

Learning to Control Complex Robots Using High-Dimensional Interfaces: Preliminary Insights

Jongmin M. Lee¹, Temesgen Gebrekristos¹, Dalia De Santis², Mahdiah Nejati-Javaremi¹, Deepak Gopinath¹, Biraj Parikh¹, Ferdinando A. Mussa-Ivaldi³, Brenna D. Argall³

^{1,3}Northwestern University, Chicago, Illinois

^{1,2,3}Shirley Ryan AbilityLab, Chicago, Illinois

²Istituto Italiano di Tecnologia, Genoa, Italy

¹{jmlee, tem, m.nejati, deepakgopinath, bparikh}@u.northwestern.edu, ²dalia.desantis@iit.it, ³{sandro, brenna.argall}@northwestern.edu

Abstract

Human body motions can be captured as a high-dimensional continuous signal using motion sensor technologies. The resulting data can be surprisingly rich in information, even when captured from persons with limited mobility. In this work, we explore the use of limited upper-body motions, captured via motion sensors, as inputs to control a 7 degree-of-freedom assistive robotic arm. It is possible that even dense sensor signals lack the salient information and independence necessary for reliable high-dimensional robot control. As the human learns over time in the context of this limitation, intelligence on the robot can be leveraged to better identify key learning challenges, provide useful feedback, and support individuals until the challenges are managed. In this short paper, we examine two uninjured participants' data from an ongoing study, to extract preliminary results and share insights. We observe opportunities for robot intelligence to step in, including the identification of inconsistencies in time spent across all control dimensions, asymmetries in individual control dimensions, and user progress in learning. Machine reasoning about these situations may facilitate novel interface learning in the future.

Introduction

Motion sensor technologies have been used to interface a person's body movements to control machines such as assistive and rehabilitation devices and robots (Casadio, Ranganathan, and Mussa-Ivaldi 2012; Jain et al. 2015), drones (Miehlbradt et al. 2018), and quadcopters (Macchini, Schiano, and Floreano 2019). A common strategy for robot control using body motions is to engineer a decoder designed to map the high-dimensional body motion to a lower-dimensional robot control signal space. Whenever the body motion has an intrinsic dimension higher than the device to be controlled, dimensionality reduction techniques, such as principal component analysis (PCA) (Wold, Esbensen, and Geladi 1987) or autoencoders (Kramer 1991) can be used to implement efficient simultaneous and continuous control of lower-dimensional devices (Pierella et al. 2018; Ranganathan et al. 2019; Rizzoglio et al. 2021; Thorp et al. 2015). However, the design and the operation of such interfaces become challenging when redundancy of the body

signals is reduced due to pathological conditions that impact mobility or when controlling complex multi-articulated robotic devices (Chau et al. 2017; Ison et al. 2015).

These challenges provide a ripe opportunity for robotics autonomy to assist the user (Gopinath, Jain, and Argall 2016; Losey et al. 2020; Muelling et al. 2017). For instance, identifying circumstances when a user is control deficient and offering support, may not only benefit both long- and short-term performance, but also help to build trust in assistive and rehabilitation machines (Fasola and Mataric 2012; Langer et al. 2019).

In this short paper we present preliminary observations, analyses, and insights on data gathered from two uninjured participants, within an ongoing study, in which a 7 degree-of-freedom (DoF) robot arm is controlled, using a net of sensors on the upper body. Study tasks are designed to familiarize, train, and evaluate robot arm operation via this sensor net, including on Activities of Daily Living (ADL) functional tasks. We describe these experimental methods in the METHODS section, share our immediate results in the RESULTS section, and wrap-up our key takeaways and future work in the CONCLUSION AND FUTURE WORK section.

Methods

Materials. The sensor net consists of four inertial measurement unit (IMU) sensors (Yost Labs, Ohio, USA), placed bilaterally on the scapulae and upper arms and anchored to a custom shirt designed to minimize movement artifacts. This is the essence of what is known as the body-machine interface (Casadio, Ranganathan, and Mussa-Ivaldi 2012). The relative quaternion orientation of the four IMUs in the net (16-dimensional) is mapped to a 6-dimensional subspace using PCA. The PCA map is precomputed using data from an experienced user, performing a predefined set of movements, and this same map is used for all participants. The lower-dimensional subspace consists of 6D velocity commands—3D position (x, y, z) and 3D rotation (*roll, pitch, yaw*)—which are used online to control a 7-DoF JACO robotic arm (Kinova Robotics, Quebec, Canada). A GUI is displayed on a tablet to provide a visualization, for the participant, of the robot velocity control commands as well as a score for each trial.

Protocol. There are three phases to the study protocol: (a)




Figure 1: An overview of the interface-robot pipeline and the study tasks.

familiarization, (b) training, and (c) evaluation (Figure 1). During *familiarization*, participants are encouraged to explore and become familiar with the system on their own, with minimal constraints enforced. Both of the next phases make use of a set of ten fixed targets \mathcal{G} . During *training*, two categories of reaching tasks are employed: reaches from a fixed center position out to a target $g_i \in \mathcal{G}$, and sequential reaches between multiple targets $g_j \in \mathcal{G}$. The ordering of targets is random and balanced across days to avoid ordering effects, and it is identical across participants. The *evaluation* phase is split into a reaching and a functional task. In the reaching task, participants reach to five targets that comprise a 3D-star $g_k \in \mathcal{G}$ in fixed succession. The functional 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 mask in the trash bin.

A trial ends upon successful completion or timeout. For reaching any target $g \in \mathcal{G}$, success is defined within a strict positional (1.00 cm) and rotational (0.02 rad, or 1.14°) threshold, and the timeout is 1.5 minutes. For the functional tasks, experimenters follow codified guidelines to determine when the tasks complete and the timeout is 3 minutes. 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 intervene and teleoperate the robot to a safe position before proceeding.

Participants. Each participant completes five sessions, executed on consecutive days for approximately two hours each. All sessions are conducted with the approval of the Northwestern University IRB, and all participants provide their informed consent. Two uninjured participants from this ongoing study are reported in this paper. P1 is a 31-year-old male, and P2 is a 29-year-old female; both participants are right-handed.

Results

Control Access and Asymmetries. We characterize the two participants’ control access by tracking the time spent moving the robot in each of the task space control dimensions ($x, y, z, roll, pitch, yaw$) during the 3D-star task. The percentage of time spent along positive and negative directions of each of the six observed control dimensions is shown in Figure 2.

P1 initially, on Day 1, spends a majority of time in only two control dimensions (31% in x , and 44% in $roll$); this is not as apparent for P2. Access to control dimensions, for both participants, is generally asymmetric, with each participant

largely accessing either positive direction or negative direction of each control dimension.

Over time, the distribution of control access tends to equalize across dimensions for P1 only. The evolution is not necessarily smooth, as swift changes can be observed between consecutive days (Day 1 \rightarrow Day 2). The final distributions themselves differ markedly between the two participants. Most striking is the difference in y and $roll$ access on Days 4 and 5. Recall that the task itself is identical for each participant, and the workspace is obstacle-free. While it is possible that the paths planned by each participant would differ even under perfect control execution, most likely spurious movements are happening as the participants learn the control mapping and interface operation. (This is further supported by the differences in the end-effector trajectories depicted in Figure 4.)

From Figure 2, we can also observe that the evolution of access asymmetries follows a distinct pattern for P1 and P2. P1 tends to reduce access asymmetries within a given dimension, as access of positive and negative commands becomes more balanced for all dimensions, as early as Days 2 and 3. However, for P2, only some control dimensions become more balanced over time (x, z), while others maintain asymmetry (all rotational dimensions) or become more asymmetric (y). In addition, the direction of the bias (positive versus negative) is not always consistent.

We further examine the distribution of directional access for each control dimension, along with the command magnitude, in Figure 3. Between Days 1 and 5, the histogram supports generally widen, and, with the exception of $roll$, each control dimension, furthermore, exhibits an increase in variance between Days 1 and 5. Each of these trends are visible in the box plots and the changes in the first- and second-order statistics (mean and variance) of the observed commands between days.

Task Performance. To better visualize human learning over the multiple study sessions, we plot the robot end-effector position on Days 1 and 5 in Figure 4. In general, for both participants, movements become more successful in reaching (or reaching closer to) the target, more directed (closer to the shortest path), and temporally front-loaded, with the bulk of the distance traveled occurring early. Much of the execution time is spent either in late-execution recovery or in the achievement of final orientation, which requires fine-motor commands (e.g., target 5 on Day 5, P1). P2 sometimes reaches near targets early, and then falls into traps of recovery because of P2’s difficulty with issuing $+y$ (needed for targets 2 and 4) and $-z$ (needed for target 1) commands. Each par-




Figure 2: Five-day evolution of proportion of time spent, in each control dimension, for participants P1 (top) and P2 (bottom) performing the 3D-star task. (Zero commands not included.)

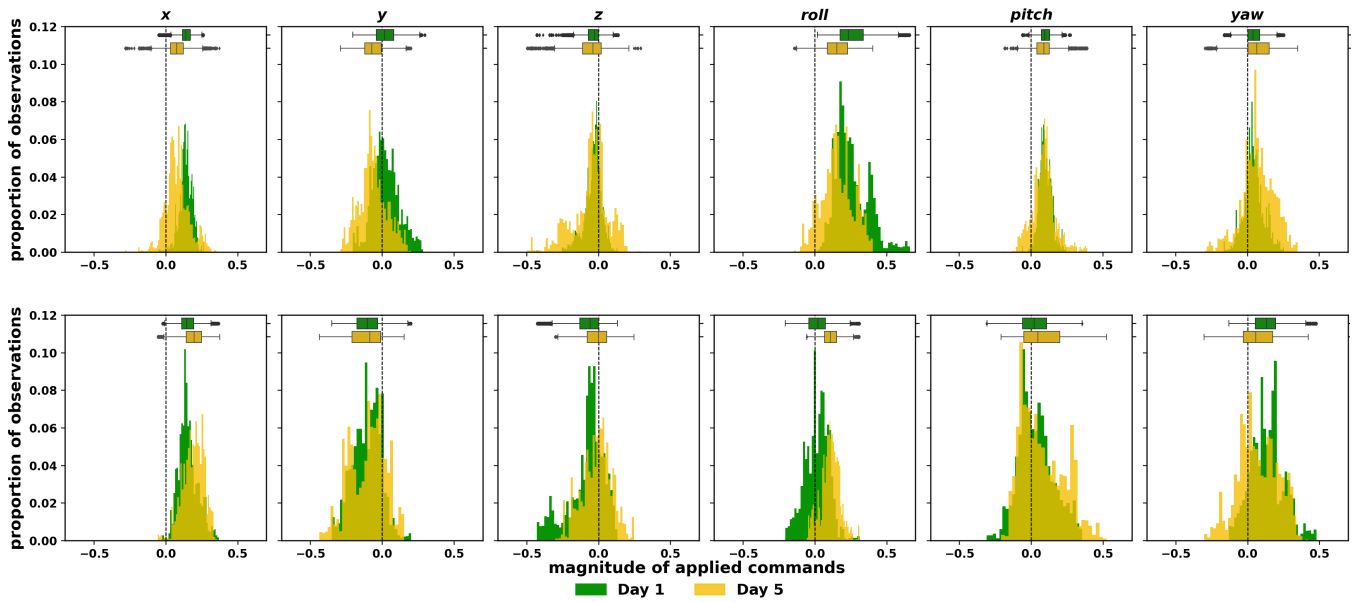


Figure 3: First and last day comparison of histograms of observed robot commands within each control dimension, as each participant (P1 top, P2 bottom) executes the 3D-star task. Standard box plots of robot commands are presented above each respective histogram.

participant, furthermore, exhibits an increase in the number of control commands (a decrease in zero commands) between the first and last days, observed in Figure 4 and from the raw counts of robot commands (not shown). There is also generally a trend of increasing variance across control dimensions between Days 1 and 5; whether this trend is a sign of increased control access and learning or from spurious movements is presently unclear, however.

Opportunities for Robotics Intelligence. Learning to con-

trol complex robots using novel high-DoF interfaces presents many challenges that can be mitigated with the support of robotics intelligence that is designed to be aware and adaptive to the user. Such intelligence might compensate for characteristics of control asymmetries or deficiencies, inconsistencies in time spent across control dimensions, or short- and long-term interface learning, for instance.

Although not presented in this short paper, subjective feedback gathered via questionnaire indicates that the robot

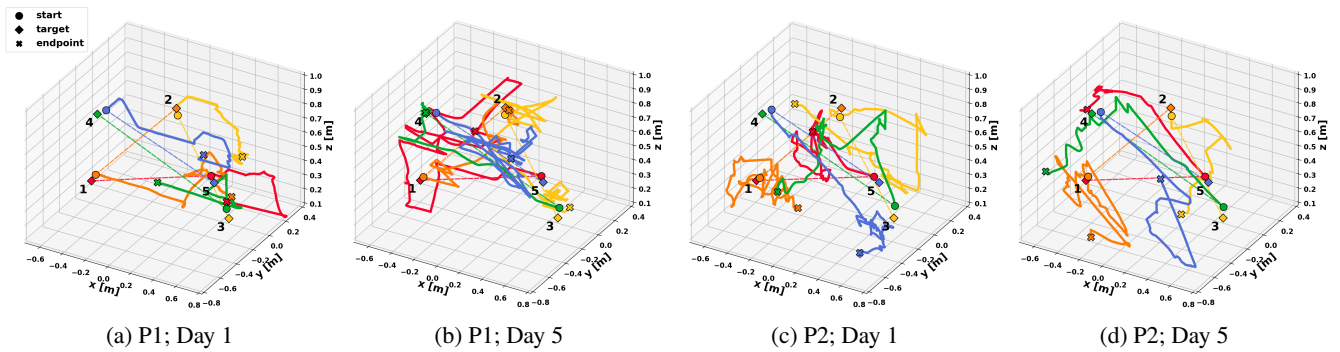


Figure 4: Trajectory plots of robot end-effector position during the 3D-star task on Days 1 and 5, for both participants. The task consists of reaching to five different targets (◆) in succession. Start (●) and end (✕) points for each reach, and the straight-line path between them (dotted line), are shown. Each target, start, end, and straight-line path for a single reach are the same color.

control is unintuitive at times. There are instances when participants feel uncertain about how to move the robot in certain dimensions, despite having become familiar with the dynamics of the robot, and other instances where slight differences in a participant’s movements lead to the robot moving in unexpected directions. As a result, we observe participants regularly issuing unintended commands through the interface—either by moving in the undesired direction of an intended control dimension or activating an unintended dimension altogether. The result is time spent attempting corrections and recovery instead of progressing towards task goals. The use of interface-aware autonomy (Gopinath, Nejati-Javaremi, and Argall 2021) that infers about and prevents these unintended commands in a shared-control framework could not only prevent the subsequent need for corrective action, but can also be used within a training and rehabilitation framework to aid in learning to provide control commands through the interface.

Conclusion and Future Work

In this short paper, we presented preliminary results from a study with two uninjured participants in which they controlled a high-DoF robotic arm, using limited upper body movements, to perform a variety of reaching tasks. We presented some key insights on the typical control asymmetries that arise as well as observations on human learning in the context of high-DoF robot control. We also identified intervention opportunities for robotics autonomy.

In the future, we will use the data and insights collected to inform the development of an assistive autonomy paradigm. The role of the paradigm will be to facilitate the user’s learning of the interface, while adapting to the user’s improvement and compensating for any deficits in control. We plan to use the results from this study and the developed autonomy paradigm to conduct a long-term user study, where participants with spinal cord injury evaluate the efficacy of this assistance paradigm.

Acknowledgements

Research reported in this publication was supported by the National Institute of Biomedical Imaging and Bioengi-

neering (NIBIB) under award number R01-EB024058; 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 Skłodowska-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 National Institutes of Health.

References

- Casadio, M.; Ranganathan, R.; and Mussa-Ivaldi, F. A. 2012. The body-machine interface: a new perspective on an old theme. *Journal of Motor behavior* 44(6): 419–433.
- Chau, S.; Aspelund, S.; Mukherjee, R.; Lee, M.-H.; Ranganathan, R.; and Kagerer, F. 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)*, 3877–3882. IEEE.
- Fasola, J.; and Mataric, M. J. 2012. Using Socially Assistive Human-Robot Interaction to Motivate Physical Exercise for Older Adults. *Proceedings of the IEEE* 100(8): 2512–2526. ISSN 0018-9219. doi:10.1109/JPROC.2012.2200539.
- Gopinath, D.; Jain, S.; and Argall, B. D. 2016. Human-in-the-loop optimization of shared autonomy in assistive robotics. *IEEE Robotics and Automation Letters* 2(1): 247–254.
- Gopinath, D.; Nejati-Javaremi, M.; and Argall, B. D. 2021. Customized Handling of Unintended Interface Operation in Assistive Robots. *IEEE International Conference on Robotics and Automation (ICRA)*.
- Ison, M.; Vujaklija, I.; Whitsell, B.; Farina, D.; and Artemiadis, P. 2015. High-density electromyography and motor skill learning for robust long-term control of a 7-DoF robot arm. *IEEE Transactions on Neural Systems and Rehabilitation Engineering* 24(4): 424–433.

Jain, S.; Farshchiansadegh, A.; Broad, A.; Abdollahi, F.; Mussa-Ivaldi, F.; and Argall, B. 2015. Assistive robotic manipulation through shared autonomy and a body-machine interface. In *2015 IEEE International Conference on Rehabilitation Robotics (ICORR)*, 526–531. IEEE.

Kramer, M. A. 1991. Nonlinear principal component analysis using autoassociative neural networks. *AIChE journal* 37(2): 233–243.

Langer, A.; Feingold-Polak, R.; Mueller, O.; Kellmeyer, P.; and Levy-Tzedek, S. 2019. Trust in socially assistive robots: Considerations for use in rehabilitation. *Neuroscience and biobehavioral reviews* 104: 231–239. ISSN 0149-7634. doi:10.1016/j.neubiorev.2019.07.014.

Losey, D. P.; Srinivasan, K.; Mandlekar, A.; Garg, A.; and Sadigh, D. 2020. Controlling assistive robots with learned latent actions. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*, 378–384. IEEE.

Macchini, M.; Schiano, F.; and Floreano, D. 2019. Personalized telerobotics by fast machine learning of body-machine interfaces. *IEEE Robotics and Automation Letters* 5(1): 179–186.

Miehlbradt, J.; Cherpillod, A.; Mintchev, S.; Coscia, M.; Artoni, F.; Floreano, D.; and Micera, S. 2018. Data-driven body-machine interface for the accurate control of drones. *Proceedings of the National Academy of Sciences* 115(31): 7913–7918.

Muelling, K.; Venkatraman, A.; Valois, J.-S.; Downey, J. E.; Weiss, J.; Javdani, S.; Hebert, M.; Schwartz, A. B.; Collinger, J. L.; and Bagnell, J. A. 2017. Autonomy infused teleoperation with application to brain computer interface controlled manipulation. *Autonomous robots* 41(6): 1401–1422. ISSN 0929-5593. doi:10.1007/s10514-017-9622-4.

Pierella, C.; Sciacchitano, A.; Farshchiansadegh, A.; Casadio, M.; and Mussa-Ivaldi, S. 2018. Linear vs non-linear mapping in a body machine interface based on electromyographic signals. In *2018 7th IEEE International Conference on Biomedical Robotics and Biomechatronics (Biorob)*, 162–166. IEEE.

Ranganathan, R.; Lee, M.-H.; Padmanabhan, M. R.; Aspelund, S.; Kagerer, F. A.; and Mukherjee, R. 2019. Age-dependent differences in learning to control a robot arm using a body-machine interface. *Scientific reports* 9(1): 1–9.

Rizzoglio, F.; Casadio, M.; De Santis, D.; and Mussa-Ivaldi, F. A. 2021. Building an adaptive interface via unsupervised tracking of latent manifolds. *Neural Networks* 137: 174–187.

Thorp, E. B.; Abdollahi, F.; Chen, D.; Farshchiansadegh, A.; Lee, M.-H.; Pedersen, J. P.; Pierella, C.; Roth, E. J.; González, I. S.; and Mussa-Ivaldi, F. A. 2015. Upper body-based power wheelchair control interface for individuals with tetraplegia. *IEEE Transactions on Neural Systems and Rehabilitation Engineering* 24(2): 249–260.

Wold, S.; Esbensen, K.; and Geladi, P. 1987. Principal component analysis. *Chemometrics and Intelligent Laboratory Systems* 2(1-3): 37–52.