

The Sense of Agency in Assistive Robotics Using Shared Autonomy

Maggie A. Collier
Robotics Institute
 Carnegie Mellon University
 Pittsburgh, PA, USA
 macollie@andrew.cmu.edu

Rithika Narayan
Robotics Institute
 Carnegie Mellon University
 Pittsburgh, PA, USA
 rithikan@andrew.cmu.edu

Henny Admoni
Robotics Institute
 Carnegie Mellon University
 Pittsburgh, PA, USA
 hadmoni@andrew.cmu.edu

Abstract—Sense of agency is one factor that influences people’s preferences for robot assistance and a phenomenon from cognitive science that represents the experience of control over one’s environment. However, in assistive robotics literature, we often see paradigms that optimize measures like task success and cognitive load, rather than sense of agency. In fact, prior work has found that participants sometimes express a preference for paradigms, such as direct teleoperation, which do not perform well with those other metrics but give more control to the user. In this work, we focus on a subset of assistance paradigms for manipulation called shared autonomy in which the system combines control signals from the user and the automated control. We run a study to evaluate sense of agency and show that higher robot autonomy during assistance leads to improved task performance but a decreased sense of agency, indicating a potential trade-off between task performance and sense of agency. From our findings, we discuss the relation between sense of agency and optimality, and we consider a proxy metric for a component of sense of agency which might enable us to build systems that monitor and maintain sense of agency in real time.

Keywords—Assistive Robotics; Shared Autonomy; Sense of Agency; Sense of Control

I. INTRODUCTION

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 if the robot has a higher number of degrees of freedom (DOFs) than the input interface used for teleoperation, which is common among assistive systems that can be controlled by the user [1]. Fortunately, automation can help. That said, having *full* automation is not always feasible, and operators often want to maintain their sense of control over the system rather than having a fully automated robot [2].

One way to provide automated assistance to people while maintaining some of their control is via a shared autonomy approach [3], [4]. Shared autonomy 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 how they use it (see Fig. 1).

The underlying assumption of many existing shared autonomy paradigms is that users will prefer more autonomous assistance if it improves task performance metrics such as task




Fig. 1. A 24-participant study was run to explore questions on people’s sense of agency and preferences for the amount of assistance during grasping tasks. During the study, participants controlled a 7-DOF robot arm with a joystick and set the amount of assistance they wanted with a dial. We find that higher levels of robot assistance led to decreased sense of agency among users.

completion time, success rate, number of mode switches, etc. However, sometimes users prefer assistance that is not optimal with respect to these metrics, and sometimes they prefer the challenge of teleoperation to losing their sense of control when assistance is applied [2], [4], [6], [7]. For example, when users were able to customize the arbitration function in a shared autonomy study, they chose functions that did not perform optimally with respect to performance metrics, even after experiencing the optimal function [6]. 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 shared autonomy, user preferences have been studied in the context of learning a mapping between input interface actions and robot actions [8], [9] or trajectories to a goal location [10]. 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 [2], [4], we observe the users’ preferences directly by letting them control the arbitration level at any point during a task.

One factor we believe has an effect on people’s preferences for assistance is sense of agency (SoA), a phenomenon from cognitive science. SoA is the sense that one can take actions

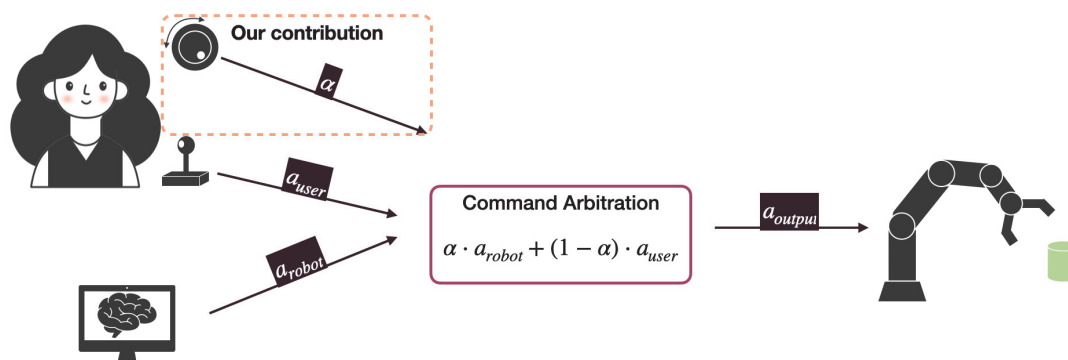


Fig. 2. Many shared autonomy paradigms work by arbitrating between input commands from the user and commands from the assistive system. In this work, we enable users to set the value of the arbitration parameter α at any point during a task, which is inspired by the approach in [5]. This framework enables us to uncover users’ preferences for autonomous assistance throughout tasks and study the effect of robot automation on the user’s sense of agency.

to change external events [11], [12]. SoA experiences are prevalent in human-machine interactions, especially in semi-autonomous systems in which some form of autonomous assistance and user input is being used to control the system [13]–[15]. SoA is an important factor that can determine technology adoption or use [16]. Additionally, SoA may be particularly important to assistive robot users with disabilities, since people with disabilities and neurodivergences experience agency and SoA differently than non-disabled, neurotypical populations [17]. Despite these facts and the prevalence of SoA experiences in semi-autonomous systems [13]–[15], assistive robotics literature does not usually explicitly address the user’s SoA. In order to create assistive systems that will actually be adopted by users, we must understand how our autonomous assistance affects the user’s SoA.

In this work, we deploy a common paradigm of robotic assistance in the form of goal-directed shared autonomy. We enable participants to operate an assistive robot arm under this shared autonomy in object-grasping tasks. We then conduct a 24-participant study recruited from a broad population in which we explore to what extent participants can perceive the autonomous assistance, how this kind of autonomous assistance affects SoA, how SoA relates to task performance metrics, and a potential proxy metric for SoA. Finally, we explore questions on people’s preferences for the amount of autonomous assistance they receive by letting them control the amount of assistance provided at any point throughout the task. Through this work, we make the following contributions:

- We demonstrate via a 24-participant study that higher goal-directed robot autonomy adversely affects SoA.
- We demonstrate an inverse relationship between task performance and SoA in the context of this form of robot autonomy.
- We posit that the angle of disagreement (see Eq. 2) between the executed command and the user’s command could be a proxy metric for a component of SoA that can be computed online in shared autonomy settings.
- We implement the first shared autonomy approach that

enables the user to directly set the arbitration level (see Fig. 2) at any point during a task.

- We are among the first to integrate the cognitive science theories of SoA, which is a prevalent phenomenon in human-machine interactions [13]–[15], into the context of physically assistive robotics.

II. RELATED WORK

This research draws from cognitive science work on SoA and research on user preferences for arbitration in shared autonomy.

A. SoA Definitions and Theories

SoA refers to the sense that one can take actions to change external events and is an experience of control over one’s environment [11], [12]. We discuss one of the predominant theories of SoA [18]–[20]. In the comparator model [18], [19], the human compares their expected outcome of an action they took with the actual outcome. When these outcomes mismatch in a perceptually significant manner for the user, their SoA is diminished. Later, we propose on an online proxy metric for SoA based on this model.

A recent review argues that assistive technology can act on SoA in a variety of ways, by augmenting the user’s actions, the user’s body, or the environment in order to achieve the intended outcome [13]. Crucially, it categorizes semi-autonomous systems as affecting SoA by augmenting a user’s action. Indeed, many forms of shared autonomy, which is a form of semi-autonomy, modify the user’s action [4]. Despite the prevalence of SoA experiences in semi-autonomous systems and human-machine interactions [13]–[15], SoA is not typically explicitly studied in assistive robotics, but terms like “sense of control” appear in the literature which may be related to sense of agency although those works do not draw from cognitive science theories (e.g. [21]). Here, we contribute one of the first assistive robotics studies to explicitly address the user’s SoA using theories from cognitive science.

B. Measuring SoA

There are only a few ways that SoA has been measured. Sometimes, researchers will craft their own SoA surveys to get subjective SoA responses, which corresponds to explicit measures of SoA because it focuses on the explicit *judgement* of agency, rather than the implicit *feeling* of agency. However, the psychology literature and some of the human computer interaction literature make use of a phenomenon called the intentional binding (IB) effect [22]. IB can be used to implicitly study SoA as it focuses on measuring the *feeling* of agency. In IB, participants must estimate the amount of delay between their control input and stimulus movement. Interestingly, people are more likely to estimate a shorter time than the actual delay duration when they did indeed cause the movement.

If we want to maintain or improve the user's SoA during a human-robot interaction, we need to have metrics for SoA or its components that we can compute in an online fashion. We contribute a proxy metric called the disagreement angle θ_d (see Eq. 2), an aspect of SoA that can be computed online. This proxy metric is based on predominant theories of SoA, such as the comparator model [18], [19]. Importantly, real-time SoA estimation methods do not yet exist, to our knowledge.

C. SoA in Robotics and Automation

Some foundational reviews of SoA in human-machine interactions and robotics are contributed in [14], [15]. Numerous studies indicate that, depending on the form of the robot, a robot collaborator tends to decrease the user's SoA in collaborative tasks [23]–[25]. These findings are consistent with other cognitive science literature showing that people experience diminished SoA when they perform a task in collaboration with another person [26]. These works indicate that a robot's involvement in a human-robot collaborative setting diminishes the user's SoA, underscoring the need to study this phenomenon in human-robot interactions and assistive robotics.

Some work, including this one, considers how automation technology affects SoA or vice versa. It may seem intuitive to assume that higher automation will adversely affect SoA, and this intuition is true some of the time (e.g. [27]). However, sometimes more automation can improve SoA [28]–[31]. These studies indicate that, although automation can at times adversely affect SoA, it is possible to build automation that improves or maintains the user's SoA.

Importantly, findings from the cognitive science literature on SoA and automation cannot easily be applied to robot autonomy. Cognitive science studies do not have a consistent definition of autonomy and deployed various forms of autonomy in a variety of non-robotics contexts, such as air traffic control. Similarly, autonomy in robotics studies varies widely [32], [33]. We need to study the effects of various forms of robot autonomy on the user's SoA. We contribute an exploration of the effects of a common form of robot autonomy in a shared autonomy system on SoA, and define that system in Section III-B1 using modern robot autonomy taxonomies.

D. Preferences in Shared Autonomy

Shared autonomy involves combining the commands from the user with an autonomous robot policy (see Fig. 2) [4]. Most literature on user preferences in shared autonomy deals with studying people's preferences for the actual shared autonomy paradigm itself [34]–[36], rather than the amount of assistance. For instance, in [21], user preferences for different modes of autonomy are explored in the context of a robot-assisted feeding task. Interestingly, participants tended to prefer the paradigms that required more effort but left them in greater control of the system.

Within the domain of shared autonomy, user preferences have been studied in the context of the mapping between input interface actions and robot actions [8], [9] and trajectories to a goal location [10]. The two studies involving user preferences most similar to our work both evaluate users' preferences for arbitration. In [4], researchers formalize shared autonomy as a blending between the robot's policy and the user's policy and present a study on users' preferences for arbitration in shared autonomy. Trends from the study suggest that users' preferences for the aggressiveness of assistance depends on the accuracy of the prediction of the users' intent and the complexity of the task as perceived by the user.

The second study most similar to our work had participants with spinal cord injury compare two forms of arbitration over the course of three weeks: manual teleoperation or full robot autonomy [2]. Interestingly, although full autonomy was associated with significantly less user effort, the mean satisfaction score at the end of the study for manual teleoperation was higher than the score for full autonomy. Instead of running a comparison study across different kinds of arbitration functions like in past work, we observe the users' preferences by letting them control the arbitration step at any point throughout a task.

E. Capturing Preference via User-controlled Arbitration

A shared autonomy paradigm that is similar to ours is proposed in [5]. The system enables users to directly set the arbitration parameter with a potentiometer, a slider interface, or a set of buttons, which is what our system does too but with a dial. Rather than evaluating our system like another shared autonomy paradigm as proposed in [5], we want to use the system to answer questions about human behavior in the presence of robot autonomy, users' preferences for robot autonomy, and users' SoA in the presence of robot autonomy. Importantly, the approach in [5] has yet to be implemented. Thus, something we contribute is the first implementation of a shared autonomy system that enables the user to directly set the arbitration parameter at any point during the task.

III. SYSTEM IMPLEMENTATION

In order to run a study on users' preferences for assistance and SoA, we developed a shared autonomy framework (see Fig 2) with three major components: user input, autonomous command generation, and arbitration. We implemented the

system on a Kinova 7-DOF Gen3 robot arm [37] with a Robotiq 2F-85 gripper [38].

A. User Input

Users can provide commands for the robot’s end effector via a 3-DOF joystick (the joystick can twist along its z -axis). The robot’s end effector has 6-DOFs of movement: x , y , z , *roll*, *pitch*, and *yaw*. However, our study only requires the translational DOFs, x , y , z . We map these three DOFs to each DOF of the joystick. Forward and backward movement of the joystick controls the end effector in x (i.e. forward and backward from the participant’s perspective). Left and right movement of the joystick controls the end effector in y (i.e. left and right from the participant’s perspective). Twisting the joystick along its z -axis controls the end effector in z (i.e. up and down from the participant’s perspective).

B. Autonomous Command Generation

We can set up our problem like a Markov Decision Process (MDP) with a single goal, with that goal being a pre-grasp pose in front of an object on the table that the robot arm is mounted to. Note that this MDP could be formulated as a multi-goal problem as well. The states S include the set of all environment states $s \in S$. The MDP includes the set of all robot’s actions $a_{robot} \in A_r$ (i.e. end effector velocity commands) and the set of all user’s actions $a_{user} \in A_u$ (i.e. end effector velocity commands), and we assume a deterministic transition function for simplicity. Additionally, we assume that users are stochastically following some policy $\pi^u(a_{user}|s)$ to which the robot does not have access. We compute some policy $\pi^r(a_{robot}|s, a_{user})$ that minimizes some cost function $C^r(s, a_{robot}, a_{user})$.

A common cost function for goal-based shared autonomy is to use some measure of the Euclidean distance between the goal and end effector pose. Because we are using grasping tasks in this study, we too can use this measure in our cost function, which enables us to solve for the MDP analytically. Essentially, we compute a straight-line trajectory between the current end effector pose and the goal pose. The direction of that straight-line path is taken as the direction of the autonomous velocity command a_{robot} . The magnitude of this command is the maximum velocity magnitude of the end effector (0.1 m/s) unless the end effector is very close to the goal pose, in which case it slows down until stopping at the goal pose. We choose goal poses and initial end effector poses such that the robot can reach the goal poses without reaching a joint limit or experiencing a kinematic failure.

1) *Defining Our Form of Robot Autonomy*: Based on the robot autonomy taxonomy in [32], our robot autonomy fits the definition of shared autonomy. The system is also low on the operational autonomy scale which is defined by [32] as “the degree of human operator disinvolvement at runtime.” Our robot autonomy is also high on the intentional autonomy scale which is defined as “the degree of robot goal-oriented involvement at runtime,” as it computes the shortest path to

the goal pose. Finally, our robot autonomy is high on the non-deterministic autonomy scale which is defined as “the degree to which a robot’s behavior is not specified prior to runtime,” as it *adaptively* computes the shortest path to the goal pose. We refer to this kind of robot autonomy throughout the manuscript as “goal-directed robot autonomy.”

C. Command Arbitration

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

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

where a_{output} is the output command, a_{robot} is the autonomous command, and a_{user} is the user’s command (see Fig. 2). An $\alpha = 0$ corresponds to teleoperation and $\alpha = 1$ to full autonomy.

Many past shared autonomy approaches compute α based on the system’s confidence in the inferred goal [4]. As confidence increases, so does the value of α , leading to more autonomous assistance. Importantly, user preferences are not considered for computing α . In contrast, we enable users to set the value of α at any point *throughout* the task which is inspired by [5] (see Fig. 2).

Our system arbitrates by sharing the velocity *direction* of the commands but not the magnitude. The velocity magnitude is determined by the joystick, unless the end effector is very close to the goal pose. The robot does not move unless the user is using the joystick and the end effector is not within some Euclidean threshold to the goal pose. Said differently, unless the end effector is at the goal pose, $|a_{output}| = |a_{user}|$, but the direction of a_{output} is determined via the arbitration function in Eq. 1.

IV. STUDY DESIGN

Our study has two major components: an assistance perception test and a preference study wherein participants freely adjust the level of assistance they want with a dial. In the assistance perception test, we have four research questions:

RQ1: *How far apart do levels of automation (i.e. different α values) have to be in order to be perceived as different by participants?*

RQ2: *What is the effect of increasing goal-directed robot autonomy (i.e. the value of α) on users’ SoA survey responses?*

RQ3: *What is the relationship between users’ SoA survey responses and task performance, as measured by trajectory length?*

RQ4: *What are potential proxy metrics for measuring SoA in real time in assistive robotics settings?*

In the preference study, we have two research questions:

RQ5: *What is the average preferred level of robot autonomy (i.e. α) among participants?*

RQ6: *How often do people change the level of robot autonomy (i.e. α) during a trial?*

A. Training

The training protocol familiarized participants with teleoperation and the assistive system.

1) *Teleoperation Training*: Participants begin the study with an introduction to the robotic arm, including use of the joystick to control the robot. The participant is asked to complete three grasping tasks with different objects: a Pringles can, an object holder built for the robot, and a small can of tomato paste. Each grasping task involves the user operating the arm under shared autonomy to guide the end effector to a pose just in front of the goal object. Once satisfied with the position of the end effector, participants press a button on the joystick that will send the end effector forward about two inches and close the robot's fingers. Participants were allowed to repeat some or all of the initial training tasks if they were not comfortable with teleoperation. Participants then attempt two more grasping tasks: a tall can and a golf ball on a tee. These tasks are used later as experimental tasks in the preference study.

2) *Assistive System Training*: The next section of training intends to acclimate the participant to the assistive system. Researchers describe the effects of using the dial throughout a task. The parameter α is described to participants as the fraction of the output command that comes from the robot's command. A diagram given to participants shows an example of differing user and robot commands and how the level of assistance affects the output command.

When given access to the dial, participants see the current level of assistance as a progress bar and a percentage on a computer monitor. The dial is initially set to an α value of 0.5 but can be set to values between 0 and 1; participants are asked to use the dial while attempting each experimental task again. They are told that the dial may be adjusted at any point during the task, including before the start. The initial α value and instructions for dial use are the same for all future tasks and training. Finally, participants get to attempt each experimental task again with the objects placed at different goal poses.

B. Assistance Perception Test

We first evaluate to what extent participants perceive differences in α or **RQ1**. Participants undergo an assistance perception test in which they perform a series of grasping actions under shared autonomy with predetermined and unknown assistance values. The dial interface is not in use during this test. The grasping tasks are completed in pairs of two across five rounds using the same object (small can of tomato paste). The second α value in each pair is chosen randomly from the set 0.00, 0.25, 0.50, 0.75, and 1.0 and not repeated. The first value in the pair is chosen by choosing a Δ from the set 0.00, 0.10, 0.20, 0.30, 0.40 with no repetitions. After each pair of grasping tasks, participants are asked to determine which task had a higher assistance level or if the levels were equal. Additionally, participants are asked a set of questions about the second task in the pair. These questions focus on the participant's SoA and causation of robot movement and

TABLE I
SENSE OF AGENCY SURVEY ITEMS FOR PERCEPTION TEST

Name in Fig 5	Likert Item
ObeY	The robot arm moved just like I wanted to, as if it was obeying my will.
Control	I felt as if I was controlling the movement of the robot arm.
Intended	I felt as if the robot arm did NOT move in the way I intended.
Movement	I felt as if I was causing the movement I saw.




Fig. 3. Experimental tasks for assistance preference study. The hard task involves grasping a golf ball off of a golf tee, and the easy task involves grasping the side of a tall cylinder.

are meant to evaluate **RQ2**, **RQ3**, and **RQ4**. Table I details each item on the SoA survey. This survey was adapted from a SoA study in virtual reality in which two people were sharing the control of a virtual avatar's arm [39].

C. Assistance Preference Study

In this study, we observe users' preferences for autonomous assistance by enabling them to set the arbitration parameter α at any point throughout an experimental task. This study is intended to address **RQ5** and **RQ6**.

1) *Experimental Tasks*: Each experimental task is a grasping task involving only translational movements. We manipulate the complexity of the task by manipulating the shape of the object to be grasped (see Fig. 3). For the lower complexity task, participants must grasp the side of a tall thin cylinder (a thin water bottle). For the higher complexity task, participants must grasp a golf ball off of a golf tee, which requires more precise alignment of the end effector. These are the same tasks used in the training section.

Participants are told that in order to succeed at the grasping task, the goal object must be inside the robot's gripper after its first closing. They are also instructed that the time taken to complete each task does not matter in order to decrease users' tendencies to attempt to complete tasks as fast as possible.

To decrease the dependence of goal position on our findings, we define two goal poses: one farther away from the base of the robot, and one closer to the base of the robot. Importantly, we chose both goal poses such that the robot's end effector could follow a straight-line path from its starting point to the goal pose without encountering a joint limit or a kinematic failure. Participants experience both experimental tasks multiple times and from each of these goal positions.

2) *Final Training*: Before collecting data, participants are given the opportunity to practice the experimental tasks with access to the dial. They attempt each task from each goal pose in a randomized order, for a total of four training trials. This gives participants the opportunity to converge to a preference for assistance if they have one and engage with the assistance according to those preferences during the data-collection trials.

3) *Data Collection*: In a randomized order, participants attempt both grasping tasks at both goal positions twice, for a total of eight recorded trials. We record the α value, joystick information, joint states of the robot, end effector information, and video of participants undergoing each trial.

D. Metrics

The value of α throughout the assistance preference study is an important metric that enables us to measure assistance preference and address **RQ5** and **RQ6**. Additionally, the angle of disagreement between the user’s command and the robot’s command (see Eq. 2) has been explored to develop a shared autonomy paradigm that allows for dimension-specific α values [40]. However, our work focuses on investigating if the disagreement between the user’s command and the executed command could serve as a proxy metric for some aspect of SoA, or **RQ4**. Having a proxy metric for SoA could enable us to monitor people’s SoA and maintain it throughout an interaction with an assistive robot. Therefore, the disagreement angle is an important metric in this work:

$$\theta_d = \arctan2(|a_{output} \times a_{user}|, a_{output} \cdot a_{user}) \quad (2)$$

where θ_d is the disagreement angle, a_{output} is the executed command, and a_{user} is the user’s command (see Fig. 2). This equation can be extended to rotational commands as well by computing the quaternion difference between the rotations in the executed command a_{output} and the user’s command a_{user} .

E. Participants & Recruitment

We recruited 24 participants from our university’s Center for Behavioral and Decision Research (CBDR) after achieving IRB approval. Around 46% of participants identify as women, and none of our participants identify as non-binary. About 58% of our participants are between the ages of 18 and 24, 25% are between 25 and 30, 4% are between 31 and 35, 4% are between 41 and 45, and 8% are at or above the age of 66. Around 67% of our participants are Asian or Asian American, 21% of our participants are Caucasian, 4% of our participants are Black, and 8% responded “Other” or “two or more races/ethnicities.”

V. RESULTS & DISCUSSION

In this section, we establish several experimental outcomes: participants can perceive different levels of goal-directed robot autonomy, higher goal-directed robot autonomy adversely affects SoA, there is a trade-off between SoA and task performance metrics, and disagreement could be a proxy metric for one aspect of SoA. Additionally, we detail the α preferences from the assistance preference study.




Fig. 4. Participants can perceive a difference in goal-directed robot automation (RQ1). The left bar at each Δ (the difference in α values of both trials in a round of the perception test) shows the proportion of participants who got the answer right (shown in purple). The right bar at each Δ show the proportion of participants who indicated the wrong condition (shown in orange) or incorrectly answered “equal” (shown in gray). Note that at $\Delta = 0$, the right answer is “equal.”

A. Participants Can Perceive a Difference in Goal-directed Robot Autonomy

According to Fig. 4, participants can perceive a difference in level of automation down to 0.2. (Recall that level of automation ranges between 0 and 1). Remember that Δ represents the difference between the two α s in a round of the perception test. At a Δ of 0, the correct answer is “equal.” The “equal” responses monotonically decrease as the Δ increases. More participants answered correctly than not for all Δ s except at 0.1. At 0.1, the majority of participants said the two conditions were equal. Thus, there appears to be a perceptual threshold between a Δ of 0.1 and 0.2. The “equal” responses show the proportion of participants who were not sensitive to that particular Δ . Many participants were sensitive to a Δ of 0.4 and 0.3. All participants got at least one round of the perception test correct, with an average performance of three of five rounds correct.

B. Higher Goal-directed Robot Autonomy Adversely Affects Sense of Agency

Boxplots for participants’ responses for each of the SoA items at α s of 0, 0.25, 0.5, 0.75, and 1 are shown in Fig. 5. (Note that the “Intended” SoA scores have been reverse coded in Fig. 5.) The Cronbach alphas for the SoA responses to each α value we tested are as follows for $\alpha = 0, 0.25, 0.5, 0.75,$ and 1 : 0.78, 0.95, 0.94, 0.78, and 0.77, respectively. These values indicate internal reliability among our four Likert items.

All SoA items decrease as the level of automation increases. Fig. 5 shows that SoA scores decrease as level of automation increases, which addresses **RQ2**. This finding implies that, as we increase goal-directed robot automation in our assistive systems, we decrease the user’s SoA. This result is not necessarily bad news for roboticists. There may be a way to

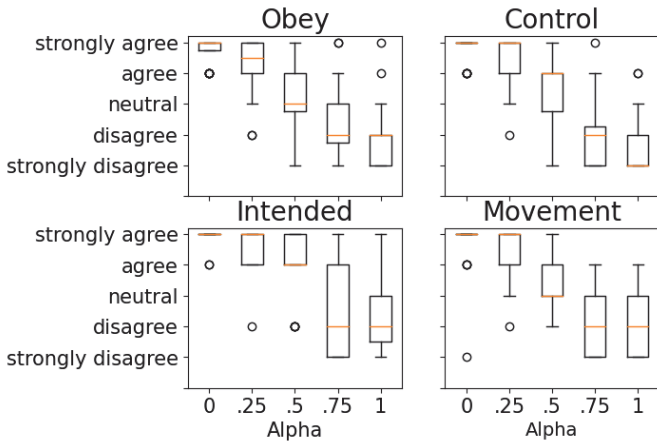


Fig. 5. Participants’ sense of agency decreases as goal-directed robot autonomy increases (RQ2). Participants’ sense of agency scores at an α of 0, 0.25, 0.5, 0.75, and 1 are shown. The data for “Intended” has been reverse coded. Table I details each sense of agency item and what plot they map to.

apply this kind of automation and still preserve the user’s SoA, which we discuss in the next section.

For statistical tests, we first ran a non-parametric version of an ANOVA called the Kruskal Wallis H-test [41] on all the data from each Likert item in Table I. Results from these tests and the post-hoc tests (see the Appendix) we ran imply that the effect of automation on SoA may not be that different for very high or very low values of α but is different in the middle of the autonomy scale.

C. Task Performance Inversely Related to Sense of Agency

When we apply automation, we are often trying to optimize for task performance metrics such as task completion time or trajectory length. In this work, we use trajectory length as our optimality metric.

In Fig. 6, the average SoA score across the survey items is plotted against trajectory length, and each marker type represents an α value of 0, 0.25, 0.5, 0.75, and 1, respectively. An α of 0 and 0.25 is associated with longer trajectory lengths. As we look from left-to-right across the average SoA scores, the trajectory lengths get longer (i.e. less optimal) as the SoA score increases (i.e. improves), and lower α s are associated with longer trajectory lengths. We ran a Pearson’s correlation on this data and found a positive correlation of $r(118) = 0.52$, $p = 8.9e^{-10}$. This finding indicates that there may be a trade-off between SoA and optimality (i.e. trajectory length) that users must consider, which addresses **RQ3**.

To summarize, as SoA improves, task performance gets worse. That said, we also see that the higher α values can, at times, achieve a close-to-optimal trajectory while still maintaining an SoA score above 3, which maps to “neutral.” This result indicates that we can achieve high-performing assistance that still maintains people’s SoA, but in general, there may be a trade-off between SoA and task performance.




Fig. 6. Task performance is inversely related to sense of agency (RQ3). Trajectory length vs avg. sense of agency score is plotted. Each marker type indicates a data point taken at an α of 0, 0.25, 0.5, 0.75, and 1. The x-axis corresponds to an avg. sense of agency score across a 5-point Likert scale.




Fig. 7. Disagreement could be a proxy metric for one component of sense of agency (RQ4). Avg. disagreement θ_d (see Eq. 2) between the executed command and the user command is plotted against the avg. sense of agency score. Each marker type indicates a data point taken at an α of 0, 0.25, 0.5, 0.75, and 1. As sense of agency score improves (increases), the avg. disagreement decreases.

D. Disagreement Could Be a Proxy Metric for a Component of Sense of Agency

In Fig. 7, the avg. disagreement angle θ_d (see Eq. 2) between the executed command and the user command is plotted against the avg. SoA score. Higher avg. disagreement angles occur for datapoints that fall below a 3 for avg. SoA score. Additionally, datapoints at higher α values tend to fall below a 3 on SoA. The data in Fig. 7 addresses **RQ4** by indicating that disagreement angle could be a proxy metric for one component of SoA. By having a proxy metric, we could monitor the user’s SoA during a task in order to maintain it.

The disagreement angle connects to a theory of SoA that accounts for the *feeling* of agency – the comparator model, which we discussed in Sect. II-A [18], [19]. The comparator




Fig. 8. Histogram of mean α values over the course of 192 trials (RQ5). The y-axis displays the proportion of all trials. The x-axis indicates mean α .

model says that people are comparing their expected outcome of an action they take to the actual outcome. Importantly, we posit that users expect the executed command a_{output} (actual outcome) to match the user’s issued command a_{user} (expected outcome) regardless of the value of α . When there is a perceptually significant mismatch between the two commands, the experience of control (i.e. SoA) is diminished. We can capture these mismatches via our disagreement metric (see Eq. 2) which corresponds nicely to the comparator model. The comparator model may account for the *feeling* of agency but does not account for the *judgment* of agency [20]. Hence, our metric may only account for one component of SoA: the implicit feeling of agency.

E. Preferences for α Vary with Most Trials Falling around 0.5

In Fig. 8, a histogram of each trials’ mean α is shown. Of the 192 trials, 53.6% have a mean α close to or at 0.5. About 20.3% of the trials had a mean α around 1, and 4.7% of trials were at an $\alpha = 0$. This data addresses **RQ5**. The default value of the dial at the start of each trial was 0.5, which we chose in order to avoid biasing participants towards direct teleoperation or full autonomy. By setting the initial α value to 0.5, we may have biased participants to that value.

Of the 192 trials, 66.7% had a standard deviation of 0, indicating those trials did not include an assistance change during the trial. This data addresses **RQ6**. We determined in the assistance perception test that participants can perceive differences in automation down to a Δ of 0.2, which is equivalent to a standard deviation of 0.1 assuming the trial only had one change. All but two of the trials in which there was an assistance change included only one change. Thus, we can calculate the percentage of trials in which a perceptually significant α change occurred, which is 15.6% of all 192 trials. We believe the cognitive load required to interact with the dial may have deterred many participants from using it during the task. However, since 15.6% of trials included a perceptually significant α change, we have evidence that a dynamic α may be preferred by some users.

F. Implications of Findings & Future Work

We are often optimizing for task performance metrics in our assistive systems. However, **our work shows that higher levels of autonomy or assistance can diminish users’ SoA and implies that optimal task performance might not actually be preferable**. What our work reveals is a tension between the experience of control over one’s assistive device (i.e. SoA) and task efficiency. Rather than setting a high level of assistance to achieve the optimal task performance, we can set the assistance to a level that maintains the user’s SoA and still achieves close-to-optimal task performance (see Fig. 6).

We cannot directly apply the findings from our study to assistive device users with disabilities and neurodivergences, because these populations may have different SoA experiences than non-disabled and neurotypical populations [11], [17]. Future work will augment this study for people with upper mobility limitations so that we can better understand the SoA of people who might use our assistive systems in the real world. We acknowledge that, as assistive technology researchers, we often make assumptions about users’ agency and build systems around those assumptions. These assumptions may be wrong and include biases, for instance, if the researchers are non-disabled. Reflexive approaches [42] are being increasingly used in the field of HCI [43] to mitigate issues with bias when researchers are working with marginalized populations.

VI. CONCLUSION

In this work, we studied people’s preferences for assistance and a factor that might influence those preferences: SoA. Many assistive systems seek to optimize task success metrics like task completion time or trajectory length, but these systems often fail to account for the user’s SoA. Through a study with 24 participants, we showed that higher levels of goal-directed robot automation adversely affects SoA. Additionally, we found that task performance is inversely related to SoA, but that a close-to-optimal performance can be achieved while still maintaining the user’s SoA. We also contribute an online proxy metric that may account for the *feeling* of agency, which is one aspect of SoA [20]. Due to the prevalence of SoA experiences in human-machine interactions [13]–[15], assistive robotics researchers should consider the effects of their forms of robot assistance on the user’s SoA, as a lack of SoA felt over an assistive device may lead to misuse or abandonment of the device [16]. This work is the first step to understanding the cognitive science phenomenon of SoA in the context of assistive robotics.

ACKNOWLEDGEMENTS

We would like to thank CBDR which enabled us to recruit from the Pittsburgh community. We would also like to thank the many labmates who piloted our system and study before we ran it. This work was funded by NSF awards IIS-1943072 and CMMI-2024794.

REFERENCES

- [1] L. V. Herlant, R. M. Holladay, and S. S. Srinivasa, "Assistive teleoperation of robot arms via automatic time-optimal mode switching," in *2016 11th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE, 2016, pp. 35–42.
- [2] 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.
- [3] S. Javdani, H. Admoni, S. Pellegrinelli, S. S. Srinivasa, and J. A. Bagnell, "Shared autonomy via hindsight optimization for teleoperation and teaming," *The International Journal of Robotics Research*, vol. 37, no. 7, pp. 717–742, 2018.
- [4] 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.
- [5] S. Jain and B. Argall, "An approach for online user customization of shared autonomy for intelligent assistive devices," in *Proc. of the IEEE Int. Conf. on Robot. and Autom., Stockholm, Sweden*, 2016.
- [6] 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.
- [7] R. Latikka, N. Savela, A. Koivula, and A. Oksanen, "Attitudes toward robots as equipment and coworkers and the impact of robot autonomy level," *International Journal of Social Robotics*, vol. 13, no. 7, pp. 1747–1759, 2021.
- [8] 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.
- [9] H. J. Jeon, D. Losey, and D. Sadigh, "Shared autonomy with learned latent actions," in *Proceedings of Robotics: Science and Systems (RSS)*, Jul. 2020.
- [10] Y. Chang, M. Kutbi, N. Agadakos, 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.
- [11] P. Haggard, "Sense of agency in the human brain," *Nature Reviews Neuroscience*, vol. 18, no. 4, pp. 196–207, 2017.
- [12] P. Haggard and V. Chambon, "Sense of agency," *Curr. Biol.*, vol. 22, pp. R390–R392, 2012.
- [13] P. Cornelio, P. Haggard, K. Hornbaek, O. Georgiou, J. Bergström, S. Subramanian, and M. Obrist, "The sense of agency in emerging technologies for human-computer integration: A review," *Frontiers in Neuroscience*, vol. 16, p. 949138, 2022.
- [14] W. Wen, Y. Kuroki, and H. Asama, "The sense of agency in driving automation," *Frontiers in Psychology*, vol. 10, p. 2691, 2019.
- [15] W. Wen and H. Imamizu, "The sense of agency in perception, behaviour and human-machine interactions," *Nature Reviews Psychology*, vol. 1, no. 4, pp. 211–222, 2022.
- [16] R. Parasuraman and V. Riley, "Humans and automation: Use, misuse, disuse, abuse," *Human factors*, vol. 39, no. 2, pp. 230–253, 1997.
- [17] C. D. Frith, S.-J. Blakemore, and D. M. Wolpert, "Abnormalities in the awareness and control of action," *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences*, vol. 355, no. 1404, pp. 1771–1788, 2000.
- [18] I. Feinberg, "Efference copy and corollary discharge: implications for thinking and its disorders," *Schizophrenia bulletin*, vol. 4, no. 4, p. 636, 1978.
- [19] C. D. Frith, *The cognitive neuropsychology of schizophrenia: Classic edition*. Psychology press, 2015.
- [20] M. Synofzik, G. Vosgerau, and A. Newen, "Beyond the comparator model: a multifactorial two-step account of agency," *Consciousness and cognition*, vol. 17, no. 1, pp. 219–239, 2008.
- [21] T. Bhattacharjee, E. K. Gordon, R. Scalise, M. E. Cabrera, A. Caspi, M. Cakmak, and S. S. Srinivasa, "Is more autonomy always better? exploring preferences of users with mobility impairments in robot-assisted feeding," in *Proceedings of the 2020 ACM/IEEE international conference on human-robot interaction*, 2020, pp. 181–190.
- [22] P. Haggard, S. Clark, and J. Kalogeras, "Voluntary action and conscious awareness," *Nature neuroscience*, vol. 5, no. 4, pp. 382–385, 2002.
- [23] F. Ciardo, D. De Tommaso, F. Beyer, and A. Wykowska, "Reduced sense of agency in human-robot interaction," in *Social Robotics: 10th International Conference, ICSR 2018, Qingdao, China, November 28-30, 2018, Proceedings 10*. Springer, 2018, pp. 441–450.
- [24] Z. Barlas, "When robots tell you what to do: Sense of agency in human- and robot-guided actions," *Consciousness and cognition*, vol. 75, p. 102819, 2019.
- [25] O. Grynspan, A. Sahaï, N. Hamidi, E. Pacherie, B. Berberian, L. Roche, and L. Saint-Bauzel, "The sense of agency in human-human vs human-robot joint action," *Consciousness and cognition*, vol. 75, p. 102820, 2019.
- [26] J. D. Loehr, "The sense of agency in joint action: An integrative review," *Psychonomic Bulletin & Review*, vol. 29, no. 4, pp. 1089–1117, 2022.
- [27] B. Berberian, J.-C. Sarrazin, P. Le Blaye, and P. Haggard, "Automation technology and sense of control: a window on human agency," *PLoS one*, vol. 7, no. 3, p. e34075, 2012.
- [28] S. Ueda, R. Nakashima, and T. Kumada, "Influence of levels of automation on the sense of agency during continuous action," *Scientific reports*, vol. 11, no. 1, p. 2436, 2021.
- [29] K. Inoue, Y. Takeda, and M. Kimura, "Sense of agency in continuous action: Assistance-induced performance improvement is self-attributed even with knowledge of assistance," *Consciousness and cognition*, vol. 48, pp. 246–252, 2017.
- [30] W. Wen, S. Yun, A. Yamashita, B. D. Northcutt, and H. Asama, "Deceleration assistance mitigated the trade-off between sense of agency and driving performance," *Frontiers in Psychology*, vol. 12, p. 643516, 2021.
- [31] S. Endo, J. Fröhner, S. Musić, S. Hirche, and P. Beckerle, "Effect of external force on agency in physical human-machine interaction," *Frontiers in Human Neuroscience*, vol. 14, p. 114, 2020.
- [32] S. Kim, J. R. Anthis, and S. Sebo, "A taxonomy of robot autonomy for human-robot interaction," in *Proceedings of the 2024 ACM/IEEE International Conference on Human-Robot Interaction*, 2024, pp. 381–393.
- [33] J. M. Beer, A. D. Fisk, and W. A. Rogers, "Toward a framework for levels of robot autonomy in human-robot interaction," *Journal of human-robot interaction*, vol. 3, no. 2, p. 74, 2014.
- [34] A. Erdogan and B. D. Argall, "The effect of robotic wheelchair control paradigm and interface on user performance, effort and preference: an experimental assessment," *Robotics and Autonomous Systems*, vol. 94, pp. 282–297, 2017.
- [35] M. N. Javaremi, M. S. Young, D. Gopinath, and B. D. Argall, "User preference in shared control of a robotic wheelchair: A longitudinal study," in *meeting of HRI Conference*, 2018.
- [36] D. Paddeu, I. Tsouros, G. Parkhurst, A. Polydoropoulou, and I. Shergold, "A study of users' preferences after a brief exposure in a shared autonomous vehicle (sav)," *Transportation Research Procedia*, vol. 52, pp. 533–540, 2021.
- [37] Kinova. (2024) Kinova. [Online]. Available: <https://www.kinovarobotics.com/product/gen3-robots>
- [38] Robotiq. (2024) Robotiq 2f-85 and 2f-140 grippers. [Online]. Available: <https://robotiq.com/products/2f85-140-adaptive-robot-gripper>
- [39] R. Fribourg, N. Ogawa, L. Hoyet, F. Argelaguet, T. Narumi, M. Hirose, and A. Lécuyer, "Virtual co-embodiment: evaluation of the sense of agency while sharing the control of a virtual body among two individuals," *IEEE Transactions on Visualization and Computer Graphics*, vol. 27, no. 10, pp. 4023–4038, 2020.
- [40] M. Bowman and X. Zhang, "Dimension-specific shared autonomy for handling disagreement in telemanipulation," *IEEE Robotics and Automation Letters*, vol. 8, no. 3, pp. 1415–1422, 2023.
- [41] W. H. Kruskal and W. A. Wallis, "Use of ranks in one-criterion variance analysis," *Journal of the American statistical Association*, vol. 47, no. 260, pp. 583–621, 1952.
- [42] J. P. Jones, H. J. Nast, and S. M. Roberts, *Thresholds in feminist geography: Difference, methodology, representation*. Rowman & Littlefield, 1997.
- [43] C. A. Liang, S. A. Munson, and J. A. Kientz, "Embracing four tensions in human-computer interaction research with marginalized people," *ACM Transactions on Computer-Human Interaction (TOCHI)*, vol. 28, no. 2, pp. 1–47, 2021.
- [44] W. J. Conover and R. L. Iman, "Multiple-comparisons procedures. informal report," Los Alamos National Lab.(LANL), Los Alamos, NM (United States), Tech. Rep., 1979.