

Uncovering People’s Preferences for Robot Autonomy in Assistive Teleoperation

Maggie Collier¹ and Henny Admoni²

Abstract—What factors influence people’s preferences for robot assistance? Answering this question can help roboticists formalize assistance that leads to higher user satisfaction and increased user acceptance of assistive technology. Often in assistive robotics literature, we see paradigms that aim to optimize task success metrics or measures of users’ perceived task complexity and cognitive load. However, frequently in this literature, participants express a preference for paradigms that do not perform optimally with respect to these metrics. Therefore, task success and cognitive load metrics alone do not encapsulate all of the factors that inform users’ needs or desires for robotic assistance. We focus on a subset of assistance paradigms for manipulation called assistive teleoperation in which the system combines control signals from the user and the automated assistance. In this work, we aim to study potential factors that influence users’ preferences for assistance during object manipulation tasks. We design a study to evaluate two factors (magnitude of end effector movement and the degrees of freedom being controlled) that may influence the amount of automated assistance the user wants.

I. OVERVIEW

Assistive robots can enable people to perform some activities of daily living independently. However, teleoperated assistive robots can be challenging for operators to control, especially when the robot has a higher number of degrees of freedom (DOFs) than the input interface used for teleoperation. Fortunately, automation can help, though having *full* automation is not ideal because operators often want to maintain their sense of control over the system [1].

One way to provide automated assistance to people while maintaining some of their control is via an assistive teleoperation method such as shared control. Shared control combines a user’s control signal with assistive policy execution in a process called arbitration. In this work, we study people’s true desire for robot assistance by putting them in control of arbitration and measuring the changes they make. Said differently, we can study users’ preferences for autonomous assistance by enabling them to set the amount of assistance they want at any point throughout a task (see Fig. 1).

The underlying assumption with many existing assistive teleoperation paradigms is that users will prefer more autonomous assistance if it improves task performance metrics like task completion time, success rate, number of mode switches, etc. However, sometimes users prefer assistance that is not optimal, or they prefer the challenge of teleoperation to losing their sense of control when assistance is

applied [1], [2], [3]. For example, when users were enabled to customize the arbitration function in a shared control study, they chose functions that did not perform optimally with respect to task performance metrics, even after experiencing the optimal arbitration function [3]. This finding, among others, underscores the need to study people’s preferences for autonomous assistance rather than assuming their preferences align with our limited definitions of optimality.

Within the domain of assistive teleoperation, user preferences have been studied in the context of learning a mapping between input interface actions and robot actions [4], [5] or trajectories to a goal location [6]. In contrast, our work is about users’ preferences for the *amount of autonomous assistance* throughout a task. Instead of comparing pre-defined arbitration functions like in past work [1], [2], we observe the users’ preferences directly by letting them control the arbitration level at any point during a task.

Further, we predict that certain features common to many object manipulation tasks may change people’s preferences for autonomous assistance. In the example of grasping, large movements of the end effector toward the goal object may be less complex for users than finer movements required to reach a precise pose near the object. This difference in **magnitude of movement** may influence the amount of assistance people want. Additionally, the DOFs of end effector movement being controlled at a given time, which we refer to as **active degrees of freedom**, may influence user preferences for assistance. For example, users may prefer more assistance when controlling the robot in rotational DOFs than in translational DOFs because they perceive rotations as more challenging than translations [7].

In this work:

- We propose a user study involving the first shared control paradigm that lets users directly control the arbitration at any point during a task.
- We investigate whether user’s preferences for assistance change throughout object manipulation tasks based on magnitude of movement and active DOFs.

II. APPROACH

Many shared control paradigms arbitrate between user and robot control using a parameter α , such that

$$a_{out} = \alpha \cdot a_{robot} + (1 - \alpha) \cdot a_{user} \quad (1)$$

where a_{out} is the output command, a_{robot} is the autonomous command, and a_{user} is the user’s command. An $\alpha = 0$ corresponds to teleoperation and $\alpha = 1$ to full autonomy.

¹Maggie Collier is with the Robotics Institute at Carnegie Mellon University macollie@andrew.cmu.edu

²Henny Admoni is with the Robotics Institute at Carnegie Mellon University henny@cmu.edu

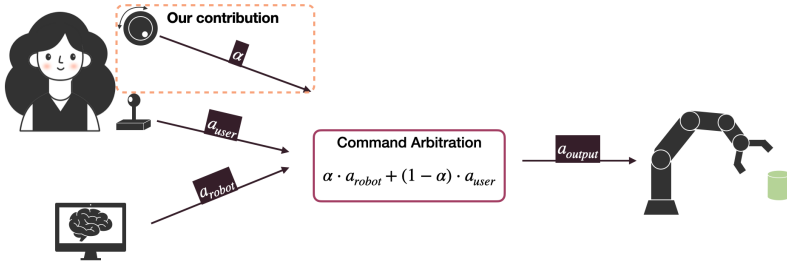


Fig. 1. Many assistive teleoperation paradigms work by arbitrating between input commands from the user and commands from the assistive system. In this work, we enable users to set the value of the arbitration parameter α at any point during a task using a dial. This framework enables us to uncover users’ preferences for autonomous assistance throughout tasks.




Fig. 2. Our approach enables users to control the way their commands from the joystick are combined with the assistive commands. We also show setup for the trajectory-following task for our second experiment.

Many past shared control approaches compute α based on a notion of the system’s confidence in the inferred goal [2]. As confidence increases, so does the value of α , leading to more autonomous assistance from the system. Importantly, user preferences are not considered for computing α .

In contrast, our work gives users direct access to α . Unlike in [3] in which users could change the arbitration function in between trials, we enable users to set the value of α with a dial interface at any point *throughout* the task (see Fig. 2).

III. EXPERIMENT PLAN

Our study has three major components: an assistance sensitivity test, an experiment on magnitude of movement, and an experiment on active DOFs.

A. Assistance Sensitivity Test

We first want to evaluate to what extent participants perceive differences in α . Participants will undergo an assistance sensitivity test in which they perform a series of grasping actions under shared control using different assistance values. For each pair of grasps, participants will be asked to report which one had more assistance, allowing us to calculate the “just noticeable difference” level for α .

B. Experiment on Magnitude of Movement

We expect that a difference in magnitude of movements (gross vs. fine) has an effect on users’ assistance preferences, and that these preferences may also be conditioned on the overall complexity of the task. This expectation comes from prior work in which users preferred more aggressive assistance during a high complexity task [2]. We design higher and lower complexity tasks by manipulating the shape of the goal object and the risk involved in the task.

Lower Complexity Grasping Task: Users will have to grasp a tall can with the robot arm. To decrease the risk, they will be allowed to attempt the task twice.

Higher Complexity Grasping Task: Users will have to grasp a very small cylinder (with height roughly equal to the width of the robot’s fingers). To increase the risk, participants will only have one attempt to complete the task.

H1: *In tasks of higher complexity, users will prefer a relative increase in assistance when switching*

from gross movements to finer movements, but in tasks of lower complexity, users will not show a consistent trend in the relative change in assistance when switching from gross to fine movements.

C. Experiment on Active DOFs

We will also investigate the effects of active DOFs on users’ preference for assistance. We predict that users will prefer more aggressive assistance when moving in rotational DOFs than in translational DOFs.

Trajectory-Following Task: Participants will have to move the robot’s end effector to follow a 3D trajectory. We use an educational toy which involves moving a looped wand along a metal wire bent into a complex shape (Fig. 2). If the wand touches the metal wire, the toy will beep.

H2: *Users will prefer more automated assistance when moving in rotational DOFs than in translational DOFs.*

D. Metrics

The dependent measure of both experiments is the assistance value α chosen by participants over the course of the task. Additionally, we will collect survey responses from participants at the end of the study asking about their sense of control and potential desire to use the system again. To define the two categories of magnitude of movement, we will take the magnitude of the end effector’s linear velocity and use a Gaussian Mixture Model (GMM) with two components to split the data into *fine* and *gross* groups.

IV. IMPLICATIONS FOR FUTURE WORK

In this work, we identify and investigate two potential features of tasks that may affect users’ preferences for autonomous assistance in an assistive teleoperation setting. These features could be considered in the design of future assistive teleoperation paradigms but also assistive systems more generally. For example, an approach that dynamically changes the α value depending on if the user is moving in translational DOFs or rotational DOFs could easily extend this work. Additionally, the user-controlled arbitration component of our shared control approach is novel and could be used to study additional features that influence assistance

preferences, such as the user's sense of control. We know that people's preferences for assistance do not always align with our limited definitions of optimality. Therefore, this work promotes the discovery of user's preferences for autonomous assistance instead.

REFERENCES

- [1] D.-J. Kim, R. Hazlett-Knudsen, H. Culver-Godfrey, G. Rucks, T. Cunningham, D. Portee, J. Bricout, Z. Wang, and A. Behal, "How autonomy impacts performance and satisfaction: Results from a study with spinal cord injured subjects using an assistive robot," *IEEE Transactions on Systems, Man, and Cybernetics-Part A: Systems and Humans*, vol. 42, no. 1, pp. 2–14, 2011.
- [2] A. D. Dragan and S. S. Srinivasa, "A policy-blending formalism for shared control," *The International Journal of Robotics Research*, vol. 32, no. 7, pp. 790–805, 2013.
- [3] D. Gopinath, S. Jain, and B. D. Argall, "Human-in-the-loop optimization of shared autonomy in assistive robotics," *IEEE Robotics and Automation Letters*, vol. 2, no. 1, pp. 247–254, 2016.
- [4] M. Li, D. P. Losey, J. Bohg, and D. Sadigh, "Learning user-preferred mappings for intuitive robot control," in *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2020.
- [5] H. J. Jeon, D. Losey, and D. Sadigh, "Shared autonomy with learned latent actions," in *Proceedings of Robotics: Science and Systems (RSS)*, Jul. 2020.
- [6] Y. Chang, M. Kutbi, N. Agadacos, B. Sun, and P. Mordohai, "A shared autonomy approach for wheelchair navigation based on learned user preferences," in *Proceedings of the IEEE International Conference on Computer Vision Workshops*, 2017, pp. 1490–1499.
- [7] M. Young, C. Miller, Y. Bi, W. Chen, and B. D. Argall, "Formalized task characterization for human-robot autonomy allocation," in *2019 International Conference on Robotics and Automation (ICRA)*. IEEE, 2019, pp. 6044–6050.