

Control-Theoretic Analysis of Shared Control Systems

Reuben M. Aronson
Department of Computer Science
Tufts University
Medford, Massachusetts, USA
reuben.aronson@tufts.edu

Elaine Schaertl Short
Department of Computer Science
Tufts University
Medford, Massachusetts, USA
elaine.short@tufts.edu

Abstract—Users of shared control systems change their behavior in the presence of assistance, which conflicts with assumptions about user behavior that some assistance methods make. In this paper, we propose an analysis technique to evaluate the user’s experience with the assistive systems that bypasses required assumptions: we model the assistance as a dynamical system that can be analyzed using control theory techniques. We analyze the shared autonomy assistance algorithm and make several observations: we identify a problem with runaway goal confidence and propose a system adjustment to mitigate it, we demonstrate that the system inherently limits the possible actions available to the user, and we show that in a simplified setting, the effect of the assistance is to drive the system to the convex hull of the goals and, once there, add a layer of indirection between the user control and the system behavior. We conclude by discussing the possible uses of this analysis for the field.

I. INTRODUCTION

Recent work [1] shows that users actively explore the behavior of assistive systems, reason about how they behave, and change their own inputs to drive the systems to their desired goals. The behavior is difficult to reconcile with assistance methods that require assumptions about user behavior, such as goal-predictive assistance systems [2]. In fact, we theorize the problem may be more serious: if users always have the “last word,” analytical tools that require assumptions about user behavior or response will encounter recursive issues as users continually adapt to changes in the system.

In this paper, we propose an alternative analysis technique based on this insight: rather than model the user directly, we model the rest of the robot and assistance systems as a dynamical system that the user is controlling. This technique is not a substitute for other methods of developing assistance. Rather, it enables examining the user’s experience independently from whatever assumptions about the user, environment, or task that the assistance relies on. In this way, all assistance can be viewed as changes to a dynamical system that the user is operating, and any adaptation or learning performed by the system becomes added dynamics for the robot.

To develop this approach, we turn to dynamical systems and control theory. These fields, which are commonly used in robotic systems, study how a system behaves over time in response to its inputs, internal dynamics, and outputs. Control theory further explores how to design rules for providing

inputs to guarantee certain behaviors of the system while knowing only the outputs. Our key insight is that when operating a system with assistance, the *user* is performing precisely this control operation however the assistance is designed. The assistance system defines the system *dynamics* that the user controls, and designing assistance behavior is equivalent to designing the plant. Thus, we can use insights from control theory about what makes dynamics easy or hard to control and apply them to evaluating the assistance itself.

In this paper, we develop this control-theoretic analysis of shared control systems by analyzing of shared autonomy [3], a popular shared control formulation. We show that the assistance acts as a pure integrator, which has some undesirable properties, and use controls analogies to propose mitigations. We next show that the assistance restricts the possible actions that the system can take no matter the user’s intentions. Finally, we discuss a simplified setting, free-space navigation, and show that the dynamics of the system reduce to the user moving a set point that the system settles at. We conclude with a discussion of next steps for this analytic approach and implications for the design of shared control systems.

II. BACKGROUND

A. Shared Control

In a shared control framework, both the user and an assistive algorithm simultaneously control a robot [4]. The assistance modifies the robot’s behavior to drive the system in a way that the user “actually means” — motivations for shared control include challenges of low-dimensional interfaces [5, 6] or safety-critical situations [7]. Goal-predictive assistance [2] breaks the problem down into *intent detection*, during which the system estimates the “true” goals of the user using various methods [8–10], and *arbitration*, during which the user’s initiative is mixed with an assistive action generated based on the detected intent. These systems improve performance in many settings [11–15]. However, they require strong assumptions about user behavior which do not necessarily hold in practice [1], thus benefiting from our approach.

Other assistance methods focus directly on the dynamics. Most directly, virtual fixtures [16] apply artificial forces to the robot’s input response. More sophisticated approaches change the behavior of the robot to encourages successful task

performance [17, 18]. Our analysis broadens this assistance strategy to characterize assistance behavior that is not designed this way explicitly.

Though shared control research generally focuses on algorithmic advances, several papers consider the human experience of controlling a robot in a shared control paradigm [1, 19]. This work includes explicitly modeling how the user will adapt to the system [20, 21], changing the system to make the user’s behavior more informative [22, 23], and adjusting the user model based on observed behavior [14, 24, 25].

B. Control Theory

Control theory analyzes how to drive a dynamical system to a particular desired output. Dynamical systems models consist of three components: the *input* u , the *state* x , and the *output* y . Then, control theory answers the following question: Assume we can provide inputs u to the system and observe output y , but we do not have direct access to x . How do we select the input u to drive the system to produce some desired output y^* ? For this paper, we examine assistance from the perspective of the user, who can provide input u to the assistance system and observe output y but does not have direct access to the internal state x . We then use insights from control theory about what makes dynamics easy or hard to control and apply them to designing the assistance itself.

Control theory techniques are common in robotics in general. Most low-level motor controllers are PID (proportional-integral-derivative) systems, while model-predictive control is used for short-horizon optimization of robot low-level commands. Here, we consider the inverse problem: how do we design the dynamics of a system so that it is easy for a controller to be designed that drives the system to an intended goal?

III. DYNAMICAL ANALYSIS OF SHARED AUTONOMY

To explore this assistance analysis, we present a worked example. We analyze the shared autonomy control paradigm as a dynamical system that the user controls and show that it acts as an integrator, which is known to cause unbounded lag when trying to change intended goals. We also show that the presence of the assistance fundamentally limits the actions that the user can take regardless of their input. Further, we show that in an obstacle-free navigation task, the system dynamics are fully determined by a set point (the expected value of the goal locations over the goal probabilities): the user adjusts the set point location by changing the goal probabilities, and the assistance drives the system to that set point.

A. Background: Shared Autonomy

We begin by briefly summarizing the shared autonomy algorithm presented in Javdani et al. [26]. The algorithm expects a pre-specified (finite) list of k goals G along with optimal policies for each goal defined by action-value functions $Q_g(x, a)$. It receives an input u_t from the user at each time step, uses a Bayesian update framework to estimate the user’s true goal within G as a probability distribution, then takes

an action that is optimal in expectation over the probability of the user’s true goal. This assistive action is then applied to the system itself. The algorithm is intended to improve performance of the system for a user who knows their goals but cannot optimally drive the robot to achieve them; the robot uses its own (presumably) optimal policy to reach them instead. This system decreases task completion time and user effort required for reaching the goal.

At each time step, the system estimates the likelihood of an observed user action u_t given each goal. Several models for the likelihood are used [9, 10]; here, we will show the Boltzmann rationality formulation, which sets

$$p(u_t|x_t, g) \propto \exp \beta Q_g(x_t, u_t) \quad (1)$$

with x_t the current state, β a “rationality” parameter that varies from $\beta = 0$, indicating random user actions, to $\beta = \infty$, indicating perfect user actions. Note that this probability is conditioned on the state x_t ; this approach treats the user’s input as informative separately from the actual robot state in which it was provided. For ease of notation, we drop x_t throughout.

Next, this likelihood is used to update the assistance system’s estimate of the user’s goal, stored as a probability distribution $p(g|u_0, \dots, u_t)$. Via Bayes’ rule, we can set

$$p(g|u_0, \dots, u_t) = \frac{1}{Z_t} p(u_t|g) p(g|u_0, \dots, u_{t-1}), \quad (2)$$

with $Z_t = \sum_g p(g|u_0, \dots, u_t)$ a normalization factor.

Finally, this system generates assistance by selecting an action a_t that has the largest expected increase in value over the goal distribution, given by

$$a_t = \arg \max_{a'} \sum_g p(g|u_0, \dots, u_t) Q_g(x_t, a_t). \quad (3)$$

This equation is derived from applying the QMDP assumption to solving a POMDP representation of the user’s behavior; see the paper for further details.

B. Dynamical Systems Representation

We now show how the same assistance system appears to the user as a dynamical system. To do so, we define the internal state of the system and its dynamics, how the state changes based on the input, and how the output of the system is derived from the state.

First, for convenience, we define $\vec{Q}_t(u) = [Q_{g_0}(u), Q_{g_1}(u), \dots, Q_{g_{k-1}}(u)]^\top$ as a vector-valued function representing the value of taking the same action u in the state x_t with respect to each of the goals in turn. This transformation turns a possible action u into a vector of length k that describes how well u aligns with each goal in turn.

Next, we define the state of the assistance system ℓ_t as

$$\ell_t = [\log p(g_0|u_0, \dots, u_t), \dots, \log p(g_{k-1}|u_0, \dots, u_t)], \quad (4)$$

the log of the goal probability vector used in each step. By taking the log of Eqn. 2 and substituting ℓ_t as appropriate, we obtain the system dynamics

$$\ell_t = \ell_{t-1} + \log p(u_t|g) - \log Z_t. \quad (5)$$

Note that this is a linear system in ℓ , but is nonlinear in the user input u_t . However, the user contribution to the update is in a single additive term, so we can perform a change in variables to remove it. We define $v_t(g) = \log p(u_t|g)$ as contribution of the user's input towards the system's selection of goal g ; using Eqn. 1, we compute

$$v_t(g) = \beta Q_g(u_t) - \log \sum_{u' \in U} \exp(\beta Q_g(u')). \quad (6)$$

Using this equation, we can express the dynamics in terms of the input vector $\vec{v}_t = [v_t(g_0), v_t(g_1), \dots, v_t(g_{k-1})]$ as

$$\ell_t = \ell_{t-1} + \vec{v}_t - \log Z_t. \quad (7)$$

Next, we consider the output of the system a_t . Using Eqn. 3 and substituting ℓ_t , we find

$$a_t = \arg \max_{a'} \exp \ell_t \cdot \vec{Q}(a'). \quad (8)$$

Now, note that the term Z_t is positive, uniform over g , and only used inside the $\arg \max$ term, so it can be dropped entirely (though normalization may be required for numerical stability). Thus, our final system is

$$\begin{aligned} \ell_t &= \ell_{t-1} + \vec{v}_t \\ a_t &= \arg \max_{a'} \exp \ell_t \cdot \vec{Q}_t(a'). \end{aligned}$$

C. Dynamics Analysis

The system dynamics for shared autonomy are relatively simple. Though (generally) nonlinear functions $v_t(g) = v(x, u, g)$ are required to process the raw user input u_t and compute the output action a_t , the system dynamics consists of a single integrator.

This representation gives us several insights into the behavior of the assistance system from the perspective of the user. Pure integrator systems are known to have infinite DC gain: if supplied a constant input, the system will increase without bound. In this case, providing a constant user command will cause the probabilities assigned to goals other than the most likely to decrease without bound. If the user has been pursuing one goal for some time and then switches to a different goal, the user must drive the system towards the new goal with an amount of effort proportional to what was spent achieving the first goal before the system will register a change in their intention. Though the original formulation of shared autonomy assumes that users will not change their goals, we can still analyze the system behavior if they do.

Dynamical systems theory also gives us insight into how to mitigate this effect. Since we have full control over the dynamics of the goal update, we can introduce a new term to Eqn. 2 that adds internal stability to the system:

$$\ell_t = k\ell_{t-1} + v_t, \quad (9)$$

where $0 \leq k \leq 1$ is a stabilizing term that limits how extreme the values of ℓ_t can become. Given a consistent user input, the probability distribution will eventually reach a steady-state value (ignoring the effect of renormalization), so the effort to




Fig. 1. While there may be many possible actions that the system can take in some state x , here plotted with their Q-values corresponding to two different goals, the assistance system only ever selects actions that lie on the Pareto frontier, indicated by the blue circles.

change the system to select a different goal remains bounded. Control theory offers other options as well, such as adding a term dependent on $\vec{v}_t - \vec{v}_{t-1}$ which monitors the *change* in user input. The analysis suggests changes in the dynamics of the system that the user is operating which can be evaluated in future user studies.

D. Experience of User Control

Now, we turn to a different question: what are the possible actions that an operator can cause the system to produce? This question makes no sense in a robot-centric perspective: the goal-specific policy defines the optimal actions for each goal, so no additional actions are needed. However, users may want to drive the system to produce other actions based on their individual information or preferences, and we can use similar analysis tools to determine their options. To examine the possible actions, we consider Eqn. 3. For the system to select an action a , the action must be the $\arg \max$ for *some* value of ℓ_t . Furthermore, since $\exp \ell_t > 0$, the term inside the $\arg \max$ is a positive combination of $\vec{Q}$; by definition, then, $\vec{Q}_t(a)$ must lie on the Pareto frontier of $\{\vec{Q}_t(a) : a \in A\}$. Therefore, not all actions A are possible. In fact, the possible actions are a function only of the state x_t . User input can only select among the Pareto-optimal actions, and other actions are impossible to achieve no matter the user input (see Fig. 1).

With simplifying assumptions, we can go further. Consider a free-space navigation system with real states $S = \mathbb{R}^n$ and real unit actions $A = \{a \in \mathbb{R}^n : |a| = 1\} \cup \{0\}$; we will denote these states and action as $\vec{x}$ and $\vec{a}$ for clarity. Further, let each goal $g \in S$ indicate a target state, which we will denote $\vec{g}$; G represents a matrix $[g_0, g_1, \dots, g_{k-1}]$ of all goals with each row corresponding to a goal. We posit a closed-form value function $Q_g(\vec{x}, \vec{a}) = (\vec{g} - \vec{x})^T \vec{a}$ which computes how well the proposed action $\vec{a}$ aligns with the direction to the goal $(\vec{g} - \vec{x})$; with this assumption, the system is a discrete-time analogue of potential function navigation¹. Then, $\vec{Q}(a) = Q^T \vec{a}$ is in fact

¹On the domain $|g - x| > 1$, with $V(x) = |g - x|$, $r(x) = (g - x)^T a - |g - x - a|$, and $T(x, a) = x + a$, this Q satisfies the Bellman equation.

a matrix multiplication with

$$Q = \begin{bmatrix} | & | & & | \\ \vec{g}_0 - \vec{x} & \vec{g}_1 - \vec{x} & \cdots & \vec{g}_{k-1} - \vec{x} \\ | & | & & | \end{bmatrix}. \quad (10)$$

In this scenario, we can solve the $\arg \max$ term directly:

$$\begin{aligned} \vec{a}^* &= \arg \max_{\vec{a}} \exp \ell \cdot \vec{Q}(\vec{a}) \\ &= \arg \max_{\vec{a}} (\exp \ell)^T Q^T \vec{a} \\ &= \frac{Q \exp \ell}{|Q \exp \ell|}. \end{aligned}$$

Thus all assistance actions $\vec{a}$ are normalized positive combinations of the columns of Q , i.e., within the positive span of the directions towards each goal.

From this analysis, we can draw several conclusions. First, the assistance can produce any action in A precisely when the positive span of Q includes the whole space. This scenario occurs when the origin lies within the convex hull of Q or, equivalently, when the state $\vec{x}$ is inside the convex hull of the goals $\vec{g}$. Thus, the assistance dynamics have two phases: when the state is outside the convex hull of the goals, the assistance will drive the state towards the convex hull no matter the user input. The user can change the direction of approach among the available goals but cannot drive the system away from the convex hull. Once the system is within the convex hull, any action can be achieved by manipulating the assistance state ℓ_t so long as the system state remains within the convex hull; the user cannot drive the system outside the hull. This hull always exists because $|G|$ is finite (by assumption), though it can be degenerate; in this case, the system is driven to and remains in the hyperplane defined by the goals and by the convex hull within that hyperplane.

Another result of this analysis is that the system can be made to stop at any point within the convex hull of the goals. Formally, for all states $\vec{x} \in \text{conv}(\vec{g})$, there is some vector v with all nonnegative elements such that $Qv = 0$. Since v is nonnegative, we can select $\exp \ell = v$; for this probability state, the optimal action $a = \text{norm}(Q \exp \ell) = 0$. Conversely, all goal prediction states have a unique stable point at $\vec{g} \exp \ell = \sum_k p(k)g_k = x^*$, the convex combination of the goal locations with weights equal to the goal probabilities; this fact follows from solving $Q \exp \ell = 0$ for $\vec{x}$.

Putting everything together, shared autonomy imposes a two-phase system. When the system state is outside the acceptable region, defined by the (possibly lower-dimensional) convex hull of the goals, the system will drive it towards the convex hull independent of user input. Within the convex hull, though, the system supports any desired action, but through a layer of indirection. The user input updates the set point x^* via the goal inference step, and then the system drives the state towards that set point. Within the convex hull, users cannot control the robot directly; rather, they move this (invisible) set point, and then system responds. Controlling this system is like driving a car and controlling its speed only through cruise

control. The layer of indirection may also explain why users remain unsatisfied with the system despite improvements in trial metrics [3] and particularly dislike pure assistance without any arbitration with the user's raw input [27]. This analysis also explains the equilibrium that users discovered intuitively in Aronson and Short [1]: by carefully altering the set point, users can cause the system to stop at any given location within the convex hull of the goals.

IV. DISCUSSION AND CONCLUSION

In this paper, we perform a dynamical systems analysis of the shared autonomy framework. Using this analysis technique, we identify a limitation of the system when given a constant input and propose alternate formulations to fix the behavior. We show that the assistance inherently restricts the user from achieving all possible actions available to the system. In a free-space navigation setting, which corresponds to the scenario in many user studies, we find that the assistance introduces a level of indirection between the user and controls, which enables intuitive identification of equilibria but may decrease user opinion.

The analysis here is an example of what this technique may provide to the field of shared control, particularly for goal-predictive assistance. Since the user experiences the robot with or without assistance as a black-box input-output system, the dynamical analysis can be applied to many other types of shared control. In general, this analysis transforms shared control algorithm design into the inverse of a controls problem: given that somebody will be controlling over a system, design the system dynamics that best enables that operator to achieve their goals. We look forward to expanding this type of analysis to other shared control algorithms and using it to develop design principles for general assistive systems.

This analytical technique is not a substitute for other ways of designing shared control algorithms. Instead, its role is to enable an analysis of the assistance system outside of its assumptions. While assumptions about user behavior are the most clearly affected by user adaptation, this technique enables examining the user's experience in the case that the user is performing a different task than the assistance is modeling or has a conflicting environment model with the assistance. It even provides a framework for considering "adversarial" user behavior from a neutral viewpoint to determine how well the assistance performs outside its intended user behavior. Future work involves using this technique to develop general analysis tools for shared control systems that apply independently of many assumptions.

This paper presents a method for analyzing shared control systems and developing design principles for useful assistance. This user-centered, control-theoretic approach promises to enable systems that expand people's capabilities, the true goal of shared control.

ACKNOWLEDGMENT

This material is based upon work supported by the National Science Foundation under grant no. IIS-2132887.

REFERENCES

- [1] R. M. Aronson and E. S. Short, "Intentional User Adaptation to Shared Control Assistance," *ACM/IEEE International Conference on Human-Robot Interaction*, pp. 4–12, 3 2024. [Online]. Available: <https://dl.acm.org/doi/10.1145/3610977.3634953>
- [2] D. P. Losey, C. G. McDonald, E. Battaglia, and M. K. O'Malley, "A review of intent detection, arbitration, and communication aspects of shared control for physical human–robot interaction," 1 2018. [Online]. Available: http://asmedigitalcollection.asme.org/appliedmechanicsreviews/article-pdf/70/1/010804/5964415/amr_070_01_010804.pdf
- [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, 6 2018. [Online]. Available: <https://doi.org/10.1177/0278364918776060>
- [4] G. Cimolino and T. C. Graham, "Two Heads Are Better Than One: A Dimension Space for Unifying Human and Artificial Intelligence in Shared Control," in *Conference on Human Factors in Computing Systems - Proceedings*. Association for Computing Machinery, 4 2022. [Online]. Available: <https://dl.acm.org/doi/10.1145/3491102.3517610>
- [5] K. Muelling, A. Venkatraman, J.-S. S. Valois, J. E. Downey, J. Weiss, S. Javdani, M. Hebert, A. B. Schwartz, J. L. Collinger, and J. A. Bagnell, "Autonomy infused teleoperation with application to brain computer interface controlled manipulation," *Autonomous Robots*, vol. 41, no. 6, pp. 1401–1422, 8 2017. [Online]. Available: <https://doi.org/10.1007/s10514-017-9622-4>
- [6] D. Gopinath, M. N. Javaremi, and B. Argall, "Customized Handling of Unintended Interface Operation In Assistive Robots," in *Proceedings - IEEE International Conference on Robotics and Automation*, vol. 2021-May, 2021, pp. 10406–10412. [Online]. Available: https://github.com/argallab/interface_assessmentshttp://arxiv.org/abs/2007.02092
- [7] A. M. Okamura, "Methods for haptic feedback in teleoperated robot-assisted surgery," *Industrial Robot*, vol. 31, no. 6, pp. 499–508, 2004. [Online]. Available: <https://pubmed.ncbi.nlm.nih.gov/16429611/>
- [8] S. Jain and B. Argall, "Probabilistic Human Intent Recognition for Shared Autonomy in Assistive Robotics," *ACM Transactions on Human-Robot Interaction*, vol. 9, no. 1, pp. 1–23, 12 2019. [Online]. Available: <http://dl.acm.org/citation.cfm?doid=3375676.3359614>
- [9] C. L. Baker, J. B. Tenenbaum, and R. R. Saxe, "Goal Inference as Inverse Planning," in *Proceedings of the Annual Meeting of the Cognitive Science Society*, no. 29, 2007, p. 29.
- [10] B. D. Ziebart, A. Maas, J. A. Bagnell, and A. K. Dey, "Maximum Entropy Inverse Reinforcement Learning," in *Proceedings of the 23rd National Conference on Artificial Intelligence - Volume 3*, ser. AAAI'08. AAAI Press, 2008, p. 1433–1438.
- [11] A. Hüntemann, E. Demeester, M. Nuttin, and H. Van Brussel, "Online user modeling with Gaussian Processes for Bayesian plan recognition during power-wheelchair steering," in *2008 IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS*, 2008, pp. 285–292.
- [12] A. D. Dragan and S. S. Srinivasa, "A policy-blending formalism for shared control," in *International Journal of Robotics Research*, vol. 32, no. 7. SAGE PublicationsSage UK: London, England, 6 2013, pp. 790–805. [Online]. Available: <http://journals.sagepub.com/doi/10.1177/0278364913490324>
- [13] S. Reddy, A. D. Dragan, and S. Levine, "Shared Autonomy via Deep Reinforcement Learning," in *Robotics: Science and Systems*, 2018. [Online]. Available: <http://www.roboticsproceedings.org/rss14/p05.pdf>
- [14] H. Jun Jeon, D. Losey, and D. Sadigh, "Shared Autonomy with Learned Latent Actions," in *Robotics: Science and Systems*. Robotics: Science and Systems Foundation, 6 2020. [Online]. Available: <https://doi.org/10.15607/rss.2020.xvi.011>
- [15] R. M. Aronson and H. Admoni, "Gaze Complements Control Input for Goal Prediction During Assisted Teleoperation," in *Robotics: Science and Systems*, 2022.
- [16] P. Marayong, M. Li, A. M. Okamura, and G. D. Hager, "Spatial motion constraints: Theory and demonstrations for robot guidance using virtual fixtures," in *Proceedings - IEEE International Conference on Robotics and Automation*, vol. 2, 2003, pp. 1954–1959.
- [17] A. Broad, T. Murphey, and B. Argall, "Learning models for shared control of human-machine systems with unknown dynamics," in *Robotics: Science and Systems*, vol. 13, 2017. [Online]. Available: <https://storage.googleapis.com/rss2017-papers/50.pdf>
- [18] T. Yoneda, L. Sun, G. Yang, B. Stadie, and M. R. Walter, "To the Noise and Back: Diffusion for Shared Autonomy," *Robotics: Science and Systems*, 2023. [Online]. Available: <https://diffusion-for-shared-autonomy.github.io>.
- [19] D. J. Rea and S. H. Seo, "Still Not Solved: A Call for Renewed Focus on User-Centered Teleoperation Interfaces," *Frontiers in Robotics and AI*, vol. 9, 2022.
- [20] S. Nikolaidis, D. Hsu, and S. Srinivasa, "Human-robot mutual adaptation in collaborative tasks: Models and experiments," *International Journal of Robotics Research*, vol. 36, no. 5-7, pp. 618–634, 6 2017. [Online]. Available: <https://journals.sagepub.com/doi/full/10.1177/0278364917690593>
- [21] S. Parekh, S. Habibian, and D. P. Losey, "RILI: Robustly Influencing Latent Intent," in *IEEE International Conference on Intelligent Robots and Systems*, vol. 2022-Octob, 2022, pp. 2135–2142. [Online]. Available: <https://youtu.be/IYsWM8An18g>

- [22] D. E. Gopinath and B. D. Argall, "Mode switch assistance to maximize human intent disambiguation," in *Robotics: Science and Systems*, vol. 13, 2017. [Online]. Available: <http://www.roboticsproceedings.org/rss13/p46.pdf>
- [23] A. Jonnavittula and D. P. Losey, "Communicating Robot Conventions through Shared Autonomy," in *Proceedings - IEEE International Conference on Robotics and Automation*, 2022, pp. 7423–7429. [Online]. Available: <http://arxiv.org/abs/2202.11140>
- [24] M. Zurek, A. Bobu, D. S. Brown, and A. D. Dragan, "Situational Confidence Assistance for Lifelong Shared Autonomy," in *2021 IEEE International Conference on Robotics and Automation (ICRA)*, 2021, pp. 2783–2789.
- [25] B. A. Newman, C. Paxton, K. Kitani, and H. Admoni, "Bootstrapping Linear Models for Fast Online Adaptation in Human-Agent Collaboration," in *AAMAS '24: Proceedings of the 23rd International Conference on Autonomous Agents and Multiagent Systems*, 2024. [Online]. Available: <https://dl.acm.org/doi/10.5555/3635637.3663006>
- [26] S. Javdani, S. Srinivasa, and A. Bagnell, "Shared Autonomy via Hindsight Optimization," in *Robotics: Science and Systems XI*, vol. 11. Robotics: Science and Systems Foundation, 7 2015. [Online]. Available: <http://www.roboticsproceedings.org/rss11/p32.pdf>
- [27] B. A. Newman, R. M. Aronson, S. S. Srinivasa, K. Kitani, and H. Admoni, "HARMONIC: A Multimodal Dataset of Assistive Human-Robot Collaboration," *The International Journal of Robotics Research*, vol. 41, no. 1, pp. 3–11, 7 2022. [Online]. Available: <http://journals.sagepub.com/doi/10.1177/02783649211050677>