training on motor control and functional performance in chronic stroke: a randomized controlled study. *Neurorehabilitation and Neural Repair*, 24(1):42–51, 2010.

[69] Albert C Lo, Peter D Guarino, Lorie G Richards, Jodie K Haselkorn, George F Wittenberg, Daniel G Federman, Robert J Ringer, Todd H Wagner, Hermano I Krebs, Bruce T Volpe, et al. Robot-assisted therapy for long-term upper-limb impairment after stroke. *New England Journal of Medicine*, 362(19):1772–1783, 2010.

[70] Claudio Loconsole, Stefano Dettori, Antonio Frisoli, Carlo Alberto Avizzano, and Massimo Bergamasco. An emg-based approach for on-line predicted torque control in robotic-assisted rehabilitation. In *2014 IEEE Haptics Symposium (HAPTICS)*, pages 181–186. IEEE, 2014.

[71] Dylan P Losey, Laura H Blumenschein, Janelle P Clark, and Marcia K OMalley. Improving short-term retention after robotic training by leveraging fixed-gain controllers. *Journal of rehabilitation and assistive technologies engineering*, 6:2055668319866311, 2019.

[72] Dylan P Losey and Marcia K O'Malley. Trajectory deformations from physical human–robot interaction. *IEEE Transactions on Robotics*, 34(1):126–138, 2018.

[73] Martin Lotze, Christoph Braun, Niels Birbaumer, Silke Anders, and Leonardo G Cohen. Motor learning elicited by voluntary drive. *Brain*, 126(4):866–872, 2003.

[74] Peter S. Lum, Charles G. Burgar, Peggy C. Shor, Matra Majmundar, and Machiel Van der Loos. Robot-assisted movement training compared with conventional therapy techniques for the rehabilitation of upper-limb motor function after stroke. *Archives of Physical Medicine and Rehabilitation*, 83(7):952–959, 2002.

[75] Ying Mao, Xin Jin, Geetanjali Gera Dutta, John P Scholz, and Sunil K Agrawal. Human movement training with a cable driven arm exoskeleton (carex). *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 23(1):84–92, 2014.

[76] Laura Marchal-Crespo and David J Reinkensmeyer. Haptic guidance can enhance motor learning of a steering task. *J. Motor Behavior*, 40(6):545–557, 2008.

[77] Laura Marchal-Crespo and David J Reinkensmeyer. Review of control strategies for robotic movement training after neurologic injury. *J. NeuroEngineering and Rehabilitation*, 6(1):20–35, 2009.

[78] Laura Marchal-Crespo, Mark van Raai, Georg Rauter, Peter Wolf, and Robert Riener. The effect of haptic guidance and visual feedback on learning a complex tennis task. *Experimental Brain Research*, 231(3):277–291, 2013.

[79] George Mathew and Igor Mezić. Metrics for ergodicity and design of ergodic dynamics for multi-agent systems. *Physica D: Nonlinear Phenomena*, 240(4):432–442, 2011.

[80] Jan Mehrholz, Bernhard Elsner, Cordula Werner, Joachim Kugler, and Marcus Pohl. Electromechanical-assisted training for walking after stroke: updated evidence. *Stroke*, 44(10):e127–e128, 2013.

[81] Lauren M Miller, Yonatan Silverman, Malcolm A MacIver, and Todd D Murphey. Ergodic exploration of distributed information. 32:36–52, 2016.

[82] Marie-Hélène Milot, Laura Marchal-Crespo, Christopher S Green, Steven C Cramer, and David J Reinkensmeyer. Comparison of error-amplification and haptic-guidance training techniques for learning of a timing-based motor task by healthy individuals. *Experimental Brain Research*, 201(2):119–131, 2010.

[83] Alexander Mörtl, Martin Lawitzky, Ayse Kucukyilmaz, Metin Sezgin, Cagatay Basdogan, and Sandra Hirche. The role of roles: Physical cooperation between humans and robots. *Int. J. Robotics Research*, 31(13):1656–1674, 2012.

[84] Salam Moubarak, Minh Tu Pham, Richard Moreau, and Tanneguy Redarce. Gravity compensation of an upper extremity exoskeleton. In *2010 Annual International Conference of the IEEE Engineering in Medicine and Biology*, pages 4489–4493. IEEE, 2010.

[85] Marcia K O'Malley, Abhishek Gupta, Matthew Gen, and Yanfang Li. Shared control in haptic systems for performance enhancement and training. *Journal of Dynamic Systems, Measurement, and Control*, 128(1):75–85, 2006.

[86] Stefano Panzeri, Simon R Schultz, Alessandro Treves, and Edmund T Rolls. Correlations and the encoding of information in the nervous system. *Proc. Royal Society of London B: Biological Sciences*, 266(1423):1001–1012, 1999.

[87] James L Patton, Mark Kovic, and Ferdinando A Mussa-Ivaldi. Custom-designed haptic training for restoring reaching ability to individuals with poststroke hemiparesis. *J. Rehabilitation Research and Development*, 43(5):643–656, 2006.

[88] James L Patton and Ferdinando A Mussa-Ivaldi. Robot-assisted adaptive training: custom force fields for teaching movement patterns. *IEEE Transactions on Biomedical Engineering*, 51(4):636–646, 2004.

[89] James L Patton, Mary Ellen Stoykov, Mark Kovic, and Ferdinando A Mussa-Ivaldi. Evaluation of robotic training forces that either enhance or reduce error in chronic hemiparetic stroke survivors. *Experimental Brain Research*, 168(3):368–383, 2006.

[90] Shahram Payandeh. Application of shared control strategy in the design of a robotic device. In *American Control Conference*, volume 6, pages 4532–4536, 2001.

[91] Ali Utku Pehlivan, Dylan P Losey, and Marcia K O'Malley. Minimal assist-as-needed controller for upper limb robotic rehabilitation. *IEEE Trans. Rob.*, 32(1):113–124, 2016.

[92] Carlos Jesús Pérez-del Pulgar, Jan Smisek, Victor F Munoz, and André Schiele. Using learning from demonstration to generate real-time guidance for haptic shared

control. In *Int. Conf. Systems Man and Cybernetics*, pages 003205–003210. IEEE, 2016.

[93] Rodrigo Pérez-Rodríguez, Carlos Rodríguez, Úrsula Costa, César Cáceres, Josep M Tormos, Josep Medina, and Enrique J Gómez. Anticipatory assistance-as-needed control algorithm for a multijoint upper limb robotic orthosis in physical neurorehabilitation. *Expert Systems with Applications*, 41(8):3922–3934, 2014.

[94] L Pignolo, G Dolce, G Basta, LF Lucca, S Serra, and WG Sannita. Upper limb rehabilitation after stroke: Aramis a robo-mechatronic innovative approach and prototype. In *2012 4th IEEE RAS & EMBS International Conference on Biomedical Robotics and Biomechatronics (BioRob)*, pages 1410–1414. IEEE, 2012.

[95] Dane Powell and Marcia K O'Malley. The task-dependent efficacy of shared-control haptic guidance paradigms. *IEEE Trans. Haptics*, 5(3):208–219, 2012.

[96] GB Prange, MJA Jannink, CGM Groothuis-Oudshoorn, HJ Hermens, and MJ Ijzerman. Systematic review of the effect of robot-aided therapy on recovery of the hemiparetic arm after stroke. *J. Rehabilitation Research and Development*, 43(2):171–184, 2009.

[97] Tommaso Proietti, Vincent Crocher, Agnes Roby-Brami, and Nathanael Jarrasse. Upper-limb robotic exoskeletons for neurorehabilitation: a review on control strategies. *IEEE reviews in biomedical engineering*, 9:4–14, 2016.

[98] Rodrigo Quian Quiroga and Stefano Panzeri. Extracting information from neuronal populations: information theory and decoding approaches. *Nature Reviews Neuroscience*, 10(3):173, 2009.

[99] R Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria, 2016.

[100] Daniel Rakita, Bilge Mutlu, Michael Gleicher, and Laura M Hiatt. Shared dynamic curves: A shared-control telemanipulation method for motor task training. In *Int. Conf. on Human-Robot Interaction*, pages 23–31. ACM, 2018.

[101] Ramya Ramakrishnan, Chongjie Zhang, and Julie Shah. Perturbation training for human-robot teams. *J. Artificial Intelligence Research*, 59:495–541, 2017.

[102] Georg Rauter, Joachim von Zitzewitz, Alexander Duschau-Wicke, Heike Vallery, and Robert Riener. A tendon-based parallel robot applied to motor learning in sports. In *IEEE Int. Conf. on Biomedical Robotics and Biomechatronics (BioRob)*, pages 82–87. IEEE, 2010.

[103] David J Reinkensmeyer and Volker Dietz. Introduction: Rational for machine use. In *Neurorehabilitation Technology*, pages xvii–xxii. Springer, 2016.

[104] David J Reinkensmeyer, Jeremy L Emken, and Steven C Cramer. Robotics, motor learning, and neurologic recovery. *Annu. Rev. Biomed. Eng.*, 6:497–525, 2004.

[105] David J Reinkensmeyer, Eric Wolbrecht, and James Bobrow. A computational model of human-robot load sharing during robot-assisted arm movement training after stroke. In *Int. Conf. of the IEEE Engineering in Medicine and Biology Society*, pages 4019–4023. IEEE, 2007.

[106] David J Reinkensmeyer, Eric T Wolbrecht, Vicky Chan, Cathy Chou, Steven C Cramer, and James E Bobrow. Comparison of 3d, assist-as-needed robotic arm/hand movement training provided with pneu-wrex to conventional table top therapy following chronic stroke. *American Journal of physical medicine & rehabilitation/Association of Academic Physiatrists*, 91(11 0 3):S232, 2012.

[107] Robert Riener, Lars Lunenburger, Saso Jezernik, Martin Anderschitz, Gery Colombo, and Volker Dietz. Patient-cooperative strategies for robot-aided treadmill training: first experimental results. 13(3):380–394, 2005.

[108] Giulio Rosati, Paolo Gallina, and Stefano Masiero. Design, implementation and clinical tests of a wire-based robot for neurorehabilitation. 15(4):560–569, 2007.

[109] Louis B Rosenberg. Virtual fixtures: Perceptual tools for telerobotic manipulation. In *IEEE Virtual Reality Annual International Symposium*, pages 76–82, 1993.

[110] Dorsa Sadigh, Anca D Dragan, Shankar Sastry, and Sanjit A Seshia. Active preference-based learning of reward functions. In *Robotics: Science and Systems*, 2017.

[111] Takeshi Sakurada, Toshihiro Kawase, Kouji Takano, Tomoaki Komatsu, and Kenji Kansaku. A bmi-based occupational therapy assist suit: asynchronous control by ssvep. *Frontiers in neuroscience*, 7:172, 2013.

[112] Richard A Schmidt and Robert A Bjork. New conceptualizations of practice: Common principles in three paradigms suggest new concepts for training. *Psychological Science*, 3(4):207–218, 1992.

[113] Nicolas Schweighofer, Younggeun Choi, Carolee Winstein, and James Gordon. Task-oriented rehabilitation robotics. *American Journal of Physical Medicine & Rehabilitation*, 91(11):S270–S279, 2012.

[114] Sherry E Scott. Different perspectives and formulas for capturing deviation from ergodicity. *Journal on Applied Dynamical Systems*, 12(4):1948–1967, 2013.

[115] Sherry E Scott, Thomas C Redd, Leonid Kuznetsov, Igor Mezić, and Christopher KRT Jones. Capturing deviation from ergodicity at different scales. *Physica D: Nonlinear Phenomena*, 238(16):1668–1679, 2009.

[116] Reza Shadmehr and Ferdinando A Mussa-Ivaldi. Adaptive representation of dynamics during learning of a motor task. *J. Neuroscience*, 14(5):3208–3224, 1994.

[117] Claude E Shannon. A mathematical theory of communication. *The Bell System Technical Journal*, 27(3):379–423, 1948.

[118] Chris R. Sims. Efficient coding explains the universal law of generalization in human perception. *Science*, 360(6389):652–656, 2018.

[119] V Squeri, L Masia, P Giannoni, G Sandini, and P Morasso. Wrist rehabilitation in chronic stroke patients by means of adaptive, progressive robot-aided therapy. *IEEE Trans. Neural Syst. Rehabil. Eng.*, 22(2):312–325, 2014.

[120] Patricia Staubli, Tobias Nef, Verena Klamroth-Marganska, and Robert Riener. Effects of intensive arm training with the rehabilitation robot armin ii in chronic stroke patients: four single-cases. *Journal of neuroengineering and rehabilitation*, 6(1):46, 2009.

[121] Micah Steele and R Brent Gillespie. Shared control between human and machine: Using a haptic steering wheel to aid in land vehicle guidance. In *Proceedings of the human factors and ergonomics society annual meeting*, volume 45, pages 1671–1675, 2001.

[122] Arno HA Stienen, Jacob G McPherson, Alfred C Schouten, and Jules PA Dewald. The ACT-4D: a novel rehabilitation robot for the quantification of upper limb motor impairments following brain injury. In *IEEE Int. Conf. on Rehabilitation Robotics*, pages 1–6, 2011.

[123] Stefanie Tellex, Nakul Gopalan, Hadas Kress-Gazit, and Cynthia Matuszek. Robots that use language. *Annual Review of Control, Robotics, and Autonomous Systems*, 3:25–55, 2020.

[124] Andrea Thomaz, Guy Hoffman, and Maya Cakmak. Computational human-robot interaction. *Foundations and Trends in Robotics*, 4(2-3):105–223, 2016.

[125] Kurt A Thoroughman and Reza Shadmehr. Learning of action through adaptive combination of motor primitives. *Nature*, 407(6805):742, 2000.

[126] Emanuel Todorov. Optimality principles in sensorimotor control. *Nature Neuroscience*, 7(9):907, 2004.

[127] Emanuel Todorov and Michael I Jordan. Optimal feedback control as a theory of motor coordination. *Nature Neuroscience*, 5(11):1226–1235, 2002.

[128] E. Tzorakoleftherakis, M.C. Bengtson, F.A. Mussa-Ivaldi, R.A. Scheidt, and T.D. Murphey. Tactile proprioceptive input in robotic rehabilitation after stroke. In *IEEE Int. Conf. on Robotics and Automation (ICRA)*, pages 6475–6481, 2015.

[129] Emmanouil Tzorakoleftherakis and Todd D Murphey. Controllers as filters: Noise-driven swing-up control based on maxwell's demon. In *IEEE Conf. on Decision and Control (CDC)*, pages 4368–4374, 2015.

[130] Emmanouil Tzorakoleftherakis and Todd D Murphey. Iterative sequential action control for stable, model-based control of nonlinear systems. *IEEE Trans. Automatic Control*, 2018.

[131] Heike Vallery, Edwin HF Van Asseldonk, Martin Buss, and Herman Van Der Kooij. Reference trajectory generation for rehabilitation robots: complementary limb motion estimation. 17(1):23–30, 2009.

[132] Janne Marieke Veerbeek, Erwin van Wegen, Roland van Peppen, Philip Jan van der Wees, Erik Hendriks, Marc Rietberg, and Gert Kwakkel. What is the evidence for physical therapy poststroke? A systematic review and meta-analysis. *PloS one*, 9(2):e87987, 2014.

[133] Massimo Vergassola, Emmanuel Villermaux, and Boris I Shraiman. infotaxis as a strategy for searching without gradients. *Nature*, 445(7126):406, 2007.

[134] Bruce T Volpe, Mark Ferraro, Daniel Lynch, Paul Christos, Jennifer Krol, Christine Trudell, Hermano I Krebs, and Neville Hogan. Robotics and other devices in the treatment of patients recovering from stroke. *Current Neurology and Neuroscience Reports*, 5(6):465–470, 2005.

[135] Joachim von Zitzewitz, Peter Wolf, Vladimir Novaković, Mathias Wellner, Georg Rauter, Andreas Brunschweiler, and Robert Riener. Real-time rowing simulator with multimodal feedback. *Sports Technology*, 1(6):257–266, 2008.

[136] Yorai Wardi and Magnus Egerstedt. Algorithm for optimal mode scheduling in switched systems. In *American Control Conference*, pages 4546–4551, 2012.

[137] Carolee J Winstein and Dorsa Beroukhim Kay. Translating the science into practice: shaping rehabilitation practice to enhance recovery after brain damage. In *Progress in brain research*, volume 218, pages 331–360. Elsevier, 2015.

[138] Carolee J Winstein, Patricia S Pohl, and Rebecca Lewthwaite. Effects of physical guidance and knowledge of results on motor learning: support for the guidance hypothesis. *Research Quarterly for Exercise and Sport*, 65(4):316–323, 1994.

[139] Kimberly J Wisneski and Michelle J Johnson. Quantifying kinematics of purposeful movements to real, imagined, or absent functional objects: implications for modelling trajectories for robot-assisted adl tasks. *J. NeuroEngineering and Rehabilitation*, 4(1):7, 2007.

[140] Eric T Wolbrecht, Vicky Chan, David J Reinkensmeyer, and James E Bobrow. Optimizing compliant, model-based robotic assistance to promote neurorehabilitation. 16(3):286–297, 2008.

[141] Daniel M Wolpert and Zoubin Ghahramani. Computational principles of movement neuroscience. *Nature Neuroscience*, 3:1212–1217, 2000.

[142] Howard G Wu, Yohsuke R Miyamoto, Luis Nicolas Gonzalez Castro, Bence P Ölveczky, and Maurice A Smith. Temporal structure of motor variability is dynamically regulated and predicts motor learning ability. *Nature Neuroscience*, 17(2):312–321, 2014.

[143] Hideyoshi Yanagisawa. A computational model of human perception with prior expectation: Bayesian integration and efficient coding. In *Int. Design Eng. Tech. Conf. and Comp. and Info. in Eng. Conf.* American Society of Mechanical Engineers, 2015.

[144] Wentao Yu, Redwan Alqasemi, Rajiv Dubey, and Norali Pernalete. Telemanipulation assistance based on motion intention recognition. In *IEEE Int. Conf. Robotics and Automation*, pages 1121–1126. IEEE, 2005.

[145] Manfred Zimmermann. The nervous system in the context of information theory. In *Human physiology*, pages 166–173. Springer, 1989.