Shared Control in Modal Teleoperation

Laura Herlant Carnegie Mellon University, *herlant@cmu.edu*
Ben Weinstein-Raun Virginia Tech, *bwr@vt.edu*
Siddhartha Srinivasa Carnegie Mellon University, *siddh@cs.cmu.edu*

I. Introduction

Assistive devices, like powered wheelchairs and other mobility aids, have greatly enhanced the quality of life for individuals with disabilities. While assistive robotic arms increase the accessibility of physical manipulation tasks, they are difficult to teleoperate effectively. When controlling an arm with many degrees of freedom using a lower-dimensional input, even simple tasks can become tedious [1] and cognitively burdensome [2]. A common method for dealing with the mismatch in dimensionality is through modal control, in which the user selects one of several control modes, each corresponding to some mapping from input dimensions to a subset of the arm's degrees of freedom. This incurs added complexities to predicting users' intentions, but also introduces interesting opportunities for assistance.

In this paper, we will describe how users typically employ assistive robotic arms under modal control, discuss the results of interviews with users to determine their opinions on this method of controlling the arm, propose some possible assistance techniques in this domain, and lay out our plan to test the effectiveness of each technique.

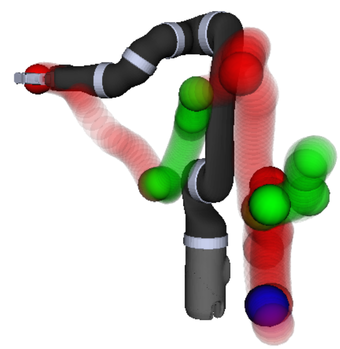

(a) Visualization of robot arm trajectory

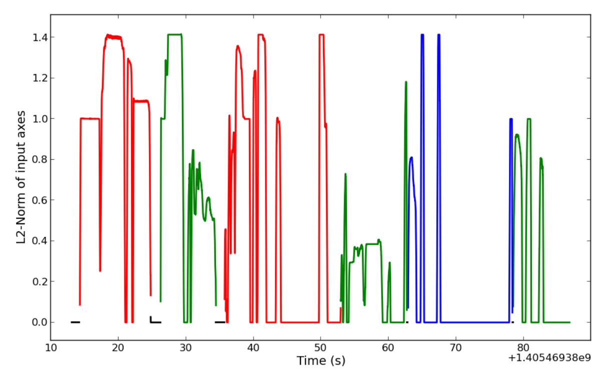

(b) Time spent in each mode

Fig. 1: A recorded trajectory showing modal control on an arm similar to the JACO: x-y-z is red, ψ-θ-ϕ is green, and gripper is blue

II. Modal Control

We will discuss what we mean by the term *modal control*. When mapping from a low-dimensional input onto a higher-DOF robotic arm, in order to get a full specification of the motion, the extra DOFs must be accounted for. The naïve solution is to control each DOF independently. Extending this, we can give the user control of some subset of the configuration space. In both scenarios, users employ a button or signal to cycle through available mappings from input- to output-space. Because the user is choosing between these different "modes," we call this "modal control."

The number of modes required grows quickly for complicated robots. Even if the control input is specified by the Cartesian location and orientation of the end effector, 6 DOFs are needed. Many such arms will require an additional DOF to close and open the end manipulator. For a 2-DOF input, such as the joystick that might be used to control an electric wheelchair, this requires at least 4 different control modes, and some method for signaling mode transitions. One example partition of the space would be to have an X-Y mode, a Z-roll mode, a pitch-yaw mode, and a grasp-release mode. There exist several possible mode partitions, and some may be better or worse for certain applications.

III. User-Centric Design

Motivated by the high rejection rate among assistive technologies, we employ a user-centric design wherein we interviewed users of the JACO assistive robot arm to assess their needs. Studies are often limited to asking users about potential experiences and future abilities [6], [5], without giving these users the opportunity to work with the technology for an extended period of time. The JACO has been commercially released and is currently used by over 150 people. There is thus a user population rich in real-world experiences controlling a robotic arm.

We contacted English speakers who are currently living with the JACO arm, and asked them about their general experience with the arm, what they use it for, and what they would like to use it for. We also interviewed employees of Kinova, the makers of the JACO. Every use has been trained by these employees, and they have considered many features in the JACO's interface design. Talking with employees also gave us some exposure to the larger user population than did the limited number of users we spoke with.

Through the interviews, we determined that speed is a governing factor in determining the success of a feature: even if it is feasible to use the arm to eat from a plate, if it takes an hour for each bite, this is not a practical or acceptable solution and will not be used in real life. Conversely, certain tasks require varying types and degrees of precision. Moving a cup of coffee without spilling requires one kind of precision, while waving to a friend, a highly personalized motion, requires another.

All of the users expressed difficulty concerning the mode switching process. The number of small adjustments needed while controlling the arm requires switching modes often. Keeping track of what control mode the arm is in requires noticeable cognitive effort, and even pressing a button to change modes is difficult or impossible for some users. Based on this finding, we are investigating assistance techniques that can be used to reduce the occurrence of mode switching.

IV. Modal Assistance Challenges

We apply policy blending as a way to keep the user in control while still providing assistance. Policy blending is a general framework that can describe several experiments looking at shared control. A detailed table of how related work fit into this framework can be found in [3]. Policy blending [4] mixes the optimal action under the robot's policy P with the optimal action under the user's policy U at every time step via an arbitration function α, which is determined by user preference and scales with the robot's confidence in its prediction as shown in (1). This introduces some difficulties when applied to modal control.

$$(1-\alpha)U + \alpha P \tag{1}$$

With modal control, the actions from the user's policy and the robot's policy have different dimensions. There are two clear ways of handling this mismatch: (a) limit the robot's policy P to taking actions only in the user's current control mode, or (b) allow the robot's policy to span the full action space and expand the user's action in the non-controlled dimensions.

Another assistance technique, which can be implemented alongside either policy-blending options (a) or (b), would be to have the robot predict when the user is likely to change control modes, and then perform the switch for the user.

All of these assistance techniques are sensitive to the quality of the predictions made about the user's goals. If the robot's prediction is incorrect but has a high confidence value, the human will not be able to overpower the "assistive" action prescribed by the robot's policy. This will occur with higher probability in control method (b), wherein the robot can exert its predicted intent on more dimensions than are available for the user to control.

The difficulties with modal assistance that we would like to explore can be summarized in the following research questions:

1) Do users prefer to have the robot assist in all dimensions or only those which the user also controls?
2) How do users react to the robot performing mode switching?
3) How does a wrong prediction effect the users' preference for type of assistance?

V. Research Plan

To implement the assistance methods outlined in IV, we need a means of predicting the user's goal, and a policy directing the robot to an arbitrary goal. Each of the assistance methods and failure cases will be applied in a within-subjects user study.

A. Learning the User's Policy for Prediction

We have created a testbed with the intention of first modeling users' motion without any mode switching. For this setup we needed to constrain the task in such a way that all inputs can be specified using a 3-axis joystick. We will ask users to control a robot arm to press buttons, each at a different x, y, and z location relative to the robot arm. The orientation of the hand is defined ahead of time and remains the same throughout the trials, and the user has control over only the x,y,and z velocities. Some interesting questions include how many axes are used simultaneously, and how much the trajectories differ from a canonical straight-line path.

Once the user's policy within a mode is modeled, we will run another experiment requiring two control modes. We will examine the cost of switching modes, and how mode switching modifies the user's policy. With a policy describing the user's actions, we can perform prediction by determining which goal minimizes the cost of the user's current trajectory [8].

Given a predicted goal, the robot can employ any motion planning algorithm to acquire a path to the predicted goal. The ideal policy is that of the user if he or she had full control over all degrees of the system at all times. Since we do not have direct access to this policy, a cost function can be constructed from the lower-dimensional user trajectories via maximum entropy inverse optimal control [7]. The robot's policy can then be determined from the cost function in a full-dimensional action space.

B. Modal Assistance Experiments

We will run a within-subjects study. The two independent measures are the assistance type and whether the robot succeeds or fails in its prediction. The objective dependent measures will be the time required for task completion, the magnitude of the difference between the actions of the user and those of the robot, and the number of mode switches. The subjective dependent measures will be a self-reported survey afterward in which the user will be asked about their preference for assistance given successful and unsuccessful prediction, and other qualitative responses about the assistance methods.

VI. Conclusion

Shared control offers an opportunity to increase the efficiency of task performance when the user's input is limited or unreliable in some way. Due to the user-oriented nature of shared control, the different techniques for assistance need to ultimately be evaluated by users for their preference as well as by performance metrics. We will test two policy-blending techniques, under successful and failed prediction, and evaluate them both objectively and subjectively.

References

[1] J. Y. C. Chen, E. C. Haas, and M. J. Barnes. Human performance issues and user interface design for teleoperated robots. *Systems, Man, and Cybernetics, Part C: Applications and Reviews, IEEE Transactions on*, 37(6):1231–1245, 2007.

[2] J. W. Crandall and M. A. Goodrich. Characterizing efficiency of human robot interaction: A case study of shared-control teleoperation. In *Intelligent Robots and Systems, 2002. IEEE/RSJ International Conference on*, volume 2, pages 1290–1295. IEEE, 2002.

[3] A. Dragan and S. Srinivasa. Formalizing assistive teleoperation. In *Robotics: Science and Systems*, July 2012.

[4] Anca D Dragan and Siddhartha S Srinivasa. A policy-blending formalism for shared control. *The International Journal of Robotics Research*, 32(7):790–805, 2013.

[5] Michael Hillman, Karen Hagan, Sean Hagen, Jill Jepson, and Roger Orpwood. The weston wheelchair mounted assistive robot - the design story. *Robotica*, 20:125–132, 2002.

[6] S. D. Prior. An electric wheelchair mounted robotic arm-a survey of potential users. *Journal of medical engineering & technology*, 14(4):143–154, 1990.

[7] Brian D Ziebart, Andrew L Maas, J Andrew Bagnell, and Anind K Dey. Maximum entropy inverse reinforcement learning. In *AAAI*, pages 1433–1438, 2008.

[8] Brian D Ziebart, Nathan Ratliff, Garratt Gallagher, Christoph Mertz, Kevin Peterson, James A Bagnell, Martial Hebert, Anind K Dey, and Siddhartha Srinivasa. Planning-based prediction for pedestrians. In *Intelligent Robots and Systems, 2009. IROS 2009. IEEE/RSJ International Conference on*, pages 3931–3936. IEEE, 2009.