

Visualization of Intended Assistance for Acceptance of Shared Control

Connor Brooks¹ and Daniel Szafir²

Abstract—In shared control, advances in autonomous robotics are applied to help empower a human user in operating a robotic system. While these systems have been shown to improve efficiency and operation success, users are not always accepting of the new control paradigm produced by working with an assistive controller. This mismatch between performance and acceptance can prevent users from taking advantage of the benefits of shared control systems for robotic operation. To address this mismatch, we develop multiple types of visualizations for improving both the legibility and perceived predictability of assistive controllers, then conduct a user study to evaluate the impact that these visualizations have on user acceptance of shared control systems. Our results demonstrate that shared control visualizations must be designed carefully to be effective, with users requiring visualizations that improve both legibility and predictability of the assistive controller in order to voluntarily relinquish control.

I. INTRODUCTION

Through *shared control* [1], human intelligence can be combined with robot autonomy in order to aid direct human teleoperation. This control paradigm involves two separate controllers — a human and an assistive controller — working in concert to operate the same physical robot. While such systems are designed to employ advances in autonomous robotics in order to improve human-in-the-loop control, this requires the human and assistive controller to share a common goal. Any mismatch in goals or intended trajectories may result in these controllers working against one another. Thus, cooperation between these two controllers is critical for successful operation.

If a human’s goal for a robotic system is known in advance, assistance can be provided without concern for goal misidentification [2], [3]. For more general applicability, the user’s goal must first be inferred by an assistive system. Inferring a human’s goal from their control inputs allows an assistive controller to estimate the target of assistance through the user’s natural interactions with the system. Existing research has established various methods for performing such inference, ranging from proximity-based estimation [4], [5], [6] to the use of an Inverse Reinforcement Learning framework for reward function estimation [7], [8], [9].

Although shared control has been shown to improve efficiency, human teleoperators often display a hesitancy to use these systems [10], [11]. We believe that one potential reason for this hesitancy is a lack of understanding of the assistive controller’s intent. For example, an assistive controller




Fig. 1. Augmented Reality visualizations demonstrate the assistive controller’s intended target and future trajectory keypoints.

might generate behavior that could be perceived as counter-productive (e.g. following a trajectory that temporarily moves away from the user’s intended goal to overcome a joint limitation) and cause uncertainty for the user about whether the assistive controller has correctly inferred their goal.

In this work, we investigate how improving the communication between assistive controllers and users may improve user acceptance and adoption of shared control systems. In particular, we investigate how Augmented Reality (AR) can provide *in situ* visualizations that directly communicate assistive controller goals and strategies to users during shared control of a robotic arm. Furthermore, we develop multiple levels of visualization in order to assess how much information is necessary to help users understand the behavior of an assistive controller. We evaluate these AR visualizations in a user study involving multiple types of assistive controllers in order to generalize this evaluation across different types of shared control systems. Finally, we present and discuss our results demonstrating that these visualizations make the shared control system more predictable to users and users are more willing to hand over control to the assistive controller when they are able to see its intended trajectory and goal.

II. RELATED WORK

Despite assistive teleoperation systems improving user efficiency, prior work has found mixed results on human teleoperator preference for these systems versus those that give the user more control [8], [10], [11]. Dragan and Srinivasa [10] suggest that, based on user feedback, it is important for the assistive controller to exhibit *legible* behavior — movement that helps the user understand its goal [12]. Our work follows directly from this idea through evaluation of two types of AR visualizations: one designed to make the assistive controller’s goal legible and the other to make the

¹University of Colorado Boulder, Department of Computer Science
connor.brooks@colorado.edu

²University of Colorado Boulder, Department of Computer Science and ATLAS Institute
daniel.szafir@colorado.edu

robot’s movement *predictable* — movement that matches the user’s expectations.

Visualizations of future robot movement are notably useful in control systems that allow users to select from high-level robot behaviors [13]. Weisz et al. [14] use AR to display possible robotic arm grasps that a user iterates through and selects in order to issue commands. A similar design is behind the *Director* control system designed by Marion et al. for use in the DARPA Robotics Challenge [15] that uses visualizations of robot commands to help operators choose appropriate actions. In contrast to these systems, we are interested in assistive teleoperation involving continuous blending of direct user commands and assistive controller inputs throughout the task.

We use AR for displaying visualizations of assistive controller intent because this modality enables visualizations of information situated in the relevant real world locations. This situated information has been shown to be useful for human-robot interaction by improving human understanding of future robot movement [16], limiting context switching during teleoperation [17], and conveying environmental information to human users [18] (for a further review of AR and human-robot interaction, see [19]). In particular, Rosen et al. [20] have shown that mixed-reality visualizations of trajectories enable users to better understand future robotic arm movement than 2D displays. Motivated by these strong results, we use AR visualizations to help users in understanding the assistive controller during shared control.

Our work is partially inspired by Zolotas et al. [21], who developed an AR system for shared control of an assistive wheelchair. Their system is intended to communicate the information that the assistive controller uses for planning in order to improve explainability. In further work, Zolotas and Demiris [22] evaluate the performance of users operating the wheelchair with and without their AR explanations and find that these explanations benefit user performance. However, the shared control systems investigated in these works do not contain goal inference, instead augmenting human control by handling complimentary tasks such as obstacle and collision avoidance. This results in shared control systems with a more indirect collaboration than those in which we are interested. Furthermore, our work extends this existing research by evaluating multiple levels of visualization across different types of assistive controllers. Finally, we investigate factors not addressed previously such as the impact of AR visualizations on the user’s acceptance of shared control.

Contributions: Our research addresses questions about shared control raised in prior work by evaluating the impact of visualizing an assistive controller’s intent during shared control involving continuous collaboration between a human user and an assistive controller. Our contributions include an assessment of AR visualizations for assistive teleoperation of a robotic arm, the creation and assessment of multiple levels of visualization across different types of assistive controllers, and an analysis of how users handle control authority in response to AR visualizations.

III. HYPOTHESES

Informed by past research, we form three hypotheses about the impact of assistive controller intent visualizations on shared control:

H1: Visualizing assistive controller intent will increase user willingness to relinquish control in shared control.

H2: Visualizing assistive controller intent will improve user assessment of the shared control system’s usability.

H3: Visualizing assistive controller intent will improve user assessment of the shared control system’s predictability.

IV. SYSTEM DESIGN

To test these hypotheses, we create an assistive controller for an object grasping task. Providing assistance in a cluttered environment requires inference of the user’s goal, which is complicated by the possibility of different possible grasp positions for a given goal. We model the user’s goal using a Hidden Markov Model (HMM) that allows us to take into account multiple grasps for a single goal.

A. HMMs for Goal Inference

Creating an HMM requires the determination of a set of possible hidden states X , the transition probabilities between each pair of these states $p(x'_{t+1}|x_t)$, and the observation probabilities at each state $p(u_t|x_t)$ given a set of possible observations U . Prior work in assistive teleoperation systems has utilized HMMs by letting the hidden state represent the goal of interest in the user’s underlying intent while the observations at each state correspond to user commands [23]. With this approach, the transition probabilities between states model the likelihood of changing goals during the task [24].

We apply and extend this idea by creating separate states for each possible grasp of a goal object. Our hidden states can be grouped into classes corresponding to the enumeration of possible grasps for a particular goal. Thus, each hidden state, x , is a member of a single class C_ϕ corresponding to the goal ϕ for which that state represents a particular grasp. The probability of a goal given a sequence of observations $\xi_{0 \rightarrow t} = (u_0, \dots, u_t)$ is found by summing over the probabilities of each grasp for that goal:

$$p(\phi_t|\xi_{0 \rightarrow t}) = \sum_{x \in C_\phi} p(x_t|\xi_{0 \rightarrow t}) \quad (1)$$

B. Determining State Transition Probabilities

In this model, a transition probability between members of the same class versus different classes represents the probability of switching intended grasps on an object versus switching to a particular grasp of a new goal, respectively. In order to specify these probabilities from one timestep to the next, we use two parameters: T_{grasp} is the probability of switching intended grasps on the same goal, and T_{goal} is the probability of switching goals (where we constrain the parameter choices such that $T_{grasp} + T_{goal} \leq 1$). Then, we can calculate the transition probabilities going out of a state x using the following equations:




Fig. 2. A sphere around the target of assistance is intended to make the assistive controller's intent legible while keypoints of the assistive controller's intended trajectory are intended to make the system's future movement more predictable.

$$p(x'_{t+1}|x_t) = \begin{cases} \frac{T_{grasp}}{|C_\phi|-1}, & \text{where } x, x' \in C_\phi \text{ and } x \neq x' \\ \frac{T_{goal}}{|\Phi|-1} \cdot \frac{1}{|C_{\phi'}|}, & \text{where } x' \in C_{\phi'}, x \notin C_{\phi'} \\ 1 - (T_{grasp} + T_{goal}), & \text{where } x = x' \end{cases} \quad (2)$$

C. Determining Observation Probabilities

In order to perform filtering in this HMM, we need to specify the probability distribution over user actions for a given hidden state. That is, we need a way to calculate $p(u_t^H|x_t)$, where u_t^H is the user's action at timestep t .

Each grasp x consists of a particular position and orientation of the robot's end effector. Let the state of the robot's end effector at timestep t be given by s_t . Then, we define a reward function $R_x(s, u)$ that estimates the value of choosing action u in state s while attempting to reach x .

The implementation of this reward function can be designed to incorporate desired characteristics as our framework is agnostic to the specific implementation choice. We use a reward function defined by the anticipated difference in distance to the intended grasp after taking action u :

$$R_x(s, u) = \text{dist}(s, x) - \text{dist}(\tau(s, u), x) \quad (3)$$

where $\tau(s, u)$ produces the resulting end effector state from applying action u in state s .

We use a Boltzmann distribution to model the operator as exponentially more likely to choose actions with increasing reward, following prior work in Inverse Reinforcement Learning [25] and assistive teleoperation [24]:

$$p(u_t^H|x_t) = \frac{\exp(R_{x_t}(s_t, u_t^H))}{\sum_{u \in U} \exp(R_{x_t}(s_t, u))} \quad (4)$$

D. Belief State Updating and Action Selection

The use of an HMM allows us to perform iterative updating of the probability assigned to each particular grasp using the forward algorithm:

$$p(x_t|\xi_{0 \rightarrow t}) = p(u_t^H|x_t) \sum_{x'} p(x'_{t-1}|\xi_{0 \rightarrow t-1})p(x_t|x'_{t-1}) \quad (5)$$

Combining this with Eq. 1 produces an updated estimate of the probability of each goal after each user action.

In our system, if any goal is assigned a probability greater than 50%, the assistive controller selects the action u_r that minimizes the expected distance toward an expert-demonstrated grasp of this target goal. In practice, these grasps could be generated by a planner using a robot's onboard perception system. The expert-demonstrated grasps we use are multi-point trajectories consisting of two to three keypoints that demonstrate a safe, reliable way to grasp this particular goal.

E. Blending of Assistive Controller and Human Commands

The final command sent to the robot is determined through the linear blending:

$$u^* = \alpha u_r + (1 - \alpha)u_h$$

where u_r is the command issued by the assistive controller, u_h is the command issued by the user, and α is an *arbitration* factor. This strategy of linear blending has an extensive history of usage in shared control [10].

Determining this arbitration factor impacts the final behavior significantly by setting the weighting between human and assistive controller. Another impact of this blending scheme is an effect on the robot's speed; if the assistive controller and user have opposite commands or if the user inputs no command, the robot's final speed will be slower in systems where α is not sufficiently close to 0 or 1. Consequently, we recognize that the effects of AR visualizations may differ based on the human-controller interactions produced by different arbitration factor values. We use multiple systems with differing values of α in this research in order to draw conclusions about the impact of our visualizations that are applicable to a range of shared control systems.

F. Visualizations of Assistive Intent

The visualizations that our system generates are focused on two goals: (1) improving legibility of the assistive controller's intent and (2) enhancing predictability of the shared controller's future movement. While there are many possible ways to design visualizations for these two goals, for this study, we chose two visualizations that explicitly convey this information while requiring minimal user interpretation.

In order to improve legibility, our system displays a sphere around the inferred goal. This makes the assistive controller's inference immediately legible by explicitly calling out its goal of assistance. This sphere appears any time a goal hits 50% probability in our goal inference system, centered on the newly inferred assistance target. Fig. 2 shows an example of this sphere around the houseplant object.

To make the shared control system's movement predictable, we use virtual images of the robot's end effector that trace out the remaining keypoints in the expert-demonstrated grasp of this object. These visualizations directly illustrate the path that the assistive controller will attempt to follow. Fig. 2 also contains an example of this type of visualization.

V. USER STUDY

We collect data to test our hypotheses through a mixed-design 3 (visualization [between participants]: none, goal only, goal + trajectory) \times 3 (controller [within participants]: low assistance, medium assistance, high assistance) user study with 27 participants (19 male, 8 female). In the study, participants interact with several shared control systems while operating a Fetch robot to grasp objects on a nearby table. Throughout this process, participants are shown one of three types of AR visualizations. The independent variable in which we are primarily interested is the visualization shown to participants, consisting of three conditions:

1) *No Visualization*: The participants are not shown any AR visualizations. This is our baseline condition representing how shared control systems work in practice today.

2) *Goal Only Visualization*: The participants are shown the goal sphere visualization whenever a goal is inferred.

3) *Goal + Trajectory Visualization*: The participants are shown both the goal sphere visualization and the virtual images of the assistive controller's intended path.

Our second independent variable represents the shared control systems with which the user interacts. The three shared controllers use the system described in Section IV with three different α -levels: $\alpha = 0.25$, $\alpha = 0.5$, $\alpha = 0.99$. We use $T_{grasp} = .01$ and $T_{goal} = 0$ to allow for transitions between target grasps while assuming a fixed goal. Participants use an $\alpha = 0$ (pure teleoperation) system first in order to establish a baseline level of control and subjective judgement of the robotic system, followed by the other three controllers in randomized order to mitigate any potential transfer effects (e.g., learning, fatigue).

A. Experiment Setup

Participants operate the robot with a PlayStation controller using a modal system consisting of two control modes (position and angular control). This type of modal control is a common paradigm for teleoperation of high-DOF systems for which direct control over all joints at once may be impractical [26], [27]. Before the experiment begins, each participant is given a tutorial on how to use the controller, followed by a 2 minute practice period. During the experiment, all participants wear a Microsoft HoloLens — a popular AR head-mounted display (ARHMD) that we use to display the relevant visualizations for each participant's condition.

For each shared controller, participants operate the robot to grasp the four objects on the table. Before each grasp, participants are shown a letter in AR to instruct them which object to grasp next (this also provides a justification to participants in the No Visualization condition for the ARHMD). Participants are allowed four failed grasps on an object before it is skipped. Figs. 1 and 2 show the experimental setup consisting of four objects spread out on the table.

The experiment consists of four rounds of grasping (once using pure teleoperation and once for each of the assistive controllers), each round including a single grasp of each of the four objects on the table. Each round of grasping is followed by completion of surveys on the shared control

system just used. The entire experiment thus consists of performing 16 successful grasps and 4 rounds of surveys, taking approximately 30 minutes per participant.

B. Measures

In order to evaluate our hypotheses, we assess both objective and subjective metrics that provide insight into the joint behavior of the human and assistive controllers.

The objective metrics that we collect data on are *completion effort*: the total number of user actions required per successful grasp, and *acceptance of assistance*: the percentage of the trajectories from each successful grasp during which the user does not actively input controls. As a manipulation check, we compare *completion effort* across controller types to test if our shared control systems improve user efficiency. We use *acceptance of assistance* to evaluate the willingness of users to relinquish control to the assistive system (**H1**).

We administer a survey after participants interact with each controller to measure subjective participant experience. We use the System Usability Scale (SUS) in order to utilize a validated metric for assessing the perceived *usability* of a technology interface (**H2**). Additionally, after participants use each controller with an $\alpha > 0$, we record responses on a 2-item scale (Cronbach's $\alpha = 0.869$), which consists of 5-point Likert-style questionnaire items ("I understood why the robot moved the way it did", "I found the robot's movement predictable"), to evaluate subjective perceptions of the shared control system's *predictability* (**H3**).

VI. ANALYSIS AND RESULTS

After 27 participants completed our study, we observed that only one participant failed the limit of four times on a particular grasp (No Visualization condition, $\alpha = 0.99$ controller). Generally, we conduct a one-way analysis of variance (ANOVA) and find no significant difference in the number of grasp failures between visualization conditions $F(2,24) = 1.09$, $p = 0.353$. We receive no successful trajectory data from this failed grasp as well as one other grasp for which our recording crashed, leaving us with data from 322 successful grasps between the 3 shared controllers.

For each of the measures discussed previously, we conduct a mixed-design ANOVA, treating visualization type as the fixed effect and controller type as the within-subjects variable. If the ANOVA demonstrates a significant effect of the independent variable of interest on a measure, we conduct a *post hoc* Dunnett's test to identify which conditions demonstrate significant differences from the baseline.

A. Objective Results

We find that controller type has a significant effect on *completion effort* $F(3,72) = 13.05$, $p < 0.0001$. A *post hoc* Dunnett's test against the $\alpha = 0$ baseline ($M = 88.77$, $SD = 35.85$) shows significant improvements for $\alpha = 0.5$ ($M = 65.07$, $SD = 54.28$, $p = 0.014$) and $\alpha = 0.99$ ($M = 43.62$, $SD = 45.17$, $p < 0.0001$) controllers, but none for the $\alpha = 0.25$ controller ($M = 85.48$, $SD = 38.70$, $p = 0.959$). This confirms our manipulation checks that our

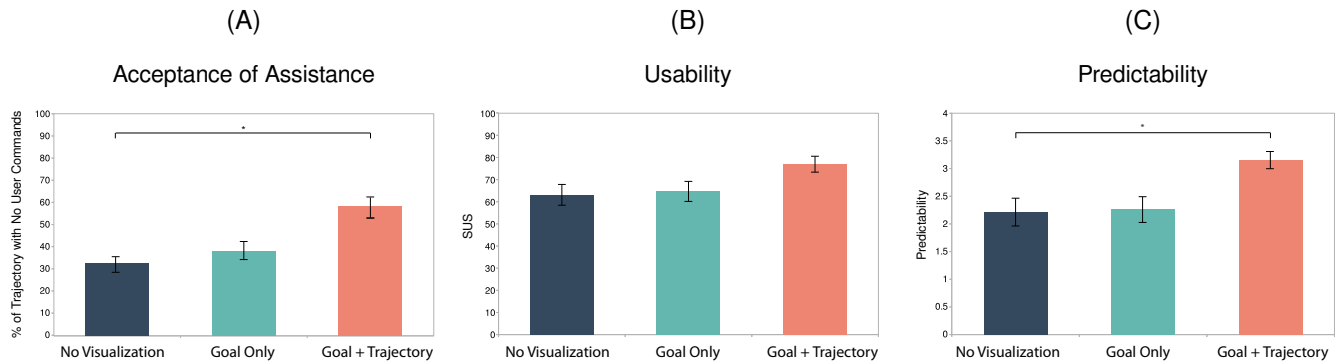


Fig. 3. (A) Users refrain from inputting control commands for a significantly larger percentage of the trajectory when shown the Goal + Trajectory visualizations. (B) Although there is an increasing average SUS score with more visualizations, we find no statistically significant effect of visualizations on perceived usability. (C) Our results shows a significant effect of visualization on user rating of the shared control system’s predictability. This improvement is only found in the Goal + Trajectory condition over the No Visualization baseline, with no significant difference between the baseline and the Goal Only visualization. In each of these subgraphs, * signifies a results with a significance level of $p < 0.05$.

assistive controllers improve user efficiency at higher levels of α irrespective of visualization condition.

While the previous measure assesses overall user control efficiency, **H1** is concerned with how the human user cooperates with the assistive controller. Our results demonstrate that visualization has a significant effect on *acceptance of assistance* $F(2,24) = 5.10$, $p = 0.014$. A post-hoc Dunnett’s test against the No Visualization ($M = 32.73$, $SD = 18.32$) baseline demonstrates that the Goal + Trajectory condition ($M = 58.19$, $SD = 22.64$) significantly improved user acceptance of assistance ($p = 0.010$) while Goal Only ($M = 38.77$, $SD = 19.59$) does not ($p = 0.694$). These results show support for **H1** when using both visualizations to increase goal legibility and trajectory predictability while not finding any such support with visualizations designed to increase legibility alone. Fig. 3 (A) illustrates these results.

B. Subjective Results

We use the validated SUS scale to measure subjective ratings of the shared control systems’ *usability*. We do not find a significant effect of visualization on this measure $F(2, 24) = 1.84$, $p = 0.181$ and thus do not find support for **H2**. The mean SUS scores for No Visualization ($M = 63.06$, $SD = 24.47$), Goal Only ($M = 64.72$, $SD = 23.31$), and Goal + Trajectory ($M = 77.04$, $SD = 18.89$) are shown in Fig. 3 (B).

The final measure we assess is the user’s subjective rating of the shared control system’s *predictability*. Our results reveal a significant effect of visualization on this rating $F(2, 24) = 4.25$, $p = 0.026$. A post-hoc Dunnett’s test against the No Visualization ($M = 2.22$, $SD = 1.31$) baseline condition again demonstrates that the Goal + Trajectory ($M = 3.15$, $SD = 0.806$) visualization produces a significant improvement ($p = 0.031$) while Goal Only ($M = 2.26$, $SD = 1.18$) visualization does not ($p = 0.992$). Consequently, we find support for **H3** when using both types of visualizations while again not finding any such support for legibility visualizations alone. These results are displayed in Fig. 3 (C).

C. Discussion

Our results demonstrate that users find shared control systems more predictable when given visualizations of the assistive controller’s inferred goal and intended trajectory. In these cases, participants are more willing to hand over control to the assistive system, creating trajectories with less direct human control, although they do not rate these systems subjectively higher in usability. Thus, we find evidence for **H1** and **H3** specifically when using both types of visualizations, but we do not find support for **H2**. These findings demonstrate the potential for visualizations of assistive controller intent to improve cooperation between humans and assistive controllers in shared control systems while highlighting the complexities of improving user preference for these systems.

We find no evidence supporting our hypotheses when using only visualizations intended to improve assistive controller legibility. This may show that in order to obtain the benefits of users properly understanding a shared control system, users need to understand both the target and the intended trajectory of assistance. Although further work is needed to confirm this trend, our results suggest that legibility alone may not improve user cooperation with an assistive controller unless it is also accompanied by user-predictable behavior.

VII. FUTURE WORK

Our results provide insight into some promising areas for further research. For instance, future work could investigate how our visualizations affect the actual actions taken by the human user. Shared control systems can influence users during cooperative control [28], raising the possibility that visualizing the assistive controller’s intended trajectory could affect the user’s intended trajectory. Future work investigating whether or not this takes place would help better explain the mechanism of cooperation in systems utilizing our visualizations.

Since our results demonstrate that visualizations can provide benefits to shared control, this creates an open ques-

tion about how to best implement these visualization for different types of shared control. We create and evaluate two representative types of visualizations, but a more thorough exploration of this design could be enlightening. Furthermore, future work could examine the combination of these visualizations with other common feedback systems for shared control such as haptic feedback. Finally, as our visualizations that are intended to improve predictability also increase legibility, we leave the effect of predictability visualizations that do not improve legibility as a question for future work.

VIII. CONCLUSION

Through shared control, robot autonomy can be used to empower a human user of a robotic system without removing the user from the control loop. While much work in shared control is focused on improving the assistive controller's understanding of the user's intent (e.g. goal inference), this work introduces and evaluates a way to improve the user's understanding of the assistive controller's intent. Our results demonstrate the potential for positive impact from these visualizations as well as the importance of sufficient information communication for this impact to be seen. This also has implications for shared control beyond the design of AR visualizations: the legibility of an assistive controller, without predictability, may not be sufficient for improving acceptance of shared control. We look forward to seeing these types of visualizations further explored in future work on improving shared control of robotic system.

IX. ACKNOWLEDGEMENTS

This work was supported by a grant from the NSF under award 1764092.

REFERENCES

- [1] D. A. Abbink, T. Carlson, M. Mulder, J. C. de Winter, F. Aminravan, T. L. Gibo, and E. R. Boer, "A topology of shared control systems-finding common ground in diversity," *IEEE Transactions on Human-Machine Systems*, vol. 48, no. 5, pp. 509–525, 2018.
- [2] J. Philips, J. d. R. Millán, G. Vanacker, E. Lew, F. Galán, P. W. Ferrez, H. Van Brussel, and M. Nuttin, "Adaptive shared control of a brain-actuated simulated wheelchair," in *2007 IEEE 10th International Conference on Rehabilitation Robotics*. IEEE, 2007, pp. 408–414.
- [3] H. Yu, M. Spenko, and S. Dubowsky, "An adaptive shared control system for an intelligent mobility aid for the elderly," *Autonomous Robots*, vol. 15, no. 1, pp. 53–66, 2003.
- [4] A. Goil, M. Derry, and B. D. Argall, "Using machine learning to blend human and robot controls for assisted wheelchair navigation," in *2013 IEEE 13th International Conference on Rehabilitation Robotics (ICORR)*. IEEE, 2013, pp. 1–6.
- [5] K. D. Katyal, M. S. Johannes, S. Kellis, T. Afalo, C. Klaes, T. G. McGee, M. P. Para, Y. Shi, B. Lee, K. Pejisa, *et al.*, "A collaborative bci approach to autonomous control of a prosthetic limb system," in *2014 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*. IEEE, 2014, pp. 1479–1482.
- [6] J. Kofman, X. Wu, T. J. Luu, and S. Verma, "Teleoperation of a robot manipulator using a vision-based human-robot interface," *IEEE transactions on industrial electronics*, vol. 52, no. 5, pp. 1206–1219, 2005.
- [7] A. D. Dragan and S. S. Srinivasa, *Formalizing assistive teleoperation*. MIT Press, July, 2012.
- [8] S. Javdani, S. S. Srinivasa, and J. A. Bagnell, "Shared autonomy via hindsight optimization," *Robotics science and systems: online proceedings*, vol. 2015, 2015.
- [9] C. Schultz, S. Gaurav, M. Monfort, L. Zhang, and B. D. Ziebart, "Goal-predictive robotic teleoperation from noisy sensors," in *Robotics and Automation (ICRA), 2017 IEEE International Conference on*. IEEE, 2017, pp. 5377–5383.
- [10] 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.
- [11] C. Brooks and D. Szafir, "Balanced information gathering and goal-oriented actions in shared autonomy," in *2019 14th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*. IEEE, 2019, pp. 85–94.
- [12] A. D. Dragan, K. C. Lee, and S. S. Srinivasa, "Legibility and predictability of robot motion," in *Proceedings of the 8th ACM/IEEE international conference on Human-robot interaction*. IEEE Press, 2013, pp. 301–308.
- [13] K. M. Tsui, D.-J. Kim, A. Behal, D. Kontak, and H. A. Yanco, "i want that: Human-in-the-loop control of a wheelchair-mounted robotic arm," *Applied Bionics and Biomechanics*, vol. 8, no. 1, pp. 127–147, 2011.
- [14] J. Weisz, P. K. Allen, A. G. Barszap, and S. S. Joshi, "Assistive grasping with an augmented reality user interface," *The International Journal of Robotics Research*, vol. 36, no. 5-7, pp. 543–562, 2017.
- [15] P. Marion, M. Fallon, R. Deits, A. Valenzuela, C. Pérez D'Arpino, G. Izatt, L. Manuelli, M. Antone, H. Dai, T. Koolen, *et al.*, "Director: A user interface designed for robot operation with shared autonomy," *Journal of Field Robotics*, vol. 34, no. 2, pp. 262–280, 2017.
- [16] M. Walker, H. Hedayati, J. Lee, and D. Szafir, "Communicating robot motion intent with augmented reality," in *Proceedings of the 2018 ACM/IEEE International Conference on Human-Robot Interaction*, 2018, pp. 316–324.
- [17] H. Hedayati, M. Walker, and D. Szafir, "Improving collocated robot teleoperation with augmented reality," in *Proceedings of the 2018 ACM/IEEE International Conference on Human-Robot Interaction*, 2018, pp. 78–86.
- [18] T. Yamamoto, N. Abolhassani, S. Jung, A. M. Okamura, and T. N. Judkins, "Augmented reality and haptic interfaces for robot-assisted surgery," *The International Journal of Medical Robotics and Computer Assisted Surgery*, vol. 8, no. 1, pp. 45–56, 2012.
- [19] D. Szafir, "Mediating human-robot interactions with virtual, augmented, and mixed reality," in *International Conference on Human-Computer Interaction*. Springer, 2019, pp. 124–149.
- [20] E. Rosen, D. Whitney, E. Phillips, G. Chien, J. Tompkin, G. Konidaris, and S. Tellex, "Communicating robot arm motion intent through mixed reality head-mounted displays," in *Robotics Research*. Springer, 2020, pp. 301–316.
- [21] M. Zolotas, J. Elsdon, and Y. Demiris, "Head-mounted augmented reality for explainable robotic wheelchair assistance," in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2018, pp. 1823–1829.
- [22] M. Zolotas and Y. Demiris, "Towards explainable shared control using augmented reality," 2019.
- [23] D. Aarno, S. Ekvall, and D. Kragic, "Adaptive virtual fixtures for machine-assisted teleoperation tasks," in *Proceedings of the 2005 IEEE international conference on robotics and automation*. IEEE, 2005, pp. 1139–1144.
- [24] S. Jain and B. Argall, "Recursive bayesian human intent recognition in shared-control robotics," in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2018, pp. 3905–3912.
- [25] D. Ramachandran and E. Amir, "Bayesian inverse reinforcement learning," *Urbana*, vol. 51, no. 61801, pp. 1–4, 2007.
- [26] L. V. Herlant, R. M. Holladay, and S. S. Srinivasa, "Assistive teleoperation of robot arms via automatic time-optimal mode switching," in *The Eleventh ACM/IEEE International Conference on Human Robot Interaction*. IEEE Press, 2016, pp. 35–42.
- [27] D. E. Gopinath and B. Argall, "Mode switch assistance to maximize human intent disambiguation," in *Robotics: Science and Systems*, 2017.
- [28] F. Flemisch, D. A. Abbink, M. Itoh, M.-P. Pacaux-Lemoine, and G. Weßel, "Joining the blunt and the pointy end of the spear: towards a common framework of joint action, human-machine cooperation, cooperative guidance and control, shared, traded and supervisory control," *Cognition, Technology & Work*, vol. 21, no. 4, pp. 555–568, 2019.