

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

KATHLEEN FITZSIMONS, Department of Mechanical Engineering, Pennsylvania State University, USA

TODD D MURPHEY, Department of Mechanical Engineering, Northwestern University, USA

Advances in exoskeletons and robot arms have given us increasing opportunities for providing physical support and meaningful feedback in training and rehabilitation settings. However, the chosen control strategies must 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, this leads to low engagement, invariant practice, and few errors, which are not conducive to motor learning. 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 this article, we examine a way to define tasks and close the loop 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.

CCS Concepts: • **Human-centered computing** → **Human computer interaction (HCI)**; *Empirical studies in interaction design*; *Interaction design theory, concepts and paradigms*; Laboratory experiments; Haptic devices; • **Computer systems organization** → **Robotics**.

Additional Key Words and Phrases: human-robot interaction, rehabilitation robotics, shared autonomy

## ACM Reference Format:

Kathleen Fitzsimons and Todd D Murphey. 2022. Ergodic Shared Control: Closing the Loop on pHRI Based on Information Encoded in Motion. 1, 1 (January 2022), 20 pages. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>

## 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[21, 41]. 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

---

Authors’ addresses: Kathleen Fitzsimons, k-fitzsimons@psu.edu, Department of Mechanical Engineering, Pennsylvania State University, University Park, PA, USA; Todd D Murphey, t-murphey@northwestern.edu, Department of Mechanical Engineering, Northwestern University, Evanston, IL, USA.

---

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

© 2022 Association for Computing Machinery.

Manuscript submitted to ACM

Manuscript submitted to ACM

1

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 [20]. Training and therapy is most effective when training is intense and patients are actively involved in the exercises [31]. User engagement can be promoted using game-like graphic interfaces for task-oriented training [48, 55] and by matching the control strategy to the relative difficulty of the task [30, 37]. 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 [5, 34, 44]. 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 [34, 44]. 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 [53].

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. A measure comparing the temporal statistics of trajectories to a spatial distribution defining a task—ergodicity—enables one to identify changes due to reduction of a person’s existing deficit or training with robotic assistance when other common measures failed to capture at least one of these effects [14]. Furthermore, one can easily use a set of demonstrations to learn a task definition with imitation learning [24].

In this paper, 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 1) and empirically compare training with the ergodic HSC to an assist-as-needed path controller.

The paper is organized as follows: First, in Section 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 3. Results



Fig. 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.

of applying the algorithm to simulated noise input are given in Section 4. The experimental protocol and design can be found in Section 5 followed by the results in Section 6. Finally, we provide a discussion of the results in Section 7.

## 2 RELATED WORK

### 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 [6, 8, 38, 51]. This type of guidance is beneficial for healthy participants when the relative difficulty is high and for patients in an acute stage of recovery. As participants learn or relearn motor skills, the task should be performed under shared control [20, 40]. 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 [5]. AAN is usually achieved through impedance/admittance control [44]. 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 or low skill need assistance [30, 37]. 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 [15], 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.

### 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 [4, 56]. In some cases, mean velocity [19, 25], anticipated deviation based on model of participant impairment [42], or other performance heuristics are used. This adaptation has led to moderately better therapeutic outcomes compared to conventional therapy [45]. Yet, the rate at which control parameters are adapted may affect the skill retention [30]. Regardless of the adaptive scheme or the assistive/resistive strategy, some definition of this desired behavior is required.

### 2.3 Reference Definition

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 [6, 26, 51], a healthy person [8, 43], or the unaffected limb in hemiparetic stroke [6, 32], 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 [10, 33] 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 [29], EEG [9, 17, 47], force torque sensors [28, 46], or kinematic data [57]. Proietti et. al suggest that another potential solution is to generate trajectories from statistically consistent patterns from a sufficient number of healthy subjects [44]. 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 paper, we present an approach to reference definition based on statistical task definitions and modulate robotic assistance based on concepts from information theory.

### 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  $X$ , 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) [3], 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 [5]. 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.

### 3 ERGODIC HYBRID SHARED CONTROL

Our previous work has compared error-based hybrid shared control (HSC) to unassisted practice in cart-pendulum inversion task. HSC is hybrid in the sense that it switches between two modes—one where the user has full authority and another where user inputs are actively rejected. User inputs are never replaced, so users must actively contribute to task execution. 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 [15]. 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 [14]. 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.

#### 3.1 Ergodic Metric

Ergodicity can be measured by several metrics [49, 50]. When we analyzed ergodic motion in our prior work, we used the spectral approach [35], 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 [36], 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 [27], where the time-average statistics of the trajectory are defined as a mixture distribution:

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

where  $\eta$  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  $\mathcal{S} = s_1, \dots, s_N$  randomly sampled

over the domain  $\mathcal{X}$ , the KL-divergence ergodic measure [1] is computed by,

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

The KL-divergence ergodic measure scales linearly with the number of randomly sampled points,  $N$ .

### 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 [2, 7, 11, 18, 54]. In hybrid shared control [15], 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,

$$u_2(t) = \begin{cases} u_{user} & t \leq t_0 + t_s \\ u_1 & t_0 + t_s < t \leq t_0 + T \end{cases} \quad (3)$$

We assume a system with dynamics

$$\dot{x}(t) = f(x(t), u(t), t) = g(x(t)) + B(x(t), t)u(t), \quad (4)$$

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

$$J = \varepsilon_{KL} + \int_{t_0}^{t_f} l_2(x(t), u(t)) dt, \quad (5)$$

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,

$$\frac{dJ}{d\lambda} = \cdot \rho(\tau)^T [f(x(\tau), u_2(\tau)) - f(x(\tau), u_1(\tau))]. \quad (6)$$

In (6), state  $x$  is calculated using nominal control,  $u_1$ , and  $\rho$  is the adjoint variable calculated according to Equation 7, below.

$$\begin{aligned} \dot{\rho} &= -l(x, t) - D_x f(x, u_1)^T \rho, \\ &\text{subject to } \rho(t_0 + T) = \mathbf{0} \end{aligned} \quad (7)$$

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 [23] and [15]. User inputs are then accepted when the following integral computed over a time window into the future is negative:

$$\int_{t_0}^{t_0+T} \frac{dJ}{d\lambda}(t) \delta t < 0. \quad (8)$$

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.

### 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, 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 corrective feedback has any effect on performance. 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 1.

---

#### Algorithm 1: Ergodic hybrid shared control

---

```

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


---


2 while  $t_0 < t_f$  do
3   Generate N samples of  $s_i$  within the domain  $\mathcal{X}$ 
4   Infer user input  $u_{user}$  from sensor data
5   Simulate the system (4) for  $t \in [t_0, t_s + T]$ 
6   Simulate (7) for  $t \in [t_0, t_s + T]$ 
7   Compute  $\int_{t_0}^{t_s+T} \frac{dJ}{d\lambda}(t)\delta t$  using Equation 6
8   if  $\int_{t_0}^{t_s+T} \frac{dJ}{d\lambda}(t)\delta t < 0$  then
9     if  $|u_{user}| < u_{sat}$  then
10      Use  $u_{user}$  as current input,  $u_{curr} = u_{user}$ 
11     else
12      Apply saturated user input  $u_{curr} = u_{sat}$ 
13   else
14      Completely “reject”  $u_{user}$  ( $u_{curr} = 0$ )
15   Apply  $u_{curr}$  for  $t \in [t_0, t_0 + t_s]$ 
16    $t_0 = t_0 + t_s$ 


---



```

#### 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  $u_{user} \in [-10m/s^2, 10m/s^2]$ . The cost function defined in Equation 5 was used with a running cost defined as

$$l_2(x(t), u(t)) = u^T 0.0001 \mathbf{1} u.$$

We can see that even using noise as the user input, the hybrid shared control paradigm can produce drawings that resemble the original images (Figure 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 inputs that do not satisfy the inequality in Equation 8. In prior work, we have noted the differences in performance when the robotic interface has relatively low power compared to the user [16] with learning outcomes being more significant when the system has power sufficient to mechanically alter the movements of the user [15].

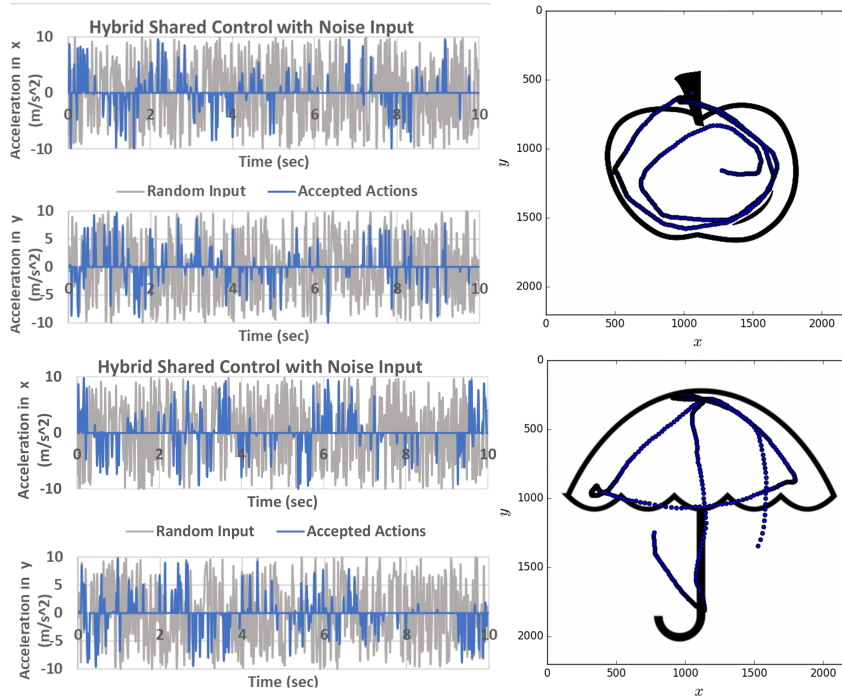


Fig. 2. When noise is used as input, we can use hybrid shared control to reject any noise that does not represent a descent direction for the ergodic cost. This enables us to transform a noisy input (gray) that would generate a random walk to the filtered input (blue). This generates a trajectory that roughly follows the lines of the original image.

#### 5 METHODOLOGY

In this study, we assess ergodic hybrid shared control in a set of timed drawing tasks, and compare it to virtual fixtures. Both hybrid shared control and virtual fixtures require the active participation of the user. Without input from the user,

neither control framework will complete the task on its own. Both control strategies are also designed to assist-as-needed—enabling some variability in the user’s task execution. HSC achieves this by instantaneously evaluating user inputs based on their impact on a cost function in the future. The virtual fixtures are only enabled when participants are far from the task path.

The scale of the drawings in this study requires the participant to reach over a large area of their workspace much like therapeutic reaching exercises in [12, 13, 22, 39, 52]. Drawing has a high level goal similar to daily activities like cleaning or food preparation, where there are many possible strategies one could use to complete the task and the order of locations visited does not necessarily 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 adds a dynamic challenge to the task such that healthy adults need practice to improve their performance. 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.

## 5.1 Study Design/Apparatus

Participants were asked to draw the four grayscale line drawings shown in Figure 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 1—sending updates on the interaction control parameters according to Equation 9 and Equation 10, defined below. The host computer also updated the visualization provided to the users.

*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 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 9,

$$F_{VF} = \frac{K_P(d_p - r)}{d_p} \begin{pmatrix} x_p - x_{user} \\ y_p - y_{user} \end{pmatrix} + K_D \mathbf{v}_{user}, \quad (9)$$

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.

*5.1.2 HSC Rejection Equations.* There are several ways to implement the rejection of user actions described in Algorithm 1. 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-irrelevant 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:

$$\begin{pmatrix} D_x \\ D_y \end{pmatrix} = \begin{pmatrix} \text{sgn}(v_x) & 0 \\ 0 & \text{sgn}(v_y) \end{pmatrix} \begin{pmatrix} \Delta v_x \\ \Delta v_y \end{pmatrix}. \quad (10)$$

## 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 placed 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 or decrease the impedance at the end effector or modify the forces at the end effector according to either Equation 9 or Equation 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.

## 5.3 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) mixed design 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.

**5.3.1 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. The error metric captures the accuracy of the participant’s movement—a common target for robotic training. In tasks such as self-feeding this accuracy is crucial, but for other tasks acceptable trajectories lie within a range as in cleaning a spill. As long as the center of one’s palm travels close enough to the spill, the task can be successfully completed.

**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 3. We use these to calculate the ergodic metric according to Equation 2—giving us the trajectory’s distance from ergodicity for each image. This distance characterizes how far a particular task execution is from the distribution of normative task executions. Therefore trajectories that are close enough to the distribution of demonstrations—represented by the pixel values in the drawings in this experiment—are not penalized for minor variations under this metric.

**5.3.3 Completeness.** Eight students and faculty with limited knowledge of the study design were recruited to rate the completeness of the drawings generated in the study. 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 was 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. Instead, the scorers were asked to rate the completeness of the image based on what percentage of the elements of the original image was completed. Due to randomization of 3000 ratings to 2880 total images, many of the drawings received only one score from the online survey. To mitigate rater bias, drawings were analyzed in sets of 10 for the baseline and post-training sets. In a typical upper limb task like target reaching, one could evaluate whether a task was successful or complete based on whether or not a certain region around the target was visited. In drawing, the completion score provides a continuous measurement of task success.

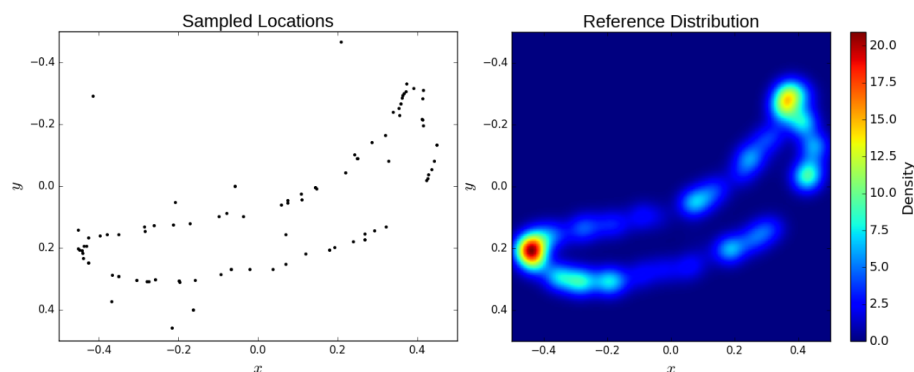


Fig. 3. Using importance sampling, 100 sample locations are randomly chosen from the original image. The color intensity at each point is used to weight a series of Gaussian distributions that generate the reference distribution.



## 6 EXPERIMENTAL RESULTS

The results were reported as follows. First, the error of each group was statistically tested in Section 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 6.2. Finally, an analysis of the completion scores of each group was performed in Section 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.

### 6.1 Error Measure

The mean error of each group in each set can be seen in Figure 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.




Fig. 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 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.

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. This statistical result is unique to the banana drawings. The grayscale area on the interior of this drawing created a narrow gap between the virtual fixtures meant to constrain participants to the outer line and the fixtures meant to constrain them to the inner gray line. Participants would oscillate between the two fixtures because when they reached the midpoint of the region between a black pixel and a gray pixel, the direction of the robotic force would change. The lower error in set 3 may be a result of participants learning to stay in the interior of the drawing where

error will be low because a black or gray pixel are frequently nearby. The interior of this drawing is relatively low area of the distribution used to represent the task, so this low error area was not targeted by the HSC assistance as the gray area would have relatively low impact on reducing the cost in (5).

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 ( $p = 0.728$ ,  $F(1, 20) = 0.124$ ). Participants generally improved as a result of increased practice in this simple line drawing as opposed to the feedback from either of the training paradigms. 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 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.*

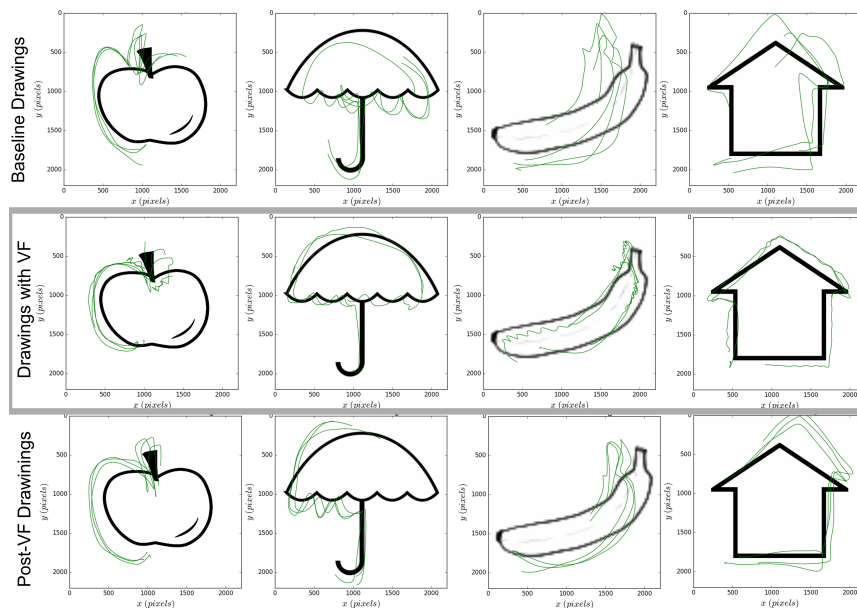


Fig. 5. Example drawings from Participant 10 training with virtual fixtures. Each plot contains the desired image as well as three drawings by the participants. In drawings with virtual fixtures, we see two notable responses to the guidance. Wall following is evident in the drawings of the apple, banana, 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.

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, 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.

## 6.2 Ergodic Measure

Two-factor mixed design ANOVAs were used to assess the effects of the group (between-subjects) and set (within-subjects) on the ergodic measure defined Section 5.3.2 for each image used in the study. The HSC group and VF 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.




Fig. 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 \* $p < 0.05$ , \*\* $p < 0.01$ , \*\*\* $p < 0.001$ .

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 ( $p = 6.17 \times 10^{-4}$ ,  $F(1, 454) = 11.889$ ). 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 ( $p > 0.05$ ). Although the error measure and VF algorithm places equal weight on the black outline and the gray interior line of this drawing, the reference distribution (Figure 1) does not. Therefore, the ergodic measure and the HSC algorithm does not improve when participants fill in these lower density areas.

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.*




Fig. 7. Representative examples of trajectories from Participant 10 using virtual fixture (top) and Participant 6 using hybrid shared control (bottom). The time-average statistics of the trajectories are plotted using the Gaussian approximation from Equation 1. The trajectories in the top row are taken from the middle row of Figure 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.

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 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 7.

### 6.3 Completion Score

The mean completion score of each group in each set can be seen in Figure 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.

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 score 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




Fig. 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 \* $p < 0.05$ , \*\* $p < 0.01$ , \*\*\* $p < 0.001$ .

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 than the VF group post-training.

The analysis of the completion score 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 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.*

## 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.<sup>1</sup> Or, when two desired points are close together on the desired path as seen in the drawings of the banana, the user can oscillate between those points, achieving low error without following the intended paths at all. This emphasizes the fact that error is a limited measure that cannot capture facets of broader 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 samples. The banana can be viewed as an example of how specifications using a single trajectory can over-prescribe a task. The gray line in the middle is a minor detail, but it is weighed equally with other dark points when the task is specified as a path. Ergodic HSC treats the gray line as less important in the context of the overall goal just as it would if one had combined multiple demonstrations to form the task definition and a majority of the demonstrations did not include this detail. The fact that a detail appears in only a few demonstrations in an ensemble implies that it may not necessary to complete the task.

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 characteristics with tasks such as cleaning, personal care, cooking, and reaching. In these activities of daily living (ADLs), completing the task in a way that is 'good enough' can be more valuable than partially completing the task with a highly accurate, precise movement. For instance, when cleaning a spill on your countertop with a sponge, you could use a back and forth scanning motion, but the precise trajectory would not be very important, and you could overreach or underreach the edge of the spill while still removing most of it from the surface. Reference distributions for these types of tasks can be developed with relatively few task demonstrations using the ergodic imitation framework presented in [24]. For activities that require higher accuracy such as self-feeding, a single reference distribution may lead to an ergodic controller that

<sup>1</sup>One could set up virtual fixtures arbitrarily close to the drawing path (i.e.  $r = 0$  in Equation 9)—reducing the error in set 2. However, the motor learning literature suggests that such a restrictive approach to managing error would lead to poor training outcomes [53].

fails to adequately support the task. If the task is segmented, task segments requiring high accuracy like reaching the mouth in self-feeding could be represented with distributions that approach a delta function as more demonstrations are included in the training data. Alternatively, one could specify accuracy-oriented tasks using trajectories; we do not claim that *all* tasks should be represented using distributions. Yet, many tasks can be and some tasks cannot be represented well using trajectories.

The combination of the control approach presented here with ergodic imitation—e.g., combining multiple demonstrations into a single distribution and executing ergodic HSC with respect to that distribution—could enable one to develop a robotic system that can rapidly take demonstrations from a teacher and return an autonomous training assistant that accounts for both the task goal and the natural variability that one should expect in typical human motion.

## ACKNOWLEDGMENTS

This work was supported by the National Science Foundation under grants 1329891 and 1637764 and by the National Defense Science and Engineering Graduate Fellowship program. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation or of the NDSEG program.

## REFERENCES

- [1] Ian Abraham, Ahalya Prabhakar, and Todd D Murphey. 2021. An Ergodic Measure for Active Learning From Equilibrium. *Transactions on Automation Science and Engineering* 18, 3 (2021), 917–931.
- [2] Alexander R Ansari and Todd D Murphey. 2016. Sequential action control: closed-form optimal control for nonlinear and nonsmooth systems. *IEEE Trans. Robot.* 32, 5 (2016), 1196–1214.
- [3] Brenna D Argall, Sonia Chernova, Manuela Veloso, and Brett Browning. 2009. A survey of robot learning from demonstration. *Robotics and Autonomous Systems* 57, 5 (2009), 469–483.
- [4] Sivakumar Balasubramanian, Ruihua Wei, Mike Perez, Ben Shepard, Edward Koeneman, James Koeneman, and Jiping He. 2008. RUPERT: An exoskeleton robot for assisting rehabilitation of arm functions. In *2008 virtual rehabilitation*. IEEE, 163–167.
- [5] Stuart A Bowyer, Brian L Davies, and Ferdinando Rodriguez y Baena. 2013. Active constraints/virtual fixtures: A survey. *IEEE Transactions on Robotics* 30, 1 (2013), 138–157.
- [6] Charles G Burgar, Peter S Lum, Peggy C Shor, and HF Machiel Van der Loos. 2000. Development of robots for rehabilitation therapy: The Palo Alto VA/Stanford experience. *J. Rehabilitation Research and Development* 37, 6 (2000), 663–674.
- [7] Timothy M Caldwell and Todd D Murphey. 2016. Projection-based optimal mode scheduling. *Nonlinear Analysis: Hybrid Systems* 21 (2016), 59–83.
- [8] Gery Colombo, Matthias Joerg, Reinhard Schreier, and Volker Dietz. 2000. Treadmill training of paraplegic patients using a robotic orthosis. *J. Rehabilitation Research and Development* 37, 6 (2000), 693–700.
- [9] Janis J Daly and Jonathan R Wolpaw. 2008. Brain–computer interfaces in neurological rehabilitation. *The Lancet Neurology* 7, 11 (2008), 1032–1043.
- [10] Alexander Duschau-Wicke, Joachim von Zitzewitz, Andrea Caprez, Lars Lunenburger, and Robert Riener. 2010. Path control: a method for patient-cooperative robot-aided gait rehabilitation. *IEEE Trans. Neural Syst. Rehab. Eng.* 18, 1 (2010), 38–48.
- [11] Magnus Egerstedt, Yorai Wardi, and Henrik Axelsson. 2006. Transition-time optimization for switched-mode dynamical systems. *IEEE Trans. Automatic Control* 51, 1 (2006), 110–115.
- [12] Michael D. Ellis, Theresa Sukal-Moulton, and Julius P. A. Dewald. 2009. Progressive Shoulder Abduction Loading is a Crucial Element of Arm Rehabilitation in Chronic Stroke. *Neurorehabilitation and Neural Repair* 23, 8 (2009), 862–869.
- [13] Michael D Ellis, Theresa M Sukal-Moulton, and Julius PA Dewald. 2009. Impairment-based 3-D robotic intervention improves upper extremity work area in chronic stroke: targeting abnormal joint torque coupling with progressive shoulder abduction loading. *IEEE Trans. Robot.* 25, 3 (2009), 549–555.
- [14] Kathleen Fitzsimons, Ana Maria Acosta, Julius P Dewald, and Todd D Murphey. 2019. Ergodicity Reveals Assistance and Learning from Physical Human-Robot Interaction. *Science: Robotics* 4 (2019).
- [15] Kathleen Fitzsimons, Aleksandra Kalinowska, Julius Dewald, and Todd D Murphey. 2020. Task-based hybrid shared control for training through forceful interaction. *Int. J. Robotics Reserach* 39, 9 (2020), 1138–1154.
- [16] Kathleen Fitzsimons, Emmanouil Tzorakoleftherakis, and Todd D Murphey. 2016. Optimal human-in-the-loop interfaces based on Maxwell’s Demon. In *American Control Conference*. 4397–4402.
- [17] Antonio Frisoli, Claudio Loconsole, Daniele Leonardis, Filippo Banno, Michele Barsotti, Carmelo Chisari, and Massimo Bergamasco. 2012. 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 (2012), 1169–1179.
- [18] Humberto Gonzalez, Ram Vasudevan, Maryam Kamgarpour, Shankar S. Sastry, Ruzena Bajcsy, and Claire Tomlin. 2010. A numerical method for the optimal control of switched systems. In *IEEE Conf. on Decision and Control*. 7519–7526.
  - [19] Marco Guidali, Alexander Duschau-Wicke, Simon Broggi, Verena Klamroth-Marganska, Tobias Nef, and Robert Riener. 2011. A robotic system to train activities of daily living in a virtual environment. *Medical & biological engineering & computing* 49, 10 (2011), 1213.
  - [20] Neville Hogan, Hermano I Krebs, Brandon Rohrer, Jerome J Palazzolo, et al. 2006. Motions or muscles? Some behavioral factors underlying robotic assistance of motor recovery. *Journal of Rehabilitation Research and Development* 43, 5 (2006), 605.
  - [21] Siddarth Jain and Brenna Argall. 2019. Probabilistic human intent recognition for shared autonomy in assistive robotics. *ACM Transactions on Human-Robot Interaction (THRI)* 9, 1 (2019), 1–23.
  - [22] Leonard E Kahn, Michele L Zygmant, W Zev Rymer, and David J Reinkensmeyer. 2006. Robot-assisted reaching exercise promotes arm movement recovery in chronic hemiparetic stroke: a randomized controlled pilot study. *J. Neuroengineering and Rehabilitation* 3, 1 (2006), 12.
  - [23] Aleksandra Kalinowska, Kathleen Fitzsimons, Julius P Dewald, and Todd D Murphey. 2018. Online User Assessment for Minimal Intervention During Task-Based Robotic Assistance. In *Robotics: Science and Systems*.
  - [24] Aleksandra Kalinowska, Ahalya Prabhakar, Kathleen Fitzsimons, and Todd D Murphey. 2021. Ergodic imitation: Learning from what to do and what not to do. In *IEEE Int. Conf. on Robotics and Automation*.
  - [25] Verena Klamroth-Marganska, Javier Blanco, Katrin Campen, Armin Curt, Volker Dietz, Thierry Ettl, Morena Felder, Bernd Fellinghauer, Marco Guidali, Anja Kollmar, et al. 2014. Three-dimensional, task-specific robot therapy of the arm after stroke: a multicentre, parallel-group randomised trial. *The Lancet Neurology* 13, 2 (2014), 159–166.
  - [26] S Kousidou, NG Tsagarakis, C Smith, and DG Caldwell. 2007. Task-orientated biofeedback system for the rehabilitation of the upper limb. In *2007 IEEE 10th International Conference on Rehabilitation Robotics*. IEEE, 376–384.
  - [27] Solomon Kullback and Richard A Leibler. 1951. On information and sufficiency. *The Annals of Mathematical Statistics* 22, 1 (1951), 79–86.
  - [28] Ming Li and Allison M Okamura. 2003. Recognition of operator motions for real-time assistance using virtual fixtures. In *Symposium on Haptic Interfaces for Virtual Environment and Teleoperator Systems*. IEEE, 125–131.
  - [29] Claudio Loconsole, Stefano Dettori, Antonio Frisoli, Carlo Alberto Avizzano, and Massimo Bergamasco. 2014. An EMG-based approach for on-line predicted torque control in robotic-assisted rehabilitation. In *2014 IEEE Haptics Symposium (HAPTICS)*. IEEE, 181–186.
  - [30] Dylan P Losey, Laura H Blumenschein, Janelle P Clark, and Marcia K O’Malley. 2019. Improving short-term retention after robotic training by leveraging fixed-gain controllers. *Journal of rehabilitation and assistive technologies engineering* 6 (2019), 2055668319866311.
  - [31] Martin Lotze, Christoph Braun, Niels Birbaumer, Silke Anders, and Leonardo G Cohen. 2003. Motor learning elicited by voluntary drive. *Brain* 126, 4 (2003), 866–872.
  - [32] Peter S. Lum, Charles G. Burgar, Peggy C. Shor, Matra Majmundar, and Machiel Van der Loos. 2002. 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 (2002), 952–959.
  - [33] Ying Mao, Xin Jin, Geetanjali Gera Dutta, John P Scholz, and Sunil K Agrawal. 2014. Human movement training with a cable driven ARm EXoskeleton (CAREX). *IEEE Transactions on Neural Systems and Rehabilitation Engineering* 23, 1 (2014), 84–92.
  - [34] Laura Marchal-Crespo and David J Reinkensmeyer. 2009. Review of control strategies for robotic movement training after neurologic injury. *J. NeuroEngineering and Rehabilitation* 6, 1 (2009), 20–35.
  - [35] George Mathew and Igor Mezić. 2011. Metrics for ergodicity and design of ergodic dynamics for multi-agent systems. *Physica D: Nonlinear Phenomena* 240, 4 (2011), 432–442.
  - [36] Lauren M Miller, Yonatan Silverman, Malcolm A MacIver, and Todd D Murphey. 2016. Ergodic Exploration of Distributed Information. *IEEE Trans. Robot.* 32 (2016), 36–52.
  - [37] Marie-Hélène Milot, Laura Marchal-Crespo, Christopher S Green, Steven C Cramer, and David J Reinkensmeyer. 2010. 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 (2010), 119–131.
  - [38] Salam Moubarak, Minh Tu Pham, Richard Moreau, and Tanneguy Redarce. 2010. Gravity compensation of an upper extremity exoskeleton. In *2010 Annual International Conference of the IEEE Engineering in Medicine and Biology*. IEEE, 4489–4493.
  - [39] James L Patton, Mark Kovic, and Ferdinando A Mussa-Ivaldi. 2006. Custom-designed haptic training for restoring reaching ability to individuals with poststroke hemiparesis. *J. Rehabilitation Research and Development* 43, 5 (2006), 643–656.
  - [40] James L Patton and Ferdinando A Mussa-Ivaldi. 2004. Robot-assisted adaptive training: custom force fields for teaching movement patterns. *IEEE Transactions on Biomedical Engineering* 51, 4 (2004), 636–646.
  - [41] Shahram Payandeh. 2001. Application of shared control strategy in the design of a robotic device. In *American Control Conference*, Vol. 6. 4532–4536.
  - [42] Rodrigo Pérez-Rodríguez, Carlos Rodríguez, Úrsula Costa, César Cáceres, Josep M Tormos, Josep Medina, and Enrique J Gómez. 2014. Anticipatory assistance-as-needed control algorithm for a multijoint upper limb robotic orthosis in physical neurorehabilitation. *Expert Systems with Applications* 41, 8 (2014), 3922–3934.
  - [43] L Pignolo, G Dolce, G Basta, LF Lucca, S Serra, and WG Sannita. 2012. 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 Biomechanics (BioRob)*. IEEE, 1410–1414.

- [44] Tommaso Proietti, Vincent Crocher, Agnes Roby-Brami, and Nathanael Jarrasse. 2016. Upper-limb robotic exoskeletons for neurorehabilitation: a review on control strategies. *IEEE reviews in biomedical engineering* 9 (2016), 4–14.
- [45] David J Reinkensmeyer, Eric T Wolbrecht, Vicky Chan, Cathy Chou, Steven C Cramer, and James E Bobrow. 2012. 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 (2012), S232.
- [46] Robert Riener, Lars Lunenburger, Saso Jezernik, Martin Anderschitz, Gery Colombo, and Volker Dietz. 2005. Patient-cooperative strategies for robot-aided treadmill training: first experimental results. *IEEE Trans. Neural Syst. Rehab. Eng.* 13, 3 (2005), 380–394.
- [47] Takeshi Sakurada, Toshihiro Kawase, Kouji Takano, Tomoaki Komatsu, and Kenji Kansaku. 2013. A BMI-based occupational therapy assist suit: asynchronous control by SSVEP. *Frontiers in neuroscience* 7 (2013), 172.
- [48] Nicolas Schweighofer, Younggeun Choi, Carolee Winstein, and James Gordon. 2012. Task-oriented rehabilitation robotics. *American Journal of Physical Medicine & Rehabilitation* 91, 11 (2012), S270–S279.
- [49] Sherry E Scott. 2013. Different perspectives and formulas for capturing deviation from ergodicity. *Journal on Applied Dynamical Systems* 12, 4 (2013), 1948–1967.
- [50] Sherry E Scott, Thomas C Redd, Leonid Kuznetsov, Igor Mezić, and Christopher KRT Jones. 2009. Capturing deviation from ergodicity at different scales. *Physica D: Nonlinear Phenomena* 238, 16 (2009), 1668–1679.
- [51] Patricia Staubli, Tobias Nef, Verena Klamroth-Marganska, and Robert Riener. 2009. 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 (2009), 46.
- [52] Theresa M Sukal, Michael D Ellis, and Julius PA Dewald. 2007. Shoulder abduction-induced reductions in reaching work area following hemiparetic stroke: neuroscientific implications. *Experimental Brain Research* 183, 2 (2007), 215–223.
- [53] Kurt A Thoroughman and Reza Shadmehr. 2000. Learning of action through adaptive combination of motor primitives. *Nature* 407, 6805 (2000), 742.
- [54] Yorai Wardi and Magnus Egerstedt. 2012. Algorithm for optimal mode scheduling in switched systems. In *American Control Conference*. 4546–4551.
- [55] Carolee J Winstein and Dorsa Beroukhi Kay. 2015. Translating the science into practice: shaping rehabilitation practice to enhance recovery after brain damage. In *Progress in brain research*. Vol. 218. Elsevier, 331–360.
- [56] Eric T Wolbrecht, Vicky Chan, David J Reinkensmeyer, and James E Bobrow. 2008. Optimizing compliant, model-based robotic assistance to promote neurorehabilitation. *IEEE Trans. Neural Syst. Rehab. Eng.* 16, 3 (2008), 286–297.
- [57] Wentao Yu, Redwan Alqasemi, Rajiv Dubey, and Norali Pernalet. 2005. Telemanipulation assistance based on motion intention recognition. In *IEEE Int. Conf. Robotics and Automation*. IEEE, 1121–1126.