Learning to Control Complex Robots Using High-Dimensional Body-Machine Interfaces

JONGMIN M. LEE, Northwestern University, USA and Shirley Ryan AbilityLab, USA
TEMESGEN GEBREKRISTOS, Northwestern University, USA and Shirley Ryan AbilityLab, USA
DALIA DE SANTIS, Shirley Ryan AbilityLab, USA
MAHDIEH NEJATI-JAVAREMI, Northwestern University, USA and Shirley Ryan AbilityLab, USA
DEEPAK GOPINATH, Northwestern University, USA and Shirley Ryan AbilityLab, USA
BIRAJ PARIKH, Northwestern University, USA and Shirley Ryan AbilityLab, USA
FERDINANDO A. MUSSA-IVALDI, Northwestern University, USA and Shirley Ryan AbilityLab, USA
BRENNA D. ARGALL, Northwestern University, USA and Shirley Ryan AbilityLab, USA

When individuals are paralyzed from injury or damage to the brain, upper body movement and function can be compromised. While the use of body motions to interface with machines has shown to be an effective noninvasive strategy to provide movement assistance and to promote physical rehabilitation, learning to use such interfaces to control complex machines is not well understood. In a five session study, we demonstrate that a subset of an uninjured population is able to learn and improve their ability to use a high-dimensional Body-Machine Interface (BoMI), to control a robotic arm. We use a sensor net of four inertial measurement units, placed bilaterally on the upper body, and a BoMI with the capacity to directly control a robot in six dimensions. We consider whether the way in which the robot control space is mapped from human inputs has any impact on learning. Our results suggest that the space of robot control does play a role in the evolution of human learning: specifically, though robot control in joint space appears to be more intuitive initially, control in task space is found to have a greater capacity for longer-term improvement and learning. Our results further suggest that there is an inverse relationship between control dimension couplings and task performance.

CCS Concepts: • **Computer systems organization → External interfaces for robotics**; **Robotic control**; • **Human-centered computing → User studies**.

Additional Key Words and Phrases: body-machine interface, assistive manipulator, motor learning

ACM Reference Format:
Jongmin M. Lee, Temesgen Gebrekristos, Dalia De Santis, Mahdieh Nejati-Javaremi, Deepak Gopinath, Biraj Parikh, Ferdinando A. Mussa-Ivaldi, and Brenna D. Argall. 2022. Learning to Control Complex Robots Using High-Dimensional Body-Machine Interfaces. *J. ACM* 37, 4, Article 111 (August 2022), 20 pages. https://doi.org/XXXXXXX.XXXXXXX

Authors' addresses: Jongmin M. Lee, jmlee@u.northwestern.edu, Northwestern University, 2145 Sheridan Rd, Evanston, Illinois, USA, 60208 and Shirley Ryan AbilityLab, 355 E Erie St, Chicago, Illinois, USA, 60611; Temesgen Gebrekristos, Northwestern University, USA and Shirley Ryan AbilityLab, USA, tem@u.northwestern.edu; Dalia De Santis, Shirley Ryan AbilityLab, USA, ddesantis@sralab.org; Mahdieh Nejati-Javaremi, Northwestern University, USA and Shirley Ryan AbilityLab, USA, m.nejati@u.northwestern.edu; Deepak Gopinath, Northwestern University, USA and Shirley Ryan AbilityLab, USA, deepakgopinath@u.northwestern.edu; Biraj Parikh, Northwestern University, USA and Shirley Ryan AbilityLab, USA, bparikh@u.northwestern.edu; Ferdinando A. Mussa-Ivaldi, Northwestern University, USA and Shirley Ryan AbilityLab, USA, sandro@northwestern.edu; Brenna D. Argall, Northwestern University, USA and Shirley Ryan AbilityLab, USA, brenna.argall@northwestern.edu.

0004-5411/2022/8-ART111 $15.00
https://doi.org/XXXXXXX.XXXXXXX

1 INTRODUCTION

The use of body motions to control machines has seen a growing interest in the neuroengineering, rehabilitation science, and robotics communities, due to their noninvasiveness, ability to adjust to an individual's available body movements, and support for lost motor function [1, 8, 18, 37, 41, 45]. In comparison to alternative, noninvasive human-machine interfaces such as surface electroencephalogram (EEG), body motions promote physical activity, with associated potential for therapeutic benefits [7, 36–38]. Body-machine interfaces (BoMIs) incentivize patients with paralysis and movement disorders to use their remaining residual mobility [6], and, with practice, they have been shown to increase patients' upper body muscle strength and upper body mobility [38]. BoMIs use motion sensor technologies to measure features of movement from the surface of the body [8]. Unlike many interfaces that are typically used to control assistive machines, BoMIs have the capacity to generate control signal inputs in high dimensions from residual body movements—even when users suffer from severe levels of paralysis [38, 39].

Despite the recent excitement in using BoMIs to control machines, demonstrations on assistive robotic systems have been predominately limited to control of low degrees-of-freedom (DoFs) systems [18, 54]. High-DoF control, however, is often needed to perform high-resolution dexterous movements in our physical world—and to efficiently achieve Activities of Daily Living (ADLs) and other complex manipulation tasks, with intention, in a timely manner [5]. While using high-dimensional interfaces (e.g., BoMI) to control high-DoF robots (e.g., robotic arms) holds the potential for positive assistive and rehabilitation outcomes [38, 41], it is unclear to what extent people are able to learn to recoordinate their body movements to issue high-dimensional control signals—*necessary* to simultaneously control all of the DoFs of a complex robot. For example, a result from a past study suggests that simply expanding control from two to three control dimensions introduces significant learning challenges to participants [42]. In addition, given that the variation of body motion data has shown to be low [40], the expectation is that the learning challenges will also dramatically increase as the need for higher-DoF control becomes greater. Taken together, these raise the questions of whether body movements can be executed with consistency and sufficient dexterity to reliably complete functional tasks, and whether BoMI control might become unintuitive or more challenging to learn as control complexity in the robot increases.

In this work, we perform a multi-session study of high-DoF robotic arm teleoperation via a BoMI as a first step to address these questions. We investigate in particular the control space mapping—that is, the selection of which robot control space (joint or task) to map the interface commands—to assess the mapping's impact on human learning and whether one space might be more intuitive to operate. While there are many indications that humans primarily plan their own hand and arm movements within the task space, it is also the case that a common characteristic found in patients with cerebellar damage is a decomposition of complex movements to individual joint movements [3]. It is possible that the novelty of the BoMI could introduce a similar simplification of the robot manipulation problem into decomposed robot movements, where controlling individual joint movements—instead of movements within the task space—is more intuitive for users.

We highlight the following contributions, first presented in our conference paper [21]:

- A multi-session study of high-DoF robotic arm teleoperation via a BoMI.
- A demonstration that body-machine interfacing is scalable to higher-DoF robots.
- An analysis of the impact of control space mappings on task performance, workload, and human learning.

In this article, we furthermore present the following new contributions:

- An analysis of the relationship between control dimension couplings and task performance.

- An evolved analysis of the impact of control space mappings on perceived workload and user subjectivity.

A review of background and related literature is provided in Section 2. Section 3 describes our experimental methods, with results shared in Section 4. A discussion of key insights and conclusions then are provided respectively in Sections 5 and 6.

2 BACKGROUND AND RELATED WORK

In this section we review related literature on assistive machines and interfaces, with a focus on BoMIs. We also provide a brief review of movement learning and functional rehabilitation.

2.1 Assistive Interfaces and Machines

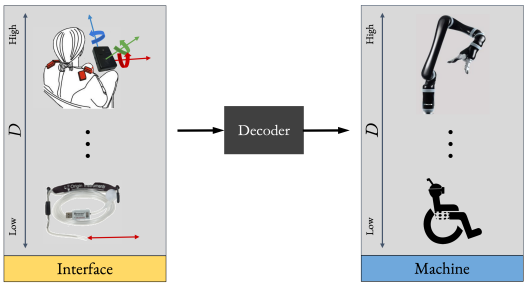

Fig. 1. Dimensionality of interface inputs and machine outputs. Interfaces might capture signals that are lower or higher dimensional than the assistive machines they intend to operate. In the latter case, input signals typically are mapped through a decoder to the control space of the machine.

Both assistive machines and interfaces have design limitations and vary in complexity or number of controllable DoFs. Assistive machines can range in control complexity from simple machines, such as power wheelchairs, to more complex machines, such as robotic arms. In general, the level of complexity varies with the number of DoFs the machine can operate. High-DoF assistive machines, such as robotic arms, can be difficult to control with commercially-available assistive interfaces (e.g., sip-and-puff, head array, and 2D joystick), which overwhelmingly are designed to control the much lower-dimensional (2D) power wheelchair. This is due to a mismatch in the number of input dimensions the operator has access to on the interface, compared to the number of output control dimensions required to operate the high-DoF machine (Figure 1) [43].

A common approach to this mismatch is modal control [17, 25, 26, 48], in which only a subset of the control dimensions of the robot are operated at a given time. While modal control does facilitate access to the full control space, it does not allow users to access all control dimensions simultaneously. An alternative strategy to tackling dimensionality mismatch is to focus on control dimensions that are only task relevant, by embedding them within a lower-dimensional space and defining latent actions to operate within that space [23]. While this is a clever approach shown to be effective with 2D joystick teleoperation, the approach requires prior knowledge of the tasks and goals and, by design, it limits the full control authority on the human operator. The *simultaneous* and *continuous* control of all translation and orientation dimensions in complex robots remains an open and challenging research area.

Intracortical Brain-Machine Interfaces (BMIs) use implants to record brain activity from the motor cortex and offer the possibility of directly issuing high-dimensional control signals to overcome the problem of dimensionality mismatch [9, 12, 33]. Intracortical BMIs have enormous potential to help

people with neurological disorders and injuries; however, they can be extremely invasive—requiring surgery to the brain and the permanent implantation of an electrode array [10]. Although BMIs can be noninvasive using surface EEG, there is less evidence that suggests the signal-to-noise ratio is sufficient for the continuous and simultaneous control of robots in high dimensions [20, 28].

2.2 Functional Mobility and Rehabilitation

People who suffer from upper and/or lower body paralysis, such as from spinal cord injury (SCI), experience loss of functional independence. Cascading effects can further lead to a disruption in quality of life (e.g., respiratory, cardiovascular; depression, substance abuse, emotional distress), leading to increased dependence on caregivers [e.g., 34]. The severity of paralysis plays a critical role in how much functionality is compromised. Even in tetraplegia, residual movements can still remain intact [e.g., 29]. Movement training and therapy designed to encourage residual movements and increase mobility has been shown to promote neuroplastic changes and improve functionality [32, 50]. Similar outcomes of enhanced muscle strength and mobility in the upper body have also been shown in studies using a BoMI for low-DoF control [37].

2.3 Body-Machine Interfacing

BoMIs are able to capture the residual movement available to patients with paralysis. They are responsive to human-to-human variability (e.g., between or within levels of SCI) [8], enhance muscle strength and mobility [38], and achieve functional rehabilitation aims [2, 38, 41]. By casting a net of sensors on the body, the BoMI captures body movements, and it can be customized to individuals through their unique availability of body movements—for example, by tuning the BoMI map's parameters (e.g., gains, offsets) session-by-session [38].

In general, the choice of sensor technology to use for body-machine interfacing is vast as long as it is capable of characterizing body kinematics. Historically, research groups have used infrared cameras [7], inertial measurement units (IMUs) [14], EMG sensors [47], or an ensemble of sensor technologies [45] to collect measurements. Similarly, body movement can be interfaced with a variety of machines such as powered wheelchairs [54], robotic arms [11, 18], and drones [24, 30].

What allows all of these hardware components to be turned into an interface between human and robot is the choice of which sensor measurements to record and how to map them to control commands for the machine can interpret. Given a set of body movements and a given machine we want to control, we need an algorithm that *decodes* a continuous body movement signal $\mathbf{x}$ to use as robot control signals $\mathbf{q}$. A classic approach to design the decoder is to describe a linear relationship $\mathbf{A}$ between $\mathbf{x}$ and $\mathbf{q}$.

More formally, given a set of sensor measurements $\mathbf{x} = \mathbf{x}_1, \ldots, \mathbf{x}_n$, we can define the feedforward linear mapping between measurements $\mathbf{x}$ and robot control commands $\mathbf{q}$ to be

$$\mathbf{q} = \mathbf{A}\mathbf{x} + \mathbf{b}, \tag{1}$$

where $\mathbf{q} = [q_1, \ldots, q_k]^T$ is a $k \times 1$ vector, $\mathbf{A}$ is a $k \times n$ matrix that maps $\mathbf{x} \rightarrow \mathbf{q}$, $\mathbf{x}$ is a $n \times 1$ vector, and $\mathbf{b}$ is a $k \times 1$ vector that is an affine offset. The shape of vector $\mathbf{q}$ is determined by the product of the number of sensors and the dimensionality of a single sensor output. For example, if we are interested in sensor measurements of orientation, represented by quaternions, and we wish to use four sensors, then $n = 4 \times 4 = 16$. The output $\mathbf{q}$ depends on the space in which robot control happens, where $k = 6$ for both 6-DoF task-space (or Cartesian-space) control ($\mathbf{q}_{TS} = [q_x, q_y, q_z, q_\theta, q_\phi, q_\psi]^T$) and 6-DoF joint-space control ($\mathbf{q}_{JS} = [q_{j1}, q_{j2}, q_{j3}, q_{j4}, q_{j5}, q_{j6}]^T$). We should note that in prior BoMI studies, $\mathbf{q}$ has predominantly been used with position- and velocity-based controllers.

Several past works use Principal Component Analysis (PCA) [35] as the approach to solve for $\mathbf{A}$ (and $\mathbf{b}$) [18, 37, 54]. More recent examples use iterative linear methods [13] and deep learning

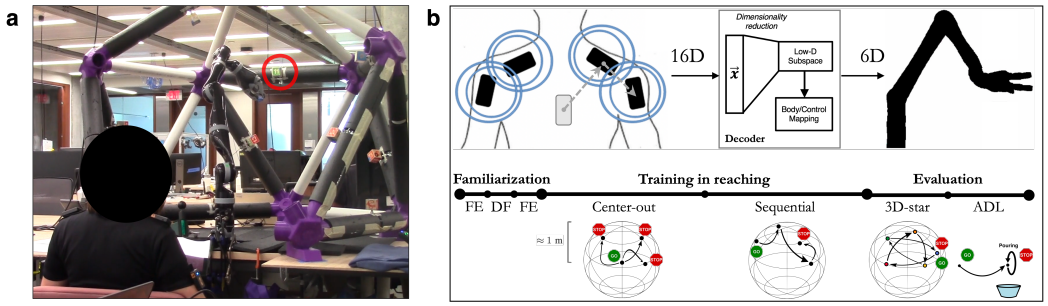

Fig. 2. An overview of the interface-robot pipeline and the study tasks. (a) Participant wearing the BoMI interacting with the JACO robotic arm and reaching targets. (b) A sensor net on the upper body controls the robotic arm to perform reaching and functional tasks. The relative quaternion orientation of the IMUs (4 IMUs × 4 quaternions = 16D combined) is mapped to a 6-dimensional linear subspace for continuous and simultaneous control of the robot, whether in task-space (3D translation + 3D orientation) or joint-space (6D joint angles). The progression of a study session tasks consists of: (1) free exploration (FE); (2) all-but-one DoF freezing (DF); (3) FE; (4) center-out reaching; (5) sequential reaching; (6) sequential reaching in a 3D-star shape; and (7) ADL-inspired tasks.

methods such as adaptive nonlinear autoencoders that relax the constraint of a linear relationship between $\mathbf{x}$ and $\mathbf{q}$ [44]. Due to our particular interest in human learning of a complex (high-DoF) system over multiple study sessions, we choose the more familiar approach of PCA, that has a rich history of success in both BoMI studies and engineering more broadly.

3 METHODS

In this section we present the experimental details of our multi-session study of robotic arm operation using a high-dimensional BoMI.

3.1 Participants

Ten uninjured adults (28 ±8 median age years; 6 males, 4 females) participated in this study. Each participant completed five sessions of approximately two hours each. Participant assignment to one of two groups was random and balanced. The joint-space (JS) group directly controlled the velocity of the robot's joints, and the task-space (TS) group controlled the velocity of the robot end-effector in translation and orientation. Participants in each group received the same ordering of tasks, number of trials, and number of sessions. All study sessions were conducted with the approval of the Institutional Review Board at Northwestern University. All participants provided their written informed consent.

3.2 Materials

Interface. A sensor net of four IMU sensors (Yost Labs, Ohio, USA) are placed bilaterally on the shoulders and upper arms and anchored to a custom shirt designed to minimize movement artifacts. Sensor placement is predetermined based on past BoMI studies. To maintain consistency between participants, we add a reference chest sensor (marked in Figure 2b, top) from which its orientation data is used to compute relative orientations through a predetermined kinematic chain (chest → shoulders → upper arms). The chest sensor data is not directly used in the decoder design, therefore, it is excluded in the online projection of the real-time data onto the control signal subspace.

A Kalman filter is applied to the IMU data and computed onboard the IMU sensors, producing real-time, orientation estimates through a fusion of accelerometer and gyroscope measurements.

The pipeline is visually represented in Figure 2b (top). The relative quaternion orientations of the four IMUs in the net (16D) are mapped to a six-dimensional (6D) subspace using PCA. To build the PCA map, we adopt a semi-supervised approach, in which we instruct an experienced user to perform a predefined, repeated set of upper-body movements (shoulders forward/backward and up/down; elbows in/out) for approximately 60 seconds each. These collected datasets, each composed of single-joint movements, are then compressed, transformed, and mapped, all via PCA, into six principal components (PCs) to produce a 6D map. To simplify the map and more evenly distribute the explained variance across the PCs, varimax rotation is applied *post hoc* [19]. The resulting BoMI map allows for a linear mapping between sensor measurements ($\mathbf{x}$) and a lower-dimensional robot control signal ($\mathbf{q}$) as: $\mathbf{q} = \mathbf{A}\mathbf{x} + \mathbf{b}$. Together, using this approach allows for a continuous output, in terms of velocity commands, to the robot—as well as a better representation of the individual movements available, with the goal of providing more intuitive and independent control of the robot's movements.

Robot Control. The lower-dimensional subspace is used online to control a 7-DoF JACO robotic arm (Kinova Robotics, Quebec, Canada). We exclude the control of the fifth joint—which is the redundant joint in the kinematic chain of the 7-DoF JACO robotic arm—from the study. We hold this joint fixed so that both the TS and JS groups are operating the same number (6) of DoFs and under the same control constraints.

The PCs of the lower-dimensional PCA subspace are mapped to the robot control space as follows. For the TS mapping, we prioritize translation control over orientation control by mapping the first three PCs to x, y, z (which by definition capture more of the body movement variance) and the next three PCs to θ (roll), ϕ (pitch), ψ (yaw). For the JS mapping, we map the PCs in the order of joints in the kinematic chain of the robotic arm. This prioritization is primarily based on prior BoMI and BMI studies, as well as study task design.

To avoid involuntary robot commands and to compensate for sensor noise during study trials, we define a control threshold linearly proportional to the applied gains, shifted by a constant offset. Though the thresholding equation remains fixed for all participants, control gains and offsets are customized daily to the individual using observation-based tuning, such that the variances across control dimensions are approximately equivalent and each control dimension is centered along its respective mean. This is to customize and maximize the utility of the map for each individual. Control signals to the robot are published at an approximate rate of 10 Hz.

In order to reduce complexity of adding grasp to the control problem, we disable control of the gripper and focus on reaching tasks.

Visual Feedback. A graphical-user interface (GUI) is displayed on a tablet to provide a real-time visualization of the robot velocity commands to the participant. Control commands are represented either in Cartesian velocities or joint velocities, depending on the group assignment.

In addition, a scoring system is displayed on the GUI to increase participant engagement and to provide trial-by-trial feedback on performance to the participant. Trial scores are calculated based on robot endpoint distance-to-target (in terms of both translation and orientation), where the orientation score is credited only when the end-effector's distance is within 50% of the reach distance. Though participants are informed that their scores would be displayed on the tablet, the aforementioned implementation details are excluded from the study instructions.

For reaching tasks, ten target blocks (marked in Figure 2a) are used as goals, and placements are determined to maximize robot workspace coverage and diversify reaching movements. The goal

locations remain fixed throughout the study for all participants. Targets are affixed to a custom-built cage constructed with PVC pipes and 3D-printed joints, as shown in Figure 2a.

3.3 Study Protocol

There are three phases to the study protocol: (a) familiarization, (b) training, and (c) evaluation (shown in Figure 2b).

Familiarization. In the first phase, participants are encouraged to explore and become familiar with the system, with minimal constraints enforced in the free exploration (FE) task. In the all-but-one DoF freezing (DF) task, participants are iteratively introduced to each control dimension one at a time, where all other dimensions are kept frozen.

Training. In the second phase, two categories of reaching tasks are employed: (1) reaches from a fixed center position out to a target and (2) sequential reaches between targets. The order of targets is randomized and balanced across days to avoid ordering effects, and this order is preserved across participants. Each trial consists of a single reach either from the center-to-target (center-out reach) or from target-to-target (sequential reach). In sequential reach trials specifically, we maintain consistency between trials and participants by positioning, at the start of each reach, the robot at predefined waypoints, in the proximity of the previous target. These two categories of reaching tasks are modeled after a standard experimental paradigm in motor learning [51] as well as movements that are more typical in real world use.

Evaluation. In the third phase, tasks are split into evaluations of reaching and ADLs. In the reaching evaluation task, participants reach to five targets that comprise a three-dimensional star in fixed succession (shown pictorially in Figure 2b, bottom). The ADL evaluation tasks are designed to emulate four ADL tasks: (a) take a cup (upside-down) from a dish rack and place it (upright) on the table, (b) pour cereal into a bowl, (c) scoop cereal from a bowl, and (d) throw away a surgical mask in the trash bin. To facilitate learning, our training protocol adopts ideas from the literature in repetition learning [49] and learning via contextual interference [4, 52]. In addition, to maintain engagement and motivation [22], we display a trial-by-trial scoring system based on performance.

A trial ends upon successful completion or timeout. For reaching any target, success is defined within strict positional (1.00 cm) and rotational (0.02 rad, or 1.14°) thresholds, and the timeout is 90 seconds. For the ADL tasks, experimenters follow codified guidelines to determine when tasks are completed with a task timeout of 180 seconds. Participants are informed of the timeouts and asked to perform tasks to the best of their ability. If there is any risk of harm to the participant or the robot, study personnel are instructed to intervene and teleoperate the robot to a safe position, as close to the original position the robot left the experiment workspace, before proceeding.

Sensors are calibrated daily and tared between trials to account for the possibility of intrinsic sensor noise and drift.

Over the course of the study, data from a total of 400 center-out, 400 sequential, and 250 3D-star reaching trials are gathered, as well as from 80 ADL task trials.

3.4 Metrics

In order to obtain a comprehensive understanding of human learning in the context of controlling a novel interface-robot system, we analyze subjective responses from questionnaires, use objective performance metrics appropriate for the specific task, and quantify control dimension couplings to capture a participant's ability to isolate robot movements.

Questionnaires. Questionnaires are provided to participants at the end of each study session. Two types of questionnaires are used. The first is the NASA task load index (NASA-TLX), which is an

assessment tool for subjective workload in human-machine interfacing contexts [16]. The second is a study-specific questionnaire consisting of questions relating to (a) task difficulty, (b) human-robot interaction, and (c) learning effects. The study-specific questions are further divided into Likert-scale questions and free form questions.

Performance Metrics. We use performance metrics specifically tailored to the study's reaching tasks and ADL tasks, defined as follows:

- Success rate

$$\mu_S = \frac{1}{N}\sum_{i=1}^{N} \mathbf{1}_{\{S\}_i} \tag{2}$$

where

$$\mathbf{1}_{\{S\}} = \begin{cases} 1, & \text{task success} \\ 0, & \text{otherwise.} \end{cases}$$

N is the total number of trials, and $\mathbf{1}_{\{S\}}$ is an indicator function for task success.

- Successful completion time

$$t_c = \mathbf{1}_{\{S\}}(t_{end} - t_{start}) \tag{3}$$

where t_{start} and t_{end} are the respective start and end times of a given trial.

- Average number of collisions

$$\mu_{collision} = \frac{1}{N}\sum_{i=1}^{N} \mathbf{1}_{\{C\}_i} \tag{4}$$

where $\mathbf{1}_{\{C\}}$ is an indicator function for a collision during a trial. A collision is marked when any part of the robot comes in contact with the physical environment.

- Normalized path length

$$\frac{\ell_{path}}{\ell_{straight}} = \frac{\sum_{i=0}^{M-1} \|\mathbf{x}_i - \mathbf{x}_{i+1}\|}{\|\mathbf{x}_{target} - \mathbf{x}_{start}\|} \tag{5}$$

where $\|\cdot\|$ is the L2 norm, and ℓ_{path} and $\ell_{straight}$ are the trial's end-effector path length and straight-line path length, respectively, to a target. M is the total number of samples (sampled at 10 Hz), $\mathbf{x}_i$ is the end-effector pose at the i^{th} sample, and $\mathbf{x}_{start} = \mathbf{x}_{i=0}$ and is the starting pose.

- Average proportion of time spent within k percent of reach distance

$$\mu_{\tau_{dist\leq k}} = \frac{1}{N}\sum_{i=1}^{N} \tau_{dist\leq k,i}, \tag{6}$$

$$\tau_{dist\leq k} = \frac{t_{dist\leq k}}{t_{end} - t_{start}},$$

where $t_{dist\leq k}$ is the time spent within k percent of the reach distance on a given trial and $0 \leq k \leq 100\%$. We furthermore also compute $\tau_{dist>100\%}$ as the average proportion of time spent beyond 100% of the reach distance, as a metric for negative progress towards the target.

Control Dimension Couplings. We use correlation analysis to compute the degree of linear coupling between pairs of control dimensions, defined as follows:

- Matrix of correlation coefficients

$$\boldsymbol{\rho}(\mathbf{q},\mathbf{q}) = \frac{\mathrm{cov}(\mathbf{q},\mathbf{q})}{\sqrt{\mathrm{var}(\mathbf{q}),\mathrm{var}(\mathbf{q})}} \tag{7}$$

  where $\mathrm{cov}(\cdot)$ and $\mathrm{var}(\cdot)$ are covariance and variance, respectively. Each element of the correlation coefficient matrix ranges between $[-1,+1]$, where $\rho_{i,j} = [-1,0)$ implies a negative (or inverse) linear correlation between control dimensions i and j, $\rho_{i,j} = (0,+1]$ implies a positive correlation, and $\rho_{i,j} = 0$ implies no correlation.
- Percentage of significant correlations

$$\rho^* = \left(\frac{2}{n(n-1)}\sum_{i=1}^{n}\sum_{j=i+1}^{n}\mathbf{1}_{\{\rho\}_{i,j}}\right)\cdot 100, \tag{8}$$

  where

$$\mathbf{1}_{\{\rho\}_{i,j}} = \begin{cases} 1, & \text{if } p < 0.05 \text{ for } \rho_{i,j} \text{ and } i < j \\ 0, & \text{otherwise.} \end{cases}$$

  Here, $\mathbf{1}_{\{\rho\}_{i,j}}$ is an indicator function for a coupling with statistical significance ($p < 0.05$) [53], and $n(n-1)/2$ is the total number of unique combinations of n control dimensions.

4 RESULTS

We share results from two categories of evaluation tasks: sequential reaching (3D-star task) and ADLs. Our results find both task performance and perceived workload to vary between control space mappings (TS control versus JS control); we also find an inverse relationship between the control dimension couplings and ADL task performance.

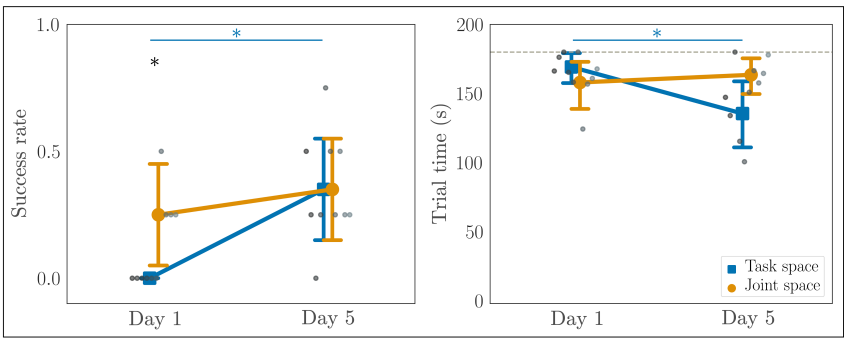

Fig. 3. Success rate (left) and trial time (right) for ADL tasks on first and last days. The dotted line represents a timeout of 180 seconds; the grayscale dots represent each participant's mean values (TS: dark; JS: light). The standard interquartile ranges are shown. $^*p < 0.05$. (*between groups, *within group).

4.1 Performance Evolutions Differ with Control Space Mappings

Intuitivity versus Learnability. Figure 3 shows the results from the ADL tasks. Both groups improve their performance in success rate between first day and last day, but only the TS group improves in trial times. We observe that the initial performance of the JS group is superior to that of the task space group. Despite the JS group starting on the first day with a higher median success rate (TS: 0.00, JS: 0.25; $p < 0.05$, Wilcoxon signed-rank test) and lower median trial time (TS: 169 s, JS: 158 s), by the final day, the success rates between groups converge (TS: 0.35, JS: 0.35) and the

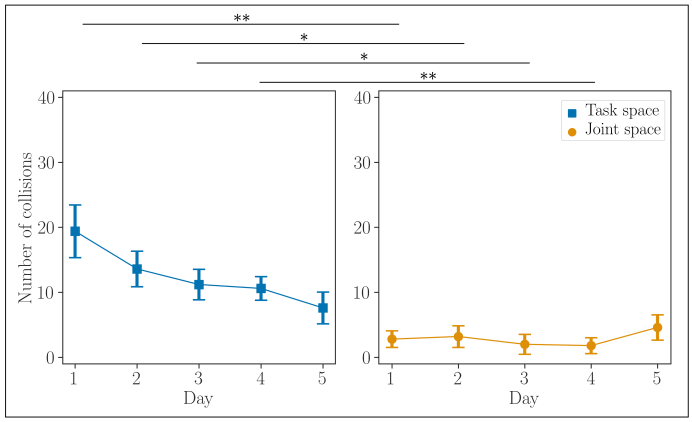

Fig. 4. Average number of collisions during the 3D-star task over five days. $^{*}p < 0.05$, $^{**}p < 0.01$.

TS group outperforms the JS group with a lower trial time (TS: 135 s, JS: 164 s). We should note that only TS demonstrates statistically significant improvements, between days 1 and 5, in success rate ($p < 0.01$) and trial time ($p < 0.01$). Thus, TS control appears to have a *greater capacity for improvement*, as measured by task success and trial time, while JS control demonstrates *greater success with naïve use.*

A similar trend is noticed when comparing the average measured collisions on the 3D-star task between groups across days (Figure 4). The TS group decreases collisions over days (by more than half between the first day and final day), whereas the JS group starts with superior performance but then shows little improvement (days 1-4, $p < 0.05$, Kruskal-Wallis H-test).

For the 3D-star task, out of the ten participants, there are no successful reaching trials (five target reaches per day; 25 reaches per person). Recall however that reaching success is defined within strict positional (1.00 cm) and rotational (0.02 rad, or 1.14°) thresholds, which are tighter position and orientation constraints than any of the ADL tasks.

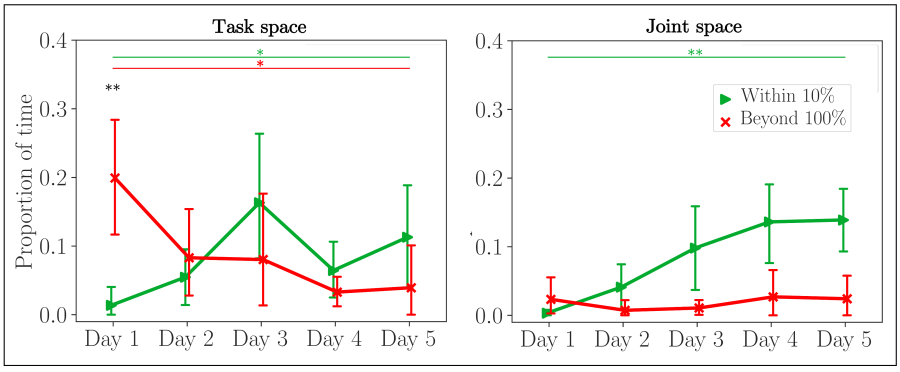

Fig. 5. Proportion of time the robot end-effector spends within 10% of reach distance (green) and outside of 100% of reach distance (red) in the 3D-star task. The proportion of time for a given reaching trial is considered out of 90 seconds (or 1.5 minutes). The standard interquartile ranges are shown. $^{*}p < 0.05$, $^{**}p < 0.01$ (*between categories, **within categories).

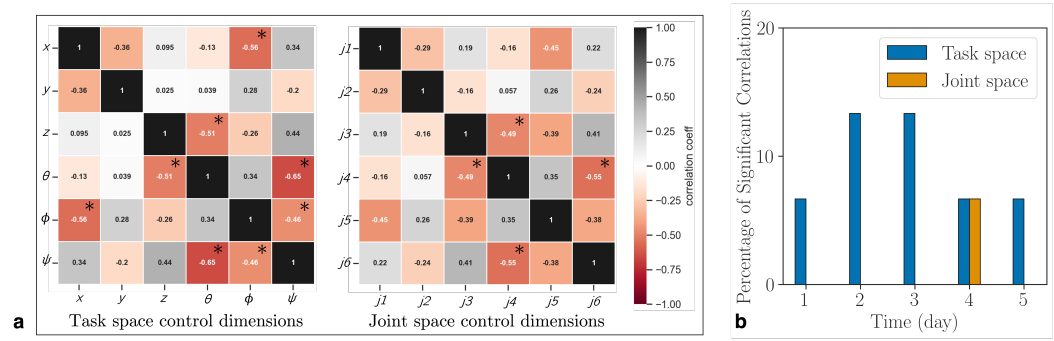

Fig. 6. Control dimension couplings between dimensions of the 6D robot control signal **q** on the 3D-star task. (a) The correlation coefficient for the task space (TS) group and joint space (JS) group. (b) Percentage of significant correlations (out of 15 possible combinations) across days. $^*p < 0.05$.

Capacity for Improvement. In the absence of successful reaches in the 3D-star task, we show in Figure 5 the proportion of time spent within 10% of the reach distance ($\tau_{dist \leq 10\%}$) and beyond 100% of the reach distance ($\tau_{dist > 100\%}$).[1] Note that while for an ideal reach $\tau_{dist > 100\%}$ would be zero and $\tau_{dist \leq 10\%}$ minimized, *during learning* an *increase* in $\tau_{dist \leq 10\%}$ is a marker of improvement when targets are not yet achievable.

Important to observe are the trends of proportion of time spent near the targets and beyond the starting distance. Specifically, on the first day, the median proportion of time both groups spend near the targets is quite low (TS: 0.0034 or 0.31 s; JS: 0.014 or 1.2 s), while the time spent beyond the starting distance is considerable for TS in particular (TS: 0.20 or 18 s; JS: 0.020 or 1.8 s). By the third day, we notice a measurable increase in $\tau_{dist \leq 10\%}$ in both groups and a decrease in $\tau_{dist > 100\%}$ in the TS group. Both TS and JS groups significantly improve the amount of time spent near the targets ($\tau_{dist \leq 10\%}$) between days 1 and 5 ($p < 0.05$, Kruskal-Wallis H-test). The TS group also significantly reduces how much time is spent beyond the starting distance ($\tau_{dist > 100\%}$), between days 1 and 5 ($p < 0.05$), whereas this stays static in JS, largely due to having less room to improve.

4.2 Impact of Control Dimension Couplings on Learning

Inherent Decoder Characteristics. In Figure 6, we evaluate the degree of control dimension coupling the two groups experience during the evaluation tasks, by computing the correlation coefficient between the decoded six-dimensional robot control signal data.

In general, it appears that the TS group experiences more control dimension couplings than the JS group, as shown both when the matrix of correlation coefficients is computed globally (all days combined) and day-by-day. Figure 6a shows four unique dimension couplings ($C_{TS} = \{\rho(x, \phi), \rho(z, \theta), \rho(\theta, \phi), \rho(\phi, \psi)\}$, $p < 0.05$) in the TS group compared to two unique dimension couplings ($C_{JS} = \{\rho(j3, j4), \rho(j4, j6)\}$, $p < 0.05$) in the JS group.

Notably, all statistically significant ($p < 0.05$) couplings show linearly *inverse* relationships between control dimensions—meaning that positive motion in one dimension is coupled with negative motion in the other dimension (and vice versa). Furthermore, despite the differences in control space mappings (TS versus JS), two coupled dimensions (i.e., corresponding to the same PCs) appear to be shared between the two groups: $C_{TS} = \{\rho(z, \theta), \rho(\theta, \psi)\} \leftrightarrow \{\rho(j3, j4), \rho(j4, j6)\} =$

[1]A simple binary result of success is not informative, as no participants achieved the target location within our positional (1.00 cm) and rotational (0.02 rad or 1.14°) constraints on success. We also find the proportion of time metrics to be more informative than path length, for which no discernible trends emerge.

C_{JS}. This suggests that either the body movements themselves are coupled or the decoder might contain intrinsic synergies (or biases) that make it prone to coactivity of multiple robot control signals at a time. As such, this result could be caused by unintended issuing of commands rather than intentional, coordinated robot motions required for task completion.

Comparing (Figure 6b) the two groups' evolution of couplings that are statistically significant (ρ^*) suggests that the user experience of couplings are more present and persistent for the TS group (1 or 2 couplings out of 15 possible combinations on each of the five days) than the JS group (1 coupling on day 4).

Control Dimension Couplings and Task Performance. To examine the importance of control dimension coupling, we evaluate its relationship to performance on the ADL tasks. We first compute the average correlation coefficient for each participant and compare the two groups in Figure 7. For the TS group in particular, we see signs of an *inverse relationship* between control dimension coupling and success rate—that is, the smaller (or less present) the average degree of control dimension coupling, the higher the success rate on the ADL tasks (and vice versa, TS: $R^2 = 0.643$, $p = 0.103$, JS: $R^2 = 0.066$, $p = 0.676$). We have shown in Figures 6 and 7 that the influence of control dimension couplings is higher in the TS group for the 3D-star task in both the total number of statistically significant couplings and how they were distributed across all five days. These could be explanations for why the effect sizes of average correlation coefficient to performance was more noticeable in the TS group than in the JS group.

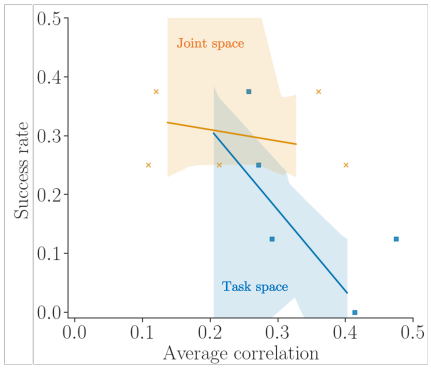

Fig. 7. Linear regression analysis of average ADL task success rate versus average control dimension coupling, for both study groups. Shaded regions represent the 95% confidence interval.

We also note that a similar analysis of trial time did not show any strong correlations with dimension coupling.

4.3 Effects of Perceived Workload on Learning

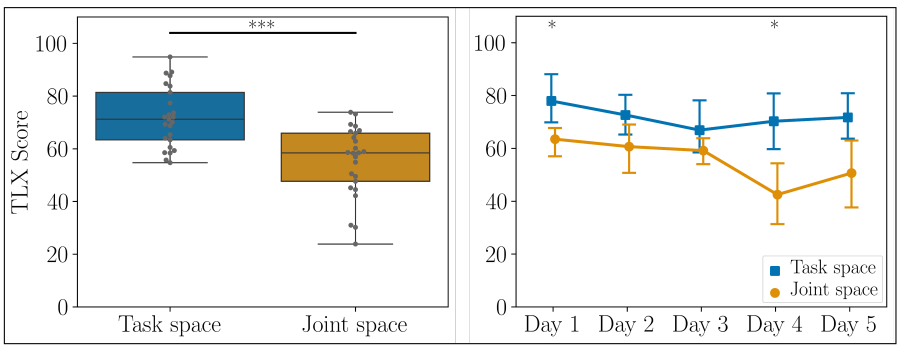

Fig. 8. Comparison of subjective workload, measured via NASA-TLX, between study groups. Summary NASA-TLX scores, averaged across study sessions (left), and evolution of NASA-TLX scores over study sessions (right). The standard interquartile ranges are shown. $^*p < 0.05$, $^{***}p < 0.001$.

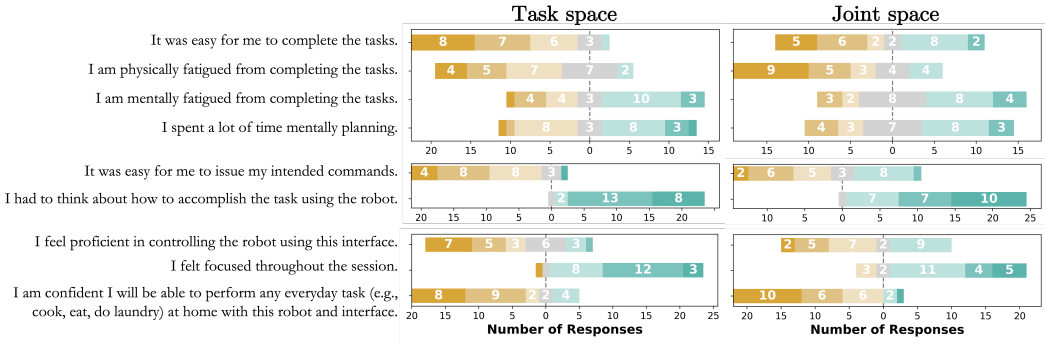

Fig. 9. Comparison of response counts from the Likert-scale questionnaire between study groups. Topics include task difficulty, human-robot interaction, and learning effects. The scale ranges from 1 (strongly disagree, brown) to 7 (strongly agree, teal), and the counts of responses are reported. Questionnaire responses are gathered at the end of each study session.

A Different Variety of Learning. The NASA-TLX assessment scores for subjective workload and their evolution across days are shown in Figure 8.

Figure 8 (left) shows that the median score (over all days and participants) of the TS group is significantly higher than that of the JS group ($p < 0.001$), indicating higher perceived workloads in the TS group than the JS group.

In Figure 8 (right), we observe a marked reduction over sessions in perceived workload (median NASA-TLX score) for the JS group. By contrast, we observe only a slight decrease in perceived workload for the TS group. Furthermore, the perceived workload of the JS group is consistently lower than that of the TS group across days—including on the first day, when the JS group's performance was higher.

To evaluate statistical significance between the two groups, we initially use the Kruskal-Wallis H-test to find a main effect, and Conover's *post hoc* pairwise test, with Bonferroni adjustments, to make appropriate corrections. Only on days one and four do we find statistically significant differences between the two groups ($p < 0.05$).

We recall that the JS group does not improve much according to either of the ADL task performance metrics of success or trial time. While the JS group does not improve significantly in task performance, the group *does* improve in perceived workload—which perhaps is indicative of learning, albeit of a different variety than task performance learning (or at the very least familiarization).

We also perform a linear regression analysis between NASA-TLX scores and ADL task performance (success rate and trial time) and find that there are no strong correlations.

Figure 9 reports a group comparison of the results from the Likert-scale study questionnaire. In general, we see that participants share relatively similar views on questions regarding difficulty, human-robot interaction, and learning effects. More specifically, participants from both groups seem to find the study tasks to be more *mentally* fatiguing than physically fatiguing, where significant mental energy is spent on planning. Additionally, it appears that the participants find it very difficult to issue intended commands (relevant to task difficulty) and thinking about how to accomplish the task using the robot (relevant to mental fatigue).

Interestingly, most participants were not confident that they would be able to perform everyday tasks with this interface-robot system at home.

Examining the trends of these questions over the study week (Figure 10 in the Appendix), shows that the mental fatigue reduces over time, physical fatigue increases, and tasks seem to get easier.

These trends also suggest that the difficulty may have decreased partly because issuing intended commands got easier and both robot embodiment and confidence increased.

Learning Takes Work. The TS group does not improve measurably with respect to perceived workload (Figure 8, right). This group, however, does improve according to both performance metrics, of success and trial time (Figure 3). Thus, a possible explanation is that the gains in performance are expensive to acquire—simply put, learning takes work. An alternate (or parallel) explanation is that the learning efforts of the JS group prioritize managing cognitive workload in lieu of task performance (perhaps because a performance ceiling has been met).

5 DISCUSSION

Body-machine interfacing provides a unique opportunity to control high-DoF robots and can play a role to help restore functionality and agency to individuals with severe paralysis. We have demonstrated that its utility is not limited to low-DoF robots alone and can scale to 6-DoF control of assistive robotic arms.

5.1 Challenges in Learning High-DoF Control

When an interface is novel, a person may not initially have the implicit knowledge of which interface commands (or how to execute them) lead to which robot consequences. As the person learns the mappings between the interface signal and robot behavior, they may have only visual feedback to assist with error correction. In addition to the novelty of the interface, the introduction of non-anthropomorphic high-DoF robot platforms contributes to the learning challenges.

There are characteristics inherent to the decoder that might contribute to the learning burden. In Figures 6 and 7, we identified a subset of control dimension couplings that are shared between study groups. If one learning strategy for an operator was to isolate robot control dimensions, they would need to devote additional time and cognitive resources to try to decouple dimensions, which in practice may not be achievable. Furthermore, the variance in the data, as explained by PCA, can be disproportionally contained within a small subset of components [40]. We also suspect that needing to learn distinct movements for each direction of a control dimension introduces an additional layer of cognitive burden to the human operator. With PCA, the reflexivity of control axes is not guaranteed to be preserved between body motion space and control space, leading to the need for operators to learn potentially twice the number of mappings (12 instead of 6). This might have contributed to the reported mental challenges in planning, trying to issue intended commands, and thinking about how to use the robot to accomplish physical tasks, which resulted in high levels of mental fatigue (Figures 8, 9, and 10).

We also have shown that the choice of control space mapping leads to different learning profiles. The mechanical simplicity of single joint movements—as opposed to movements along a single Cartesian-space dimension, which often requires a coordination of multiple joints—is a viable explanation for why people tend to intuit JS control, prior to any practice, more than TS control. This is marked by the early performance lead in the JS group observed in Figures 3, 4, and 5. Our results on robotic arm control corroborates with what is commonly observed in cerebellar ataxia, where the lesioned cerebellum causes patients to think out individual joint movements in their own arm rather than being able to coordinate multi-jointed movements, which are more complex. The biological underpinnings that lead to decomposed movements in ataxia are still debated [3, 27, 31]. We should also note that, in addition, TS control forces users to learn the forward kinematics of the robotic arm (whereas the JS control does not).

Not only do the learning profiles differ with respect to task proficiency, they also differ in regards to *what was learned*. Learning is generally assessed with respect to task performance. Equally

important within the field of assistive and rehabilitation robotics, however, is the burden on the human operator. Learning to interface with the robot with a lower workload also is learning, and it achieves one of the driving motivators for the development of assistive robots. Broadly, this is critical when such robots have had a history of acceptance issues [5]. The presented study observed mental planning and mental fatigue to be some of the leading challenges participants faced when learning to control a robot in high DoFs using the BoMI. These challenges seem to be presented as larger effects early on rather than later in the study (Figure 10).

5.2 Potential for Robot Autonomy

Controlling high-DoF robots is a challenging task, even with an interface that covers the full control space of the robot. Therefore, there is potential for introducing robot autonomy to facilitate both human learning and robot control. In order to design explicit training regimes that allow for smooth interactions between humans and robots, facilitate human learning, and cater training to the individual, it is important to acquire a baseline understanding of how well people can adapt to novel human-machine systems.

This study takes a first step towards the design of intelligent robot systems that aim to facilitate human training by (a) assessing the evolution of native human interaction and learning in both control space mapping paradigms (Figures 3, 4, 5, and 10), as well as (b) investigating the influence of control dimension coupling on task performance (Figures 6 and 7). Designing robotics autonomy to provide assistance that substitutes for and mitigates the challenges we have identified in this paper could potentially lead to higher-dimensional interfaces and assistive robots becoming more accessible to those with severe levels of paralysis or injury.

5.3 Future Work

A significant limitation to this work is that our study considers only an uninjured population. Recall that the key motivation for this work was to obtain a baseline understanding of high-DoF robot teleoperation using a BoMI.

While the BoMI previously had been shown to be effective at adjusting to the available residual movements in patients with severe paralysis (for two-dimensional control [38]), questions related to the ability of human users to learn to teleoperate a robotic arm in higher DoFs using a BoMI—and whether control space mappings have any effect towards facilitating learning—were largely unanswered. Having now determined a baseline for human learning, our next steps will be to build on this work and apply the gained knowledge to a diverse population of patients with paralysis [34].

While our investigation over five sessions did demonstrate evolving learning effects, we have yet to uncover where exactly the inflection points lie on the human learning curve. We should note that a one-session study would have shown evidence that JS control was superior to TS control for operating high-DoF robots with a BoMI, and that TS control was unlearnable (Figures 3, 4, and 5). Instead, over five sessions, we have observed that TS control was in fact *more learnable* than JS control, which remained largely static over time. It is possible that task performance, especially for the TS group, would have continued to improve in a longer study and that we have not yet observed the full capability of human learning on this system. We expect that further gains might be possible with a more regimented learning curriculum that further customizes the presentation of tasks and control access to each individual.

Lastly, while there has been a rich history of studies that used PCA to capture residual body movement for lower-dimensional control [14, 37, 41], the use of an alternative mapping paradigm could help alleviate some learning burden and improve functionality, intuitivity, and learnability. Modern machine learning approaches [15, 44] and hybrid interface systems [46] that use additional

types of sensors (e.g., EMG) hold much promise as we make improvements on our decoder design, introduce robot autonomy, and scale our investigation aimed at more diverse populations.

6 CONCLUSIONS

People with limited mobility can benefit from assistive and rehabilitation technologies that help return functionality and agency. The integration of body-machine interfacing with high-DoF assistive robots is one such example. We have presented in this paper a multi-session study of robotic arm teleoperation with a BoMI, to address the question of to what extent the control of a high-DoF robot is feasible with a high-dimensional interface. While feasibility has been demonstrated, a number of challenges were also identified—paving the way for potential assistance from robotics autonomy, more advanced decoding approaches, and longer investigations with a more diverse population. The results of our study suggest that the choice of control space impacts both the initial performance and its evolution during learning. While JS control was found to be *more intuitive* prior to training, TS control was subject to greater improvement over time and thus presented as *more learnable* with respect to task proficiency. Our results further found an inverse relationship between control dimension couplings and performance. Gains in task performance were coupled with higher cognitive load, while static task performance paired with decreasing cognitive load. These results suggest a trade-off between intuitiveness and learnability, when comparing the two control space mappings, as well as between improving task performance and reducing cognitive load. Both of these learning curves merit further investigation—with patients with motor impairments and with longer studies—that more deeply probe their potential points of inflection and plateaus.

ACKNOWLEDGMENTS

Research reported in this publication was supported by the National Institute of Biomedical Imaging and Bioengineering (NIBIB) under Award Number R01-EB024058; the Eunice Kennedy Shriver National Institute of Child Health & Human Development (NICHD) of the National Institutes of Health (NIH), Award Number T32-HD007418; the National Science Foundation (NSF), Award Number 2054406; National Institute on Disability, Independent Living and Rehabilitation Research (NIDILRR), Award Number 90REGE0005-01-00; and European Union's Horizon 2020 Research and Innovation Program under the Marie Sklodowska-Curie, Project REBoT, Award Number GA-750464. The content is solely the responsibility of the authors and does not necessarily represent the official views of the NIH.

We thank Andrew Thompson for his expertise and efforts with 3D printing and building the experiment cage. We thank Larisa Loke and Andrew Thompson for their pre-study participation.

REFERENCES

[1] F. Abdollahi, A. Farshchiansadegh, C. Pierella, I. Seanez-Gonzalez, E. Thorp, M. H. Lee, R. Ranganathan, J. Pedersen, D. Chen, E. Roth, M. Casadio, and F. Mussa-Ivaldi. 2017. Body-Machine Interface Enables People With Cervical Spinal Cord Injury to Control Devices With Available Body Movements: Proof of Concept. *Neurorehabil Neural Repair* 31, 5 (2017), 487–493. https://doi.org/10.1177/1545968317693111

[2] Kim D Anderson. 2004. Targeting recovery: priorities of the spinal cord-injured population. *Journal of neurotrauma* 21, 10 (2004), 1371–1383.

[3] Amy J Bastian, TA Martin, JG Keating, and WT Thach. 1996. Cerebellar ataxia: abnormal control of interaction torques across multiple joints. *Journal of neurophysiology* 76, 1 (1996), 492–509.

[4] William F Battig. 1972. Interference During Learning as a Sources of Facilitation in Subsequent Retention and Transfer.

[5] Elaine Biddiss and Tom Chau. 2007. The roles of predisposing characteristics, established need, and enabling resources on upper extremity prosthesis use and abandonment. *Disability and Rehabilitation: Assistive Technology* 2, 2 (2007), 71–84.

[6] M Casadio, A Pressman, S Acosta, Z Danzinger, A Fishbach, FA Mussa-Ivaldi, K Muir, H Tseng, and D Chen. 2011. Body machine interface: remapping motor skills after spinal cord injury. In *2011 IEEE International Conference on*

Rehabilitation Robotics. IEEE, 1–6.

[7] Maura Casadio, Assaf Pressman, Alon Fishbach, Zachary Danziger, Santiago Acosta, David Chen, Hsiang-Yi Tseng, and Ferdinando A Mussa-Ivaldi. 2010. Functional reorganization of upper-body movement after spinal cord injury. *Experimental brain research* 207, 3 (2010), 233–247.

[8] Maura Casadio, Rajiv Ranganathan, and Ferdinando A Mussa-Ivaldi. 2012. The body-machine interface: a new perspective on an old theme. *Journal of Motor behavior* 44, 6 (2012), 419–433.

[9] EK Chadwick, D Blana, JD Simeral, J Lambrecht, Sung-Phil Kim, AS Cornwell, DM Taylor, LR Hochberg, JP Donoghue, and RF Kirsch. 2011. Continuous neuronal ensemble control of simulated arm reaching by a human with tetraplegia. *Journal of neural engineering* 8, 3 (2011), 034003.

[10] Aswin Chari, Sanjay Budhdeo, Rachel Sparks, Damiano G Barone, Hani J Marcus, Erlick AC Pereira, and Martin M Tisdall. 2021. Brain–machine interfaces: the role of the neurosurgeon. *World Neurosurgery* 146 (2021), 140–147.

[11] Sheryl Chau, Sanders Aspelund, Ranjan Mukherjee, Mei-Hua Lee, Rajiv Ranganathan, and Florian Kagerer. 2017. A five degree-of-freedom body-machine interface for children with severe motor impairments. In *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 3877–3882.

[12] Jennifer L Collinger, Brian Wodlinger, John E Downey, Wei Wang, Elizabeth C Tyler-Kabara, Douglas J Weber, Angus JC McMorland, Meel Velliste, Michael L Boninger, and Andrew B Schwartz. 2013. High-performance neuroprosthetic control by an individual with tetraplegia. *The Lancet* 381, 9866 (2013), 557–564.

[13] Dalia De Santis and Ferdinando A. Mussa-Ivaldi. 2020. Guiding functional reorganization of motor redundancy using a body-machine interface. *J Neuroeng Rehabil* 17, 1 (2020), 61. https://doi.org/10.1186/s12984-020-00681-7

[14] A. Farshchiansadegh, F. Abdollahi, D. Chen, Lee Mei-Hua, J. Pedersen, C. Pierella, E. J. Roth, I. Seanez Gonzalez, E. B. Thorp, and F. A. Mussa-Ivaldi. 2014. A body machine interface based on inertial sensors. *Conf Proc IEEE Eng Med Biol Soc* 2014 (2014), 6120–4. https://doi.org/10.1109/EMBC.2014.6945026

[15] Joshua I Glaser, Ari S Benjamin, Raeed H Chowdhury, Matthew G Perich, Lee E Miller, and Konrad P Kording. 2020. Machine learning for neural decoding. *Eneuro* 7, 4 (2020).

[16] Sandra G. Hart. 2006. Nasa-Task Load Index (NASA-TLX); 20 Years Later. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting* 50, 9 (2006), 904–908. https://doi.org/10.1177/154193120605000909

[17] Laura V Herlant, Rachel M Holladay, and Siddhartha S Srinivasa. 2016. Assistive teleoperation of robot arms via automatic time-optimal mode switching. *2016 11th ACM/IEEE International Conference on Human-Robot Interaction (HRI)* (2016), 35–42.

[18] Siddarth Jain, Ali Farshchiansadegh, Alexander Broad, Farnaz Abdollahi, Ferdinando Mussa-Ivaldi, and Brenna Argall. 2015. Assistive robotic manipulation through shared autonomy and a body-machine interface. In *2015 IEEE international conference on rehabilitation robotics (ICORR)*. IEEE, IEEE Press, Singapore, 526–531.

[19] Henry F Kaiser. 1958. The varimax criterion for analytic rotation in factor analysis. *Psychometrika* 23, 3 (1958), 187–200.

[20] Brian Lee, Charles Y Liu, and Michael LJ Apuzzo. 2013. A primer on brain–machine interfaces, concepts, and technology: a key element in the future of functional neurorestoration. *World neurosurgery* 79, 3-4 (2013), 457–471.

[21] Jongmin M Lee, Temesgen Gebrekristos, Dalia De Santis, Mahdieh Nejati Javaremi, Deepak Gopinath, Biraj Parikh, Ferdinando A Mussa-Ivaldi, and Brenna D Argall. 2023. An Exploratory Multi-Session Study of Learning High-Dimensional Body-Machine Interfacing for Assistive Robot Control. In *IEEE-RAS-EMBS International Conference on Rehabilitation Robotics (ICORR)*. IEEE.

[22] Keith Lohse, Navid Shirzad, Alida Verster, Nicola Hodges, and HF Machiel Van der Loos. 2013. Video games and rehabilitation: using design principles to enhance engagement in physical therapy. *Journal of Neurologic Physical Therapy* 37, 4 (2013), 166–175.

[23] Dylan P Losey, Hong Jun Jeon, Mengxi Li, Krishnan Srinivasan, Ajay Mandlekar, Animesh Garg, Jeannette Bohg, and Dorsa Sadigh. 2022. Learning latent actions to control assistive robots. *Autonomous robots* 46, 1 (2022), 115–147.

[24] Matteo Macchini, Fabrizio Schiano, and Dario Floreano. 2019. Personalized telerobotics by fast machine learning of body-machine interfaces. *IEEE Robotics and Automation Letters* 5, 1 (2019), 179–186.

[25] Veronique Maheu, Philippe S Archambault, Julie Frappier, and François Routhier. 2011. Evaluation of the JACO robotic arm: Clinico-economic study for powered wheelchair users with upper-extremity disabilities. In *2011 IEEE international conference on rehabilitation robotics*. IEEE, 1–5.

[26] Richard M Mahoney. 2001. The raptor wheelchair robot system. *Integration of assistive technology in the information age* 43, 10 (2001), 135–141.

[27] Jon Marsden and Chris Harris. 2011. Cerebellar ataxia: pathophysiology and rehabilitation. *Clinical rehabilitation* 25, 3 (2011), 195–216.

[28] Michael L Martini, Eric Karl Oermann, Nicholas L Opie, Fedor Panov, Thomas Oxley, and Kurt Yaeger. 2020. Sensor modalities for brain-computer interface technology: a comprehensive literature review. *Neurosurgery* 86, 2 (2020), E108–E117.

[29] WB McKay, HK Lim, MM Priebe, DS Stokic, and AM Sherwood. 2004. Clinical neurophysiological assessment of residual motor control in post-spinal cord injury paralysis. *Neurorehabilitation and neural repair* 18, 3 (2004), 144–153.

[30] Jenifer Miehlbradt, Alexandre Cherpillod, Stefano Mintchev, Martina Coscia, Fiorenzo Artoni, Dario Floreano, and Silvestro Micera. 2018. Data-driven body–machine interface for the accurate control of drones. *Proceedings of the National Academy of Sciences* 115, 31 (2018), 7913–7918.

[31] Susanne M Morton and Amy J Bastian. 2007. Mechanisms of cerebellar gait ataxia. *The cerebellum* 6, 1 (2007), 79–86.

[32] Gillian D Muir and John D Steeves. 1997. Sensorimotor stimulation to improve locomotor recovery after spinal cord injury. *Trends in neurosciences* 20, 2 (1997), 72–77.

[33] Ferdinando A Mussa-Ivaldi and Lee E Miller. 2003. Brain–machine interfaces: computational demands and clinical needs meet basic neuroscience. *TRENDS in Neurosciences* 26, 6 (2003), 329–334.

[34] Kemal Nas, Levent Yazmalar, Volkan Şah, Abdulkadir Aydın, and Kadriye Öneş. 2015. Rehabilitation of spinal cord injuries. *World journal of orthopedics* 6, 1 (2015), 8.

[35] Karl Pearson. 1901. LIII. On lines and planes of closest fit to systems of points in space. *The London, Edinburgh, and Dublin philosophical magazine and journal of science* 2, 11 (1901), 559–572.

[36] Camilla Pierella, Farnaz Abdollahi, Ali Farshchiansadegh, Jessica Pedersen, David Chen, Ferdinando A Mussa-Ivaldi, and Maura Casadio. 2014. Body machine interfaces for neuromotor rehabilitation: a case study. In *2014 36th Annual International Conference of the IEEE Engineering in Medicine and Biology Society*. IEEE, IEEE Press, 397–401. https://doi.org/10.1109/EMBC.2014.6943612

[37] Camilla Pierella, Farnaz Abdollahi, Ali Farshchiansadegh, Jessica Pedersen, Elias B Thorp, Ferdinando A Mussa-Ivaldi, and Maura Casadio. 2015. Remapping residual coordination for controlling assistive devices and recovering motor functions. *Neuropsychologia* 79 (2015), 364–376.

[38] Camilla Pierella, Farnaz Abdollahi, Elias Thorp, Ali Farshchiansadegh, Jessica Pedersen, Ismael Seáñez-González, Ferdinando A Mussa-Ivaldi, and Maura Casadio. 2017. Learning new movements after paralysis: Results from a home-based study. *Scientific reports* 7, 1 (2017), 1–11.

[39] Camilla Pierella, Elisa Galofaro, Alice De Luca, Luca Losio, Simona Gamba, Antonino Massone, Ferdinando A Mussa-Ivaldi, and Maura Casadio. 2021. Recovery of distal arm movements in spinal cord injured patients with a body-machine interface: a proof-of-concept study. *Sensors* 21, 6 (2021), 2243.

[40] Alexandra A Portnova-Fahreeva, Fabio Rizzoglio, Ilana Nisky, Maura Casadio, Ferdinando A Mussa-Ivaldi, and Eric Rombokas. 2020. Linear and non-linear dimensionality-reduction techniques on full hand kinematics. *Frontiers in bioengineering and biotechnology* 8 (2020), 429.

[41] Rajiv Ranganathan, Mei-Hua Lee, Malavika R Padmanabhan, Sanders Aspelund, Florian A Kagerer, and Ranjan Mukherjee. 2019. Age-dependent differences in learning to control a robot arm using a body-machine interface. *Scientific reports* 9, 1 (2019), 1–9.

[42] Rajiv Ranganathan, Jon Wieser, Kristine M Mosier, Ferdinando A Mussa-Ivaldi, and Robert A Scheidt. 2014. Learning redundant motor tasks with and without overlapping dimensions: facilitation and interference effects. *Journal of Neuroscience* 34, 24 (2014), 8289–8299.

[43] Daniel J Rea and Stela H Seo. 2022. Still Not Solved: A Call for Renewed Focus on User-Centered Teleoperation Interfaces. *Front Robot AI* 9 (2022).

[44] Fabio Rizzoglio, Maura Casadio, Dalia De Santis, and Ferdinando A Mussa-Ivaldi. 2021. Building an adaptive interface via unsupervised tracking of latent manifolds. *Neural Networks* 137 (2021), 174–187.

[45] F. Rizzoglio, C. Pierella, D. De Santis, F. Mussa-Ivaldi, and M. Casadio. 2020. A hybrid Body-Machine Interface integrating signals from muscles and motions. *J Neural Eng* 17, 4 (2020), 046004. https://doi.org/10.1088/1741-2552/ab9b6c

[46] Fabio Rizzoglio, Camilla Pierella, Dalia De Santis, Ferdinando Mussa-Ivaldi, and Maura Casadio. 2020. A hybrid Body-Machine Interface integrating signals from muscles and motions. *Journal of Neural Engineering* 17, 4 (2020), 046004.

[47] Fabio Rizzoglio, Francesca Sciandra, Elisa Galofaro, Luca Losio, Elisabetta Quinland, Clara Leoncini, Antonino Massone, Ferdinando A Mussa-Ivaldi, and Maura Casadio. 2019. A myoelectric computer interface for reducing abnormal muscle activations after spinal cord injury. In *2019 IEEE 16th International Conference on Rehabilitation Robotics (ICORR)*. IEEE, IEEE Press, 1049–1054.

[48] JC Rosier, JA Van Woerden, LW Van der Kolk, BJF Driessen, HH Kwee, JJ Duimel, JJ Smits, AA Tuinhof de Moed, G Honderd, and PM Bruyn. 1991. Rehabilitation robotics: The MANUS concept. In *Fifth International Conference on Advanced Robotics' Robots in Unstructured Environments*. IEEE, IEEE Press, 893–898.

[49] Alan W Salmoni, Richard A Schmidt, and Charles B Walter. 1984. Knowledge of results and motor learning: a review and critical reappraisal. *Psychological bulletin* 95, 3 (1984), 355.

[50] Félix Scholtes, Gary Brook, and Didier Martin. 2012. Spinal cord injury and its treatment: current management and experimental perspectives. In *Advances and Technical Standards in Neurosurgery*. Springer, 29–56.

[51] Reza Shadmehr and Ferdinando A Mussa-Ivaldi. 1994. Adaptive representation of dynamics during learning of a motor task. *Journal of neuroscience* 14, 5 (1994), 3208–3224.
[52] John B Shea and Robyn L Morgan. 1979. Contextual interference effects on the acquisition, retention, and transfer of a motor skill. *Journal of Experimental psychology: Human Learning and memory* 5, 2 (1979), 179.
[53] Student. 1908. Probable error of a correlation coefficient. *Biometrika* 6, 2-3 (1908), 302–310.
[54] Elias B Thorp, Farnaz Abdollahi, David Chen, Ali Farshchiansadegh, Mei-Hua Lee, Jessica P Pedersen, Camilla Pierella, Elliot J Roth, Ismael Seáñez Gonzáles, and Ferdinando A Mussa-Ivaldi. 2015. Upper body-based power wheelchair control interface for individuals with tetraplegia. *IEEE Transactions on Neural Systems and Rehabilitation Rngineering* 24, 2 (2015), 249–260.

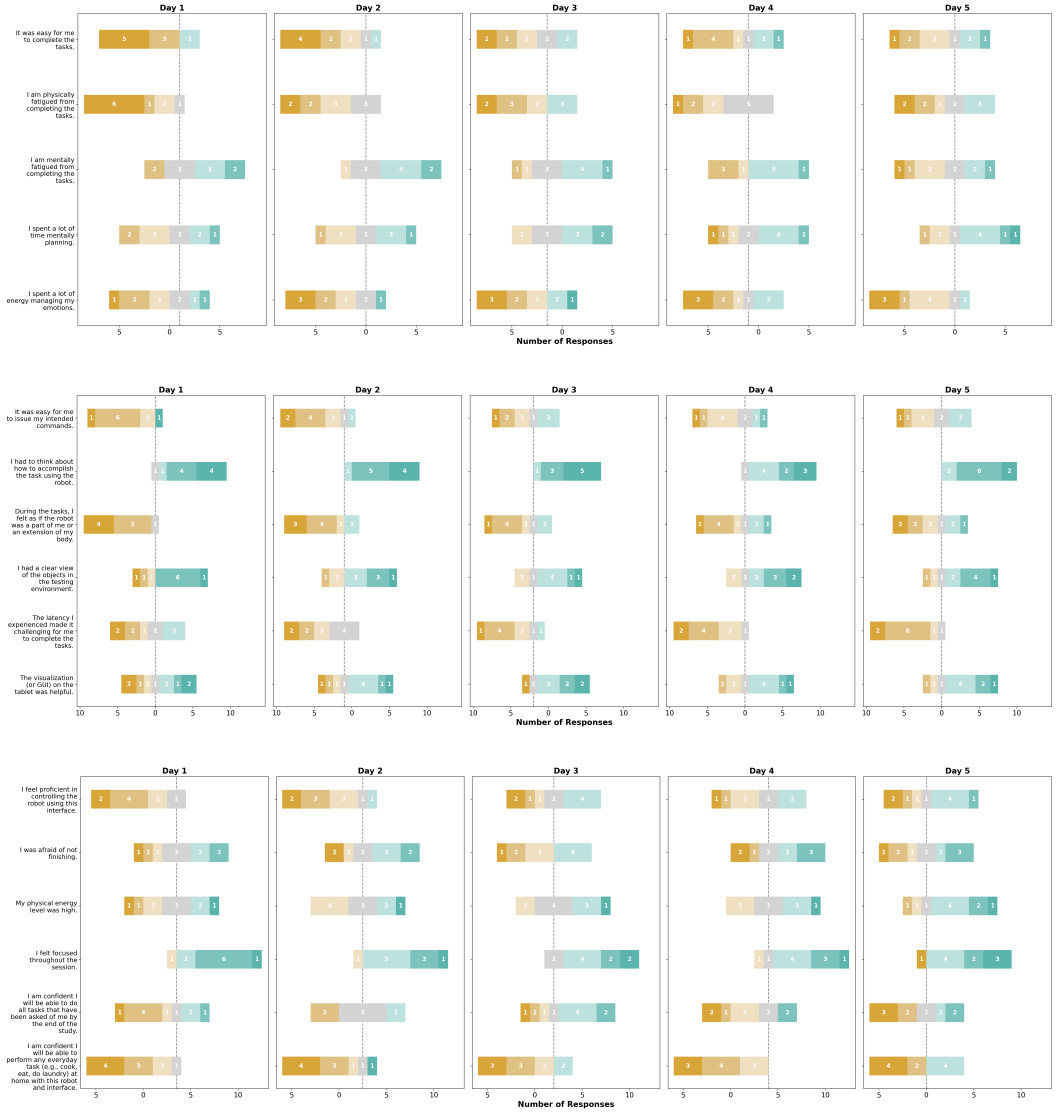

Fig. 10. Evolution of Likert-scale daily questionnaire responses over days. Themes included (top) task difficulty, (middle) human-robot interaction, (bottom) learning effects. Questionnaires were provided to participants using an electronic tablet after each study session concluded.