Robot Learning to Switch Control Modes for Assistive Teleoperation

Siddarth Jain and Brenna Argall
Northwestern University, Evanston, IL, USA 60208
Rehabilitation Institute of Chicago, Chicago, IL, USA 60611
Email: sidd@u.northwestern.edu, brenna.argall@northwestern.edu

I. Introduction

Assistive robotic arms hold great potential to assist individuals with motor impairments in performing pick-and-place tasks, object retrieval tasks, or even assist with personal hygiene and feeding. There are about 150,000 people in the United States who could benefit from using such assistive technology [4]. However, robotic arms typically have a much higher number of controllable degrees of freedom than their control interfaces, which makes teleoperation challenging. Also, the more severe the motor impairment, the more limited the control interface that the user can operate, and the greater is the difficulty to teleoperate the assistive arm.

Over a couple of decades, several assistive robotic arms have been developed and evaluated but only few of these have been commercialized. Their operation involves direct teleoperation using a traditional control interface like 2- or 3-axis joystick, which requires the user to switch between one of several control modes (*mode switching*). Mode switching partitions the controllable degrees of freedom of the robot such that each control mode maps the input dimensions to a subset of the arm's controllable degrees of freedom. Performing even simple manipulation tasks can involve multiple mode switches and the process can become extremely challenging with the more limited interfaces like a sip-and-puff or a head array that are available to individuals with severe motor impairments.

The most commonly used commercial assistive robotic arm is the JACO (Kinova Robotics, Canada). The JACO has six degrees of freedom (6-DoF) and is equipped with a multi fingered gripper. The arm is controlled by direct teleoperation using a 3-axis joystick and a mode switching paradigm in which the user can either control the position or the orientation of the robot hand in 3-D and can also control the grasp and release action of the gripper. The three control modes (Figure 1) are selected by the user using the pushbuttons on the joystick. Moreover, operating the 3-axis joystick involves twisting the shaft which requires fine motor control in the fingers and thus can become limiting to some of the targeted population. For an alternate control scheme that does not require twisting the shaft (2-axis joystick), now the required number of control modes increases to four. Performing tasks involves control of the shaft (for the robot motion control) and the pushbuttons (to switch between control modes) and thus teleoperation can become challenging and tedious.

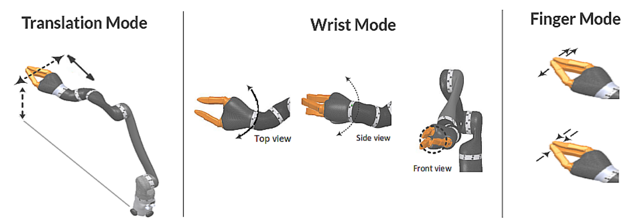

Fig. 1: Jaco arm teleoperation requires the user to switch between 3 control modes when using 3-axis joystick (images from Kinova Jaco arm user guide).

Studies have identified that the numerous and frequent mode switch operations required for performing everyday tasks with such assistive machines make them difficult to operate [3, 6]. Such mode switching operations are slow, non-intuitive and require the user to divert their attention away from accomplishing the task, which results in sustained physical and even cognitive effort from the user.

In this paper we propose a robot learning solution to provide automated mode switching in order to assist in the teleoperation of robotic arms in an efficient manner. We apply machine learning methods directly on robot teleoperation demonstrations to decide when to switch between control modes automatically—eliminating the need for manual mode switches. In our evaluation we also present pilot results of the robot's prediction of automated mode switching on unseen teleoperation data.

II. Background

The teleoperation of assistive robotic arms presents a challenge as it involves controlling more degrees of freedom than the available number of control signals from the user. Researchers have proposed control-sharing solutions to offload some control burden from the human user. The most common solutions involve control blending [2] or control partitioning [7] which keeps the high-level decisions with the human (e.g target selection) and low-level control execution with the robot autonomy. Only a handful of works focus on providing assistance for mode switching to make the teleoperation process easier. Herlant *et al.* [3] propose automatic mode switching by determining a time-optimal policy. Their experiments solve the problem for a 2D simulated robot, and they target to scale their approach to assistive robot arms in future work.

Similar to robotic arm teleoperation, different control modes are used to operate powered myoelectric prostheses [6]. Conventional controllers have the user cycle (using EMG/EEG inputs or hardware switches) through the available functions on the prosthesis, which are unacceptably slow and difficult to use. Pilarski *et al.* [5] propose a reinforcement learning approach to predict control switching functions using EMG signals and robot state-information (e.g. joint angles).

Our work targets robot learning for automated control mode switching in the case of assistive robot arms and we propose to directly learn from the teleoperation data with a supervised learning approach. Our approach enables the online prediction and automated switching of control modes for assistive teleoperation, is appropriate for higher dimensional assistive robot arms and operates in real-time. Our aim is to reduce the effort of users controlling high-dimensional robot arms.

III. Approach

The goal of our research is to enable automated mode switching for assistive robot teleoperation. We gather demonstration data from users teleoperating a robotic arm, and use machine learning to capture when users change the control modes. The robot can predict the ideal control mode and perform automated mode switching for teleoperation.

A. Robot Learning Formulation

We model the mode switching problem as a classification problem using a supervised learning approach. The aim is to train a classifier that can efficiently estimate the control mode y given the current state of the robot. Our learned classifier accepts an input vector of features **x** and returns a single scalar label y, which is the predicted control mode for the robot.

We consider the case of 3-axis joystick operation for the robot teleoperation (Cartesian control in 6D) and focus on the prediction of two control modes: *Translational Mode (TM)* and the *Wrist Mode (TM)* (Figure 1), as these are the modes that require frequent back-and-forth switching when performing tasks with the robot.[1] The label for the classifier y represent the translational (0) or the wrist mode (1). We define a two dimensional input feature-vector **x**, where x_1 is the position difference (Euclidean distance) in 3-D between the end-effector position of the robot and the goal location and x_2 is a scalar metric which is a measure of the orientation alignment (computed using quaternion difference) between the current end-effector orientation and the goal pose for a task. We assume the goal pose is known (e.g. through a combination of machine perception and human intent inference).

Provided n samples of the data [$x^{(i)}$, $y^{(i)}$], we formulate a classifier to learn the function $f(x)$ that can predict the label y, which is the desired automated prediction of the control mode for assistive teleoperation. We learn a Support Vector Machine (SVM) classifier, which is able to find non-linear separators and is computationally efficient.

[1]Finger Mode switching will be implemented in our future work, for example based on the proximity to the desired goal.

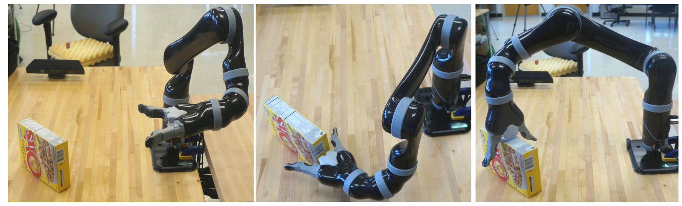

Fig. 2: Data collection Tasks. Left: Start configuration. Middle: Task1 (side grasp on box). Right: Task2 (top grasp on box).

B. Data Collection: Teleoperation Demonstrations

We demonstrate our approach on the MICO robotic arm (Kinova Robotics, Canada), a 6-DoF manipulator with a two finger gripper (Figure 2). The MICO is the research edition of the JACO arm which is used commercially within assistive domains. It is customarily controlled using the same control interface and scheme as the JACO arm (3-axis joystick operation with 3 control modes, as shown in Figure 1).

We collected teleoperation demonstration data from two users (1M, 1F). Each user teleoperated the MICO robot (mounted on a table) using the 3-axis joystick and performed two tasks (reach two different target positions in the environment). The two target positions were a side grasp pose (S) and a top grasp pose (T) on a box placed on the table (Figure 2). The robot started from the same fixed initial configuration for each run. Target S was easier (required fewer mode switches) to reach as compared to T. The robot end-effector pose and the targets were represented in 6D (position + orientation).

Each user performed 10 demonstrations for each task by teleoperating the robot. At any instant during the teleoperation, we recorded the 6D Cartesian trajectory ξ of the end-effector pose, and the current control mode $y^{(i)}$ selected by the user and computed the feature-vector $x^{(i)}$.

We plot the data (Figure 3) in the feature space x_1, x_2 and denote the labels by color codes (green for Translational Mode and blue for Wrist Mode). From the plots (Figure 3) it is observed that each user adopts a different teleoperation strategy to perform the same task, and the same user also differs in his/her teleoperation strategy for performing different tasks. For example, in Task 1, the first user operated in TM to move near the object, switched to WM to align the end-effector and then again to TM, whereas the second user adopted a reverse strategy.

C. Experiments and Results

To evaluate performance, we tested the accuracy of our approach to predict control modes on unseen teleoperation demonstrations and evaluated the performance of the learned models within and across tasks and users. An SVM classifier with an RBF kernel (C=1 and γ=10) using LibSVM [1] was used to learn the classifier models.

To evaluate task-specific performance, for each user and for each task we split the collected demonstrations into a training set (6 demonstrations) and test set (4 demonstrations). The training sets were used to learn the classifier for each user and task and the learned models were used to predict the label y on the test set demonstrations (Table I).

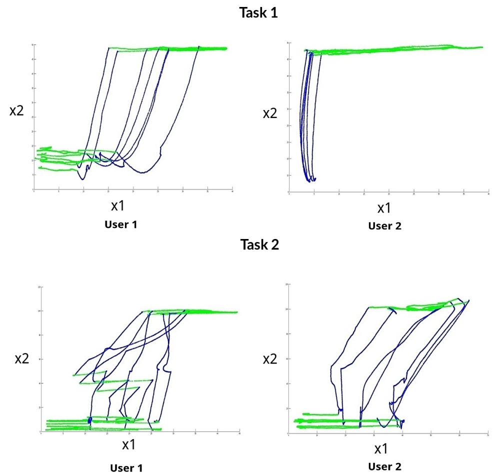

Fig. 3: Training data visualization in the feature space. TM (green) and WM (blue).

TABLE I: Control Mode Prediction

	User1		User2	
	Task1	Task2	Task1	Task2
No. of Samples	8188	12563	10231	17434
Accuracy	84.4%	82.12%	89.60%	84.26%

To test the generalization of the approach across tasks, we created a second training set (12 demonstrations) and test set (8 demonstrations) for each user using data from both tasks, and learned classifier models on the training set. Table II represents the results of the robot learning to predict the control modes on the test datasets for each user for both tasks.

TABLE II: Generalization Across Tasks

Test Set	User1	User2
No. of Samples	20751	27665
Accuracy	80.82%	82.87%

To test the generalization across users, we tested the user-specific models learned in Table II on the other user's test set. We also created a training set (24 demonstrations) that combined the training set data from both users and learned a single classifier model to test on individual user's test sets (8 demonstrations each). These results are presented in Table III.

TABLE III: Generalization Across Users

Model	User-Specific		User-Generic	
Test Set	User1	User2	User1	User2
No. of Samples	20751	27665	20751	27665
Accuracy	71.77%	61.24%	77.15%	82.54%

It is observed that the model specific to a user drops in performance when tested on another user, evident from the differences in the teleoperation strategies. However, the performance is improved when a user-generic model is tested. This demonstrates the effectiveness of the approach to perform across users but indeed the dataset is required to be significantly expanded before any concrete conclusion can be made about generalization of the learned models.

Lastly, a visualization of the automated mode switching prediction as compared to the user teleoperation for a Cartesian trajectory is shown in Figure 4 (disagreements shown in red).

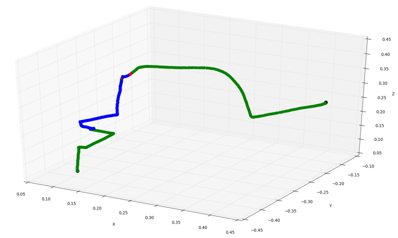

Fig. 4: Automated mode switching for a Cartesian trajectory. TM (green), WM (blue), Wrong predictions (red).

IV. Conclusions and Discussion

We have presented an approach for automated mode switching using robot learning. Our pilot results indicate that the robot is able to predict the correct control modes with the proposed features and learning formulation. We also demonstrate that our approach is able to handle mode switching for multiple tasks operating within a 6-D control space. It is observed that each user adopts a different teleoperation strategy and in our future work we will explore the categorization of users based on their teleoperation strategy and learn a single model per category, to see if such models generalize across users. We also will implement automated mode switching on the robot to examine how users react to the robot performing mode switches during the teleoperation and also to wrong control mode predictions. Our future work will extend the approach to work with 2-axis joystick control (4 control modes) and an evaluation with a larger user study that will involve subjects with motor impairments and richer teleoperation tasks.

Acknowledgments

This material is based upon work supported by the National Science Foundation under Grant CNS 15544741.

References

[1] Chih-Chung Chang and Chih-Jen Lin. LIBSVM: a library for support vector machines. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 2(3):27, 2011.

[2] Anca Dragan and Siddhartha Srinivasa. Formalizing assistive teleoperation. In *Proceedings of RSS*, July 2012.

[3] Laura V Herlant, Rachel M Holladay, and Siddhartha S Srinivasa. Assistive teleoperation of robot arms via automatic time-optimal mode switching. In *Proceedings of the ACM/IEEE HRI*, 2016.

[4] Isabelle Laffont, Nicolas Biard, Gérard Chalubert, Laurent Delahoche, Bruno Marhic, François C Boyer, and Christophe Leroux. Evaluation of a graphic interface to control a robotic grasping arm: a multicenter study. *Archives of physical medicine and rehabilitation*, 90(10):1740–1748, 2009.

[5] Patrick M Pilarski, Michael R Dawson, Thomas Degris, Jason P Carey, and Richard S Sutton. Dynamic switching and real-time machine learning for improved human control of assistive biomedical robots. In *Proccedings of the IEEE RAS & EMBS BioRob*, 2012.

[6] Erik Scheme and Kevin Englehart. Electromyogram pattern recognition for control of powered upper-limb prostheses: State of the art and challenges for clinical use. *Journal of rehabilitation research and development*, 48(6):643, 2011.

[7] Katherine M Tsui, Dae-Jin Kim, Aman Behal, David Kontak, and Holly A Yanco. I want that: Human-in-the-loop control of a wheelchair-mounted robotic arm. *Applied Bionics and Biomechanics*, 8(1):127–147, 2011.