Probabilistic Human Intent Recognition for Shared Autonomy in Assistive Robotics

SIDDARTH JAIN and BRENNAN ARGALL, Northwestern University and Shirley Ryan AbilityLab

2

Effective human-robot collaboration in shared autonomy requires reasoning about the intentions of the human partner. To provide meaningful assistance, the autonomy has to first correctly predict, or infer, the intended goal of the human collaborator. In this work, we present a mathematical formulation for intent inference during assistive teleoperation under shared autonomy. Our recursive Bayesian filtering approach models and fuses multiple non-verbal observations to probabilistically reason about the intended goal of the user without explicit communication. In addition to contextual observations, we model and incorporate the human agent's behavior as goal-directed actions with adjustable rationality to inform intent recognition. Furthermore, we introduce a user-customized optimization of this adjustable rationality to achieve user personalization. We validate our approach with a human subjects study that evaluates intent inference performance under a variety of goal scenarios and tasks. Importantly, the studies are performed using multiple control interfaces that are typically available to users in the assistive domain, which differ in the continuity and dimensionality of the issued control signals. The implications of the control interface limitations on intent inference are analyzed. The study results show that our approach in many scenarios outperforms existing solutions for intent inference in assistive teleoperation and otherwise performs comparably. Our findings demonstrate the benefit of probabilistic modeling and the incorporation of human agent behavior as goal-directed actions where the adjustable rationality model is user customized. Results further show that the underlying intent inference approach directly affects shared autonomy performance, as do control interface limitations.

CCS Concepts: • **Mathematics of computing** → **Probabilistic inference problems**; *Probabilistic algorithms*; **Bayesian computation**; • **Computing methodologies** → *Theory of mind*; • **Computer systems organization** → **Robotics**; • **Human-centered computing** → *Human computer interaction (HCI)*;

Additional Key Words and Phrases: Human intent recognition, intent inference, shared autonomy, probabilistic modeling, assistive robotics, assistive teleoperation, human-robot interaction

ACM Reference format:

Siddarth Jain and Brenna Argall. 2019. Probabilistic Human Intent Recognition for Shared Autonomy in Assistive Robotics. *ACM Trans. Hum.-Robot Interact.* 9, 1, Article 2 (December 2019), 23 pages. <https://doi.org/10.1145/3359614>

This material is based upon work carried out at Northwestern University and the Shirley Ryan AbilityLab. Research reported in this publication was supported by the National Institute of Biomedical Imaging and Bioengineering (NIBIB) and National Institute of Child and Human Development (NICHD) under award numbers R01EB019335 and R01EB024058. The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Institutes of Health. This work was also supported by grant from U.S. Office of Naval Research under the Award Number N00014-16-1-2247.

Authors' addresses: S. Jain, Mitsubishi Electric Research Labs (MERL), Cambridge, MA 02139; email: sjain@merl.com; B. Argall, Northwestern University, 2145 Sheridan Road, Evanston, IL 60208, and Shirley Ryan AbilityLab, 355 East Erie Street, Chicago, IL 60611; email: brenna.argall@northwestern.edu.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2019 Copyright held by the owner/author(s). Publication rights licensed to ACM.

2573-9522/2019/12-ART2

<https://doi.org/10.1145/3359614>

1 INTRODUCTION

Human-robot teams are common in many application areas, such as domestic service, search and rescue, surgery, and driving vehicles. In the area of assistive robotics, shared autonomy is utilized to aid the human in the operation of an assistive robot through the augmentation of human control signals with robotics autonomy [13, 19, 23, 30]. One fundamental requirement for effective human-robot collaboration in shared autonomy is *human intent recognition*. To meaningfully assist the human collaborator, the robot has to correctly infer the intended goal of the user from a number of potential task-relevant goals—known as the intent inference problem (Figure 1).

Humans are very good at anticipating the intentions of others with non-verbal communication [11]. Research studies in experimental psychology show that humans demonstrate a strong and early inclination to interpret observed behaviors of others as goal-directed actions [7]. Evidence shows that even human infants are inclined to interpret certain types of human behaviors as goal-directed actions. For example, a developmental proposition to the question of how the child develops a theory of mind posits that infants learn about the signal value of adult actions because these behaviors predict the locations of interesting objects and events [28]. Such a teleological interpretation of actions in terms of goals could inform us why the action has been performed—that is, it provides a special type of explanation for the action.

Taking inspiration from psychology and teleological interpretations of actions, we aim to explore how the robotics autonomy can implicitly observe and leverage non-verbal human actions to perform human intent recognition. We focus specifically on human commands provided to a shared-control assistive robotic system through a control interface. Such control command information is readily available, and the operation of these control interfaces is already familiar to end users and does not divert attention from the task execution. Directly leveraging the user's interaction with the control interface to perform intent recognition could pose critical challenges as a consequence of the limitations of the interface or due to the underlying motor impairment. We therefore model the user's interaction with the control interface as probabilistic goal-directed actions to allow for uncertainty in reasoning about the user's intention.

We consider the user control inputs as representative of the action the user wants to take and build a probabilistic behavior model with adjustable rationality to inform human intent recognition. We furthermore introduce a user-customized optimization of the adjustable rationality for user personalization, which admits for suboptimality in the human actions due to any number of reasons (discussed further in Section 3.5). For intent recognition, we model the uncertainty over the user's goal within a recursive Bayesian filtering framework. Our framework allows for the continuous update of the intent probability for each goal hypothesis based on the available observations. Furthermore, it enables the seamless fusion and incorporation of any number of observation sources, allowing the intent inference to leverage rich sources of information. Finally, the framework can explicitly express uncertainty in the resulting inference, which is critical to know in shared autonomy as assistance toward wrong goals could be worse than providing no assistance.¹

In this article, we perform a variety of user studies in different scenarios to evaluate intent recognition and its affect on shared autonomy performance. Our experimental work furthermore evaluates intent inference during novice teleoperation and under shared autonomy assistance,

¹This work is a continuation of the work presented in our prior conference paper [18]. In current work, we extend the intent inference evaluation across multiple interfaces that typically are available to people with motor impairments, include end users in addition to subjects without motor impairments, and provide a more extensive and detailed evaluation of intent inference. Furthermore, we optimize the adjustable rationality for each user and examine the benefit of the individualized adjustable rationality model of the human actions for inference.

Fig. 1. Schematic of human intent recognition for shared autonomy.

and using multiple control interfaces typically employed within the assistive domain, which is first within the field.

In summary, our work makes the following contributions:

- It introduces a user-customized optimization of adjustable rationality in a probabilistic model of human actions as goal-directed observations for performing intent inference, which thus can account for human action suboptimality within the assistive domain. We demonstrate the benefit of individualized adjustable rationality for intent inference.
- It presents a mathematical framework for human intent recognition that enables the seamless fusion of observations from any number of information sources to probabilistically reason about the intended goal of the user.
- It examines intent inference under multiple control interfaces that typically are utilized within the assistive domain—a first within the field.
- It provides a more extensive evaluation of intent inference than is typically seen within the literature, including a larger set of candidate goals, trials with change of intentions, an examination of the impact on shared autonomy performance, comparison to human inference baselines, and evaluation by assistive technology end users.

The remainder of the article is organized as follows. Related work is discussed in Section 2. Section 3 casts intent inference within a recursive Bayesian filtering framework and presents our algorithm. Implementation details are provided in Section 4. Details of experimental work are presented in Section 5 and discussed further in Section 6 with results. In Section 7, we present the discussion and then conclude with directions for future research in Section 8.

2 RELATED WORK

Intent inference—also referred to as inference of the desired goal, target, action, or behavior—has been investigated under various settings [26]. For instance, intent inference includes human activity recognition in the area of computer vision [1] using spatio-temporal representations and task executions in surgical telemanipulation systems using hidden Markov models [25].

Probabilistic methods often are used to infer unknown intentions from human movements. A latent variable model is proposed to infer the human intentions from an ongoing movement, which is verified for target inference in robot table tennis and movement recognition for interactive humanoid robots [35]. A data-driven approach synthesizes anticipatory knowledge of human motions and subsequent action steps for the prediction of human reaching motion

toward a target [31]. Predictive inverse optimal control methods also are utilized to estimate human motion trajectories and shown to work for anticipating the target locations and activity intentions of human hand motions [27].

Instead of human movements and motions, in this work we investigate how shared autonomy systems can take advantage of the indirect signals people implicitly provide when operating a robot through a direct control interface (e.g., a joystick), as well as the implications of the control interface limitations on intent inference. The commercial control interfaces that are accessible to those with motor impairments are limited in both the dimensionality and continuity of the control signals. Examples include 2- or 3D control joysticks, and more limited interfaces such as Head Arrays and Sip-N-Puff control interfaces (1D control). For all, the dimensionality of the control signal is considerably lower than the control space of the robot—making the intent inference problem more challenging.

In shared autonomy, many systems simplify the intent inference problem by focusing on predefined tasks and behaviors or by assuming that the robot has access to the user's intent *a priori* [19, 23]. One approach to infer the user's intent could be to have the user communicate the intended goal explicitly, such as via verbal commands as observations. There exist works that rely on explicit commands from the user to communicate intent via an interface, such as a GUI [34]. However requiring explicit communication from the user could lead to ineffective collaboration and increased cognitive load [12]. To infer intentions and maintain a common understanding of the shared goal during human-robot collaborative tasks, a number of human-robot interaction approaches investigate the use of non-verbal communication including gestures, expressions, and gaze. For example, user gaze patterns are studied to understand human control behavior for shared manipulation [3] and to perform anticipatory actions [14]. The eye gaze data quality can depend on the calibration accuracy and projection method errors of the eye-tracker [14]. Gestures, along with language referential expressions, are utilized for object inference in human-robot interactions [36].

Intent inference in shared autonomy is often represented as a confidence in the prediction of the intended goal, based on instantaneous observations [2, 9, 13]. Some shared autonomy systems compute a belief over the space of possible goals. A memory-based inference approach utilizes the history of trajectory inputs and applies Laplace's method about the optimal trajectory between any two points to approximate the distribution over goals [9]. Memory-based inference is utilized in previous works involving shared autonomy and human-robot systems [8, 9, 30]. Another approach uses Laplace's approximation [9] and formulates the problem as one of optimizing a partially observable Markov decision process (POMDP) over the user's goal to arbitrate control over a distribution of possible outcomes [20]. The approach considers user inputs for the prediction model and uses a hand-specified distance-based user cost function to achieve a closed-form value function computation.

In this work, we consider probabilistic modeling of implicit observations and compute a belief over the space of possible goals within a Bayesian filtering framework. In addition to a distance-based likelihood observation used in the prior works, our approach incorporates a fusion of multiple observations and introduces a probabilistic modeling of user control inputs as goal-directed actions that customizes the rationality index value to each individual user, and thus can account for their particular behavior. The motivation for our approach are prior studies that show that users vary greatly in their performance, preferences, and desires [9, 20, 37]—suggesting a need for assistive systems to customize to individual users.

Finally, in the majority of the literature on shared-control systems that perform intent inference, the focus of the experimental work is on the control sharing and robot assistance [2, 13, 21, 30]. The intent inference itself is mostly assessed only implicitly and seldom the central aim of a study. Inferred intent, however, can have critical impact on the shared autonomy performance: incorrect

predictions can result in unfortunate circumstances including collisions, whereas imprecise predictions impact the robot's ability to correctly assist the user toward his or her intended goal. This in turn can result in large penalties in time and user preference on assistance [9]. In this article, we therefore present a more extensive evaluation of intent inference than is typically seen within the literature.

3 FRAMEWORK FOR INTENT INFERENCE UNDER SHARED AUTONOMY

We first mathematically define the intent inference problem and then present our framework (Algorithm 1). Our target domain is assistive robots endowed with shared autonomy that assist the user towards his or her intended goal.

3.1 Problem Formulation

Let the environment have a discrete set of accessible goals² $\mathcal{g}$, known at runtime to both the user and robot. The human's intended goal is unknown to the robot, and so the robot must infer (predict) the most likely goal $g^* \in \mathcal{g}$ that the user is trying to reach: the intent inference problem. With knowledge of g^* , the robotics autonomy can meaningfully assist the user in shared autonomy. Note that the user might change his or her intended goal during the execution, and intent inference that updates in real time will allow the assistance provided to dynamically adjust.

3.2 Intent Estimation

We formulate the intent inference problem for shared autonomy as Bayesian filtering in a Markov model, which allows us to model the uncertainty over the candidate goals as a probability distribution over the goals. Bayesian models are shown to be effective for inference in cognitive science [4] and human-robot interaction research [36]. We cast the intent inference problem as a classification task where the robot aims to infer the most likely goal class g^* from the set of possible goals $\mathcal{g}$, given a set of observations (features).

We represent the goal g_t as the query variable (hypothesis) and the observed features $\Theta_0, \dots, \Theta_t$ as the evidence variables, where Θ_t is a k -dimensional vector of k observations θ_t^i , $i = 1 : k$, and t represents the current time. For compactness, we use colon notation to write $\Theta_1, \dots, \Theta_t$ as $\Theta_{0:t}$. The uncertainty over goals is then represented as the probability of each goal hypothesis. The goal probability conditioned on a single observation source θ^i over t timesteps can be represented by Bayes' rule as

$$b_t(g) = P(g_t | \theta_{0:t}) \propto P(\theta_t | g_t, \theta_{0:t-1})P(g_t | \theta_{0:t-1}), \quad (1)$$

where the superscript i has been dropped for notational simplicity, and the posterior probability $P(g_t | \theta_{0:t})$ at time t represents the belief $b_t(g)$ after taking the single observation source into account, where $b_t(g)$ is a single element of the posterior distribution b_t . Since the hidden Markov model allows for a conditional independence assumption between observations at the current and previous timesteps given the current goal estimate ($\theta_t \perp \theta_{0:t-1} | g_t$), we simplify $P(\theta_t | g_t, \theta_{0:t-1})$ to $P(\theta_t | g_t)$. Applying the law of total probability, the conditional goal probability becomes

$$b_t(g) \propto P(\theta_t | g_t) \sum_{g_{t-1} \in \mathcal{g}} P(g_t, g_{t-1} | \theta_{0:t-1}), \quad (2)$$

which becomes

$$b_t(g) \propto P(\theta_t | g_t) \sum_{g_{t-1} \in \mathcal{g}} P(g_t | g_{t-1})b_{t-1}(g_{t-1}) \quad (3)$$

²Potential goals could be detected by perception methods such as recognizing objects clusters/grasps [17] or saliency maps [6] for manipulation and navigation goals such as doorways [32] and docking locations [16].

ALGORITHM 1: Recursive Bayesian Intent Inference

Given Goals $\mathbf{g}$
Initialize $P(g_{t=0}) \forall g \in \mathbf{g}$
Initialize $b_{t=0}(g) \leftarrow P(\Theta_{t=0} \mid g_{t=0})P(g_{t=0}) \forall g \in \mathbf{g}$
Normalize $b_{t=0}$
while *executing* **do**
 Observe Θ_t
 foreach $g \in \mathbf{g}$ **do**
 $b_t(g) \leftarrow \prod_{\theta_t \in \Theta_t} P(\theta_t \mid g_t) \sum_{g_{t-1} \in \mathbf{g}} P(g_t \mid g_{t-1})b_{t-1}(g_{t-1})$
 end
 Normalize b_t
 Update intent $g_t^* \leftarrow \arg \max_{g \in \mathbf{g}} b_t(g)$
end

upon applying the definition of conditional probability to $P(g_t, g_{t-1} \mid \theta_{0:t-1})$ and the Markov assumption, $(g_t \perp \theta_{0:t-1} \mid g_{t-1})$. The computation of $b_t(g)$ thus is a recursive update and so encodes memory of prior goal distributions. Furthermore, $P(g_t \mid g_{t-1})$ is the conditional transition distribution of changing to goal g_t at time t given that the goal was g_{t-1} at time $t - 1$. The model thus encodes that the user's intent or goal can change over time.

We now take into consideration multiple observation sources $\theta^1, \dots, \theta^k$ as k evidence variables, which can derive from any number of sources—for example, control commands or cues such as eye gaze. We assume the k observations sources to be conditionally independent of each other given a goal g , $(\theta^i \perp \theta^j \mid g)$, $\forall i \neq j$. Thus, Equation (3) becomes

$$b_t(g) \propto \prod_{\theta_t \in \Theta_t} P(\theta_t \mid g_t) \sum_{g_{t-1} \in \mathbf{g}} P(g_t \mid g_{t-1})b_{t-1}(g_{t-1}). \quad (4)$$

We present our algorithm for inferring the probability distribution over goals in Algorithm 1. The posterior distribution at time t , denoted b_t , represents the belief after taking observations into account. The set of prior probabilities at $t = 0$, $P(g_{t=0})$, $\forall g \in \mathbf{g}$, represents the robot's initial belief over the goals. The beliefs then are continuously updated, by computing the posteriors $P(g_t \mid \Theta_{0:t})$, $\forall g \in \mathbf{g}$, as more observations become available.

Finally, to predict the most likely goal $g_t^* \in \mathbf{g}$, we select the goal class that is most probable according to the *maximum a posteriori* decision,

$$g_t^* = \arg \max_{g_t \in \mathbf{g}} P(g_t \mid \Theta_t). \quad (5)$$

3.3 Inference Uncertainty

Within the domain of shared-control assistive robotics, it is important to express uncertainty in the robot's prediction of the intended goal—because assisting toward the wrong goal could be worse than providing no assistance. We express prediction uncertainty as a *confidence* computed as the difference between the probability of the most probable and second most probable goals,

$$C(\mathbf{g}) = P(g^* \mid \Theta) - \arg \max_{g \in \mathbf{g} \setminus g^*} P(g \mid \Theta). \quad (6)$$

When the robot is uncertain about the intended goal of the user, a variety of behaviors might be implemented, such as to hold off on providing assistance or assist toward multiple goals simultaneously, if possible.

3.4 Human Behavior Model as an Observation

We are interested in investigating how we can utilize implicit intent information encoded in the signals people provide to operate a robot. Taking inspiration from psychology and teleological interpretations of actions [7], we aim to explore how the robotics autonomy can implicitly observe and leverage non-verbal human actions to perform human intent recognition. Specifically, we model the user's interaction with the control interface that he or she uses to operate the robot as probabilistic goal-directed actions to reason about the user's intention. We consider the user inputs as representative of the actions the user wants to take to reach a goal g . We model these actions as observations using Boltzmann rationality, which has been shown to explain human behavior on various datasets [15].

We incorporate adjustable rationality in a probabilistic behavior model such that in any state s (robot configuration) the probability that action u_g is chosen by a rational human agent to reach goal g is given as

$$P(u_g | s, g) \propto \exp(\beta \cdot Q_g(s, u_g)), \quad (7)$$

where $Q_g(s, u_g)$ denotes the Q-value when the intended goal is g . β is a rationality index (discussed further in Section 3.5) that controls how diffuse the probabilities are. We model $Q_g(s, u_g)$ as the cost of taking action u_g at configuration s and acting optimally from that point on to reach the goal g . We approximate optimal action selection with an autonomy policy and compute this cost as the *agreement* between the user control u_h and the autonomy control u_r to reach the goal g . In our implementation, a policy based on potential fields [22] is employed, and agreement is measured in terms of the cosine similarity, computed as $\arccos((u_h \cdot u_r) / \|u_h\| \|u_r\|)$, where $\cdot$ denotes the dot product and $\|\cdot\|$ denotes the vector norm. A moving average filter with a 2-second time window is applied to the observations to consider a brief history of observations and reduce any undesired oscillations due to noisy or corrective control signals.

3.5 User-Customized Optimization for Personalization

A perfectly rational model assumes that a human agent always acts optimally to reach his or her goal, when in reality a number of factors might induce suboptimality—for example, limitations and challenges imposed by the lower degree of freedom (DoF) interfaces (e.g., 3-axis joystick, Sip-N-Puff, Head Array) available to control high-DoF robotic systems, cognitive impairments, physical impairments, and environment factors such as obstacles or distractions. Furthermore, not only do users differ in their abilities but also in the manner in which they interact with the robot. Therefore, a critical detail is to include adjustable rationality in the human action model for user personalization and adaptation.

We introduce a user-customized optimization of adjustable rationality that customizes the value of β to each user. Adjustable rationality is represented by the rationality index parameter β in the Boltzmann policy model (Equation (7)). β reflects the robot's belief about the optimality of the human agent actions. By tuning β , the robot can account for human agents who may behave suboptimally in their actions and thus improve inference from their actions. To find the value of the rationality index for a human agent, we perform an optimization procedure on data gathered while the human teleoperates the robot to reach multiple goals in the environment. In particular, we find the value β^* that minimizes the average log-loss for goal inference across dataset trajectories given the true labels Y for the intended goal and the probability estimates P for the inference over goals,

$$\beta^* = \arg \min_{\beta} \mathcal{L}_{\log}(Y, P, \beta), \quad (8)$$

and the log-loss is computed as

$$\mathcal{L}_{\log}(Y, P, \beta) = -\frac{1}{M} \sum_{i=1}^M \sum_{j=1}^N y_{i,j} \log p_{i,j}(g_i | \Theta, \beta), \quad (9)$$

where M is the number of samples in the trajectory, N is the number of goals, $y_{i,j}$ is a binary indicator of whether or not the prediction j is the correct classification for goal instance i , and $p_{i,j}$ is the probability associated with the goal j at timestep i . The value of β^* returned by the optimization procedure is then used to model the human actions using Equation (7), and $P(\mathbf{u}_g | \mathbf{s}, g)$ are used as observations in the inference predictions of Equation (4).

4 IMPLEMENTATION

In this section, we detail our implementations of our intent inference algorithm and the shared autonomy that utilizes the inferred intent to provide assistance.

4.1 Autonomy Inference

In shared autonomy, oftentimes the only cues available to the system comes from the input commands received from the human user (as discussed earlier) and via sensing the robot's environment. Our approach, recursive Bayesian intent inference (RBII), allows for the seamless fusion of any number of observations to perform human intent recognition. To examine how the incorporation of different observation sources affects the intent inference and shared autonomy, we implement two different observation schemes (RBII-1 and RBII-2).

RBII-1. The first observation scheme considers a single modality, the *proximity* to a goal, as this feature is utilized most in existing shared autonomy work [9, 13, 21, 30]. We compute proximity θ^d as the Euclidean distance between the current position of the robot end-effector $\mathbf{x}$ and the goal $\mathbf{x}_g$. For Algorithm 1, we model the likelihood using the principle of maximum entropy such that given the goal g , the class conditional probability decreases exponentially as the likelihood of g decreases, $P(\theta^d | g) \propto \exp(-\kappa \cdot \theta^d)$. κ is set to the mean of the range of values that θ^d can take.

RBII-2. In the second observation scheme, we consider a probabilistic fusion of multiple observations. Specifically, in addition to proximity to the goal, we model the actions of the human agent as observations. Following a model of human action from cognitive science [4], we model the user as Boltzmann rational in his or her actions to reach a goal g (Section 3.4) and furthermore perform an optimization to customize the value of adjustable rationality for each user (Section 3.5). In our implementation, the global optimum is approximated using a probabilistic optimization algorithm (simulated annealing [24]). The individualized value of adjustable rationality returned by the optimization procedure is used to model the human actions using Equation (7) and $P(\mathbf{u}_g | \mathbf{s}, g)$ are used as observations by our algorithm to make inference predictions.

Last, our approach encodes the possibility that the human's goal might change during task execution (Section 3.2). We compute the probability of changing goals in the case of n number of goals as

$$P(g_t = g^i | g_{t-1} = g^j) = \begin{cases} 1 - \Delta & \text{if } i = j \\ \Delta/(n - 1) & \text{otherwise.} \end{cases} \quad (10)$$

Note that when $\Delta = 0$, the model represents the case when the user exclusively pursues one goal during the execution. When $\Delta = (n - 1)/n$, the model represents the possibility of choosing a new goal at random at each timestep. Our implementation sets $\Delta = 0.1$ and initializes the probability distribution over goals to be uniform.

4.2 Assistance Under Shared Autonomy

In this section, we discuss implementation details of the shared autonomy that utilizes the inferred intent. Blending-based shared autonomy methods combine the user and autonomy control of the robot by some arbitration function that determines the relative contribution of each. Although not optimal for operation in complex safety-concerned environments [33], blending is one of the most used shared autonomy paradigms due to its computational efficiency and effectiveness [9, 13, 29]. Furthermore, it has been shown in subjective evaluations that users tend to prefer blending and find it easier to learn over some other approaches of shared autonomy, for example, as compared to probabilistic shared control (PSC) [10] and a POMDP-based policy method [20]. We implement a linear blending-based paradigm to provide assistance,

$$\mathbf{u}_{blend} = \mathbf{u}_h \cdot (1 - \alpha) + \mathbf{u}_r \cdot \alpha, \quad (11)$$

where $\mathbf{u}_h$ denotes the user control command, $\mathbf{u}_r$ the autonomy control command generated under a potential field policy [22], and $\mathbf{u}_{blend}$ the shared autonomy command sent to the robot. Note that $\mathbf{u}_r$ is available in all parts of the robot state space, for every goal $g \in \mathcal{g}$ such that g is treated as an attractor and all other goals $\mathcal{g} \setminus g$ as repellers. $\alpha \in [0, 1]$ is a blending factor that arbitrates how much control remains with the human user versus the autonomy. In our implementation, α is a piecewise linear function of the confidence in the intent prediction,

$$\alpha = \begin{cases} 0 & C(g) \leq \delta_1 \\ \frac{\delta_3}{(\delta_2 - \delta_1)} \cdot C(g) & \delta_1 < C(g) \leq \delta_2, \\ \delta_3 & C(g) > \delta_2 \end{cases} \quad (12)$$

where $C(g)$ is defined as in Equation (6), the difference between the highest and second highest probable goals. δ_1 is a lower bound on $C(g)$ (set to 30%), below which assistance is not active, and δ_2 is an upper bound on $C(g)$ (set to 90%), above which assistance is maximal. The upper bound on assistance α is given by δ_3 (set to 70%). Note the following. First, different approaches to intent inference will generate different values for $C(g)$, and so the amount of assistance accordingly will differ. Second, in particular, if $C(g)$ is lower, meaning that the inference is not very certain in its prediction, the amount of assistance also will be lower. Third, if the inferred goal is wrong, the robot will assist toward the wrong goal, with potentially serious implications.

5 EXPERIMENTS ON INTENT INFERENCE

Our experimental work aims to evaluate the performance of the intent inference algorithm, as well as its impact on shared autonomy. We furthermore are interested in how the control interface used by the human affects the intent inference performance. This is particularly important in the domain of assistive robotics, as traditional joystick control interfaces are often not accessible to people with severe motor impairments who instead use limited control interfaces to operate the robot. In this work, we consider the 3-axis joystick and the 1D Sip-N-Puff as the representative assistive domain input interfaces. We performed two human subject studies, that aim (1) to characterize the complexity and variability of the intent inference problem, (2) to compare the inference performance of our approach to existing approaches utilized in prior shared autonomy work, (3) to evaluate the impact of interface on intent inference, and (4) to evaluate the impact of inference on shared autonomy assistance.

5.1 Hardware

Robot Platform. Our research platform for the designed experiments was the MICO robotic arm (Kinova Robotics, Canada), a 6-DoF manipulator with a two-finger gripper (Figure 2). The MICO is

MICO robot arm

6-DoF

Control Interfaces and Control Modes

3-axis Joystick

- Translational Mode
- Wrist Mode
- Finger Mode

Sip-N-Puff

- Forward / Backward
- Left / Right
- Up / Down
- Lateral Orientation
- Vertical Orientation
- Wrist Orientation
- Finger

Fig. 2. MICO robotic arm and the control interfaces used in the study, with illustration of the control modes for the robotic arm teleoperation.

a research edition of the JACO arm (Kinova Robotics, Canada), which is used commercially within the assistive domain.

Control Interfaces. The control interfaces used in the study were (1) a traditional 3-axis joystick (Kinova Robotics, Canada) that have typically been utilized for operating robotic arms, and (2) a 1D Sip-N-Puff interface (Origin Instruments, United States) accessible to people with severe motor impairments. Direct teleoperation using control interfaces, which are lower dimensional than the control space of the robotic arm, requires the user to switch between one of several control modes (mode switching). Modes partition the controllable DoFs of the robot such that each control mode maps the input dimensions to a subset of the arm's controllable DoFs. The control interfaces and defined control modes are shown in Figure 2. Teleoperation with the 3-axis joystick requires at minimum three control modes, and the required number of control modes increases to six in the case of the Sip-N-Puff interface.

5.2 Software

In addition to the two variants of our algorithm detailed in Section 4.1 (RBII-1 and RBII-2), for comparative purposes we also implement two approaches to intent inference utilized in previous shared autonomy works—amnesic inference [2, 9, 13] and memory-based inference [8, 9, 20, 30].

Amnesic Inference. The amnesic inference approach associates a confidence in the prediction of the user's goal as a hinge-loss function, where it is assumed that the closer a goal is, the more likely it is the intended goal,

$$c(g) = \max\left(0, 1 - \frac{d}{D}\right), \quad (13)$$

where d is the distance to the goal and D a threshold past which the confidence $c(g)$ is 0. It is possible to design richer confidence functions, but in practice this function often is used for its simplicity. The approach is termed as *amnesic prediction* [9] because it ignores all information except the instantaneous observations. In our implementation, d is the Euclidean distance $\|\mathbf{x}_g - \mathbf{x}_r\|$ between the current position of robot $\mathbf{x}_r$ and the goal $\mathbf{x}_g$, and D is set to 1.0 m (maximum reach of the MICO robot arm).

Memory-Based Inference. The memory-based prediction [9] approach is a Bayesian formulation that takes into consideration the history of a trajectory to predict the most likely goal. Let $\xi_{x \rightarrow y}$ denote a trajectory starting at pose x and ending at y . Using the principle of maximum entropy [38], the probability of a trajectory for a specific goal g is given as $P(\xi \mid g) \propto \exp(-c_g(\xi))$ —that is, the

Fig. 3. *Left*: Goal scenarios of varied complexity used in human inference study. *Right*: Example setup shows the robot executing a trajectory to reach a goal. The subject predicts the intended goal during robot motion.

probability of the trajectory decreases exponentially with cost. It is assumed that the cost is additive along the trajectory.

$$P(\xi_{s \rightarrow x} | g) = \frac{\exp(-c_g(\xi_{s \rightarrow x})) \int_{\xi_{x \rightarrow g}} \exp(-c_g(\xi_{x \rightarrow g}))}{\int_{\xi_{s \rightarrow g}} \exp(-c_g(\xi_{s \rightarrow g}))} \quad (14)$$

Such a solution becomes too expensive to compute in high-dimensional spaces (e.g., for robotic manipulation) and so [9] estimates the most likely goal by approximating the integral over trajectories using Laplace's method and a first-order approximation,

$$g^* = \arg \max_{g \in \mathcal{G}} \frac{\exp(-c_g(\xi_{s \rightarrow x}) - c_g(\xi_{x \rightarrow g}^*))}{\exp(-c_g(\xi_{s \rightarrow g}^*))} P(g). \quad (15)$$

In practice, the cost c_g is often the Euclidean distance and the goal probabilities are initialized with a uniform prior [21, 30], which we do in our implementation as well.

5.3 Human Inference Study

We first aim to ground the complexity of the intent inference problem through a study of human inference ability by having a human observer interpret the motion of a robotic arm to infer its intended goal. Humans are very good at anticipating the intentions of others with non-verbal communication [11] and by observation of goal-directed actions [7].

In our experiments, the robot motion trajectories were prerecorded demonstrations by a human expert operating a 3-axis joystick to reach a goal. Related work has addressed the topic of human interpretation of robot motion trajectories generated by autonomy policies [8] with the aim to generate intent inexpressive trajectories. We, however, are interested in robot motion generated by *human* control commands, given our target domain of assistive teleoperation. We study human inference under a variety of goal scenarios of varied complexity (Figure 3, left).

We furthermore are interested in how a *change of intent* can affect the inference. Our human demonstrator therefore provided two types of motion: (1) *no change of intent*, where the robot motion maintained a single goal from start to end, and (2) *change of intent*, where the robot motion switched during execution of the goal that it was reaching toward.

Design. The study involved a variety of goal scenarios with varied complexity: in total, eight different scenarios involving two to four potential goals (Figure 3, left). For each goal scenario, one change of intent demonstration was recorded (eight in total), as well as no change of intent

trajectories for each goal in the scene (22 trajectories in total). We chose a within-subjects design, and each subject observed the 30 recorded demonstrated trajectories replayed on the robot.

Subject Allocation. We recruited 12 subjects without motor impairments from the local community (five male, seven female, age 19 to 35 years). The subjects were novice users and had no prior experience operating a robotic arm. All participants gave their informed signed consent to participate in the studies, which was approved by the Northwestern University Institutional Review Board.

Protocol. At the start of the experiment, the subjects were given a 5-minute training period in which they observed 1 prerecorded human teleoperation trajectory toward a random goal in the environment to become familiarized with the robot motion capabilities. The 30 teleoperated trajectories for the eight goal scenarios then were executed (replayed) in counterbalanced order on the robot. For each trajectory motion, we tasked the human subjects to observe the motion of the robot and predict which object was the intended goal by verbally mentioning the “Object Name” (e.g., “cup”) or “Number” (e.g., “one”). The subjects were allowed to change their inference at any time during the robot motion. They were also given the option to express *uncertainty* about the intended goal by saying “not sure” or “zero.” Note that by default the inference is registered as uncertain (“zero”) at the start of the trajectory motion. The subject inferences were recorded by the experimenter via a button press. Figure 3 (right) shows the experimental setup for the study.

5.4 Autonomy Inference Study

Our second study aims to evaluate how well the autonomy can infer the intent of a user under *novice teleoperation* and how the inferred intent affects *assistance* in shared autonomy.

Design. We adopted a between-subjects experimental design with two conditions that differed in the control interface type used by the subjects to operate the robot: a 3-axis joystick and a 1D Sip-N-Puff interface. All other elements were the same in the two conditions. Four intent inference approaches were evaluated: amnesic inference, memory-based inference, RBII-1, and RBII-2. Each computed the inference online as the subject teleoperated the robot to complete tasks. Four different tasks of varied complexity were used to evaluate the intent inference performance (Figure 4). The first three tasks involved object retrieval, where the goals were the approach grasp pose $\in \mathbb{R}^6$ on the objects. The fourth task involved pouring and placing operations, where the goals were tied to the manipulation operations: the pose of the initiation of the pour over the bowls and placing in the dish rack. Fetching objects and manipulation for meal preparation have been identified as the top preferred ADL tasks [5] by end users.

Subject Allocation. We recruited 24 subjects without motor impairments (uninjured) from the local community (11 male, 13 female, age 19 to 40 years) and six end-user subjects (5 male, 1 female, age 18 to 47 years). Twelve of the subjects without motor impairments (5 males, 2 females) and three of the end user subjects (2 male, 1 female) used the 3-axis joystick interface, and the remaining subjects used the Sip-N-Puff interface. The end user group included four subjects with spinal cord injury (SCI), one with muscular dystrophy (MD), and one with cerebral palsy (CP). All participants gave their informed signed consent to participate in the studies, which was approved by the Northwestern University Institutional Review Board.

Protocol. At the start, the subjects were given a 10-minute training period in which they became familiar with using the control interface and the robot operation. Next, they teleoperated the robot to perform a training task (a similar setup as shown in task 3 in Figure 4 but with different positions and orientations of the objects), and this data was used to perform the optimization procedure to

Fig. 4. Task scenarios and associated grasp pose on goals used to evaluate intent inference under novice teleoperation (with and without shared autonomy). Note that for task 3, the two goals on the extreme have the same grasp pose as shown in task 2.

tune the rationality index parameter β for each subject (as discussed in Section 3.5). The optimized value β^* was used in the RBII-2 approach for intent inference.

There were two stages in the study protocol. In the first stage, the subjects teleoperated the robot *without assistance*. The subjects performed the tasks shown in Figure 4, and they were instructed to (1) complete each goal in every task setup and (2) for each task perform one additional trial in which they *changed the goal* during the task execution. The change of goal was recorded with a timestamp via a button press. All four approaches for intent inference (discussed in Sections 4.1 and 5.2) computed the inference online as the user teleoperated the robot to complete tasks. In the second stage, all previous trials were performed again but now *with assistance* under shared autonomy. Furthermore, stage 2 trials were performed twice, once with the RBII-1 approach and once with the RBII-2 approach.

6 ANALYSIS AND RESULTS

We first discuss performance measures and then present the analysis with the results. For each performance measure, one factor repeated measure analysis of variance (ANOVA) was performed to determine significant differences ($p < .05$). Once the significance was established, multiple post hoc pairwise comparisons were performed by using Bonferroni confidence interval adjustments. For all figures, the notation * implies $p < .05$, ** implies $p < .01$, and *** implies $p < .001$.

6.1 Performance Measures

Percentage of Correct Predictions. Percentage of correct predictions is a metric commonly employed in machine learning works. We compute percentage correct as the percentage of time (success rate) the inference identified the correct intended goal of the user with confidence ($C(g) > 30\%$).

Cross-Entropy Loss (Log-Loss). Assessing the uncertainty of a prediction is an important indicator of performance that is not captured by the percentage correct metric. The cross-entropy or log-loss considers prediction uncertainty by including the classification probability in its calculation. In the case of N goals, and given the true labels Y for the intended goal and the inference

Fig. 5. Human inference performance on the demonstration trajectories. Plot shows mean and standard error.

probability estimates P for inference, we calculate the average log-loss across a trajectory as

$$\mathcal{L}_{\log}(Y, P) = -\frac{1}{M} \sum_{i=1}^M \sum_{j=1}^N y_{i,j} \log p_{i,j}, \quad (16)$$

where M is the number of samples in the trajectory, $y_{i,j}$ is a binary indicator of whether or not the prediction j is the correct classification for goal instance i , and $p_{i,j}$ is the probability associated with the goal j at timestep i . Note that a perfect inference model would have a log-loss of 0, and that log-loss increases as the predicted probability diverges from the intended goal. Note also that log-loss is unable to be computed for the amnesic inference, as it is not a probabilistic method.

Task Completion Time. It is important to consider how quickly the user is able to complete a task with and without assistance. Our intuition is that the intent inference method directly affects the onset of shared autonomy assistance and thus will indirectly affect the task completion time. In other words, better inference will result in correct, earlier, and stronger assistance.

Number of Control Mode Switches. Teleoperation of robotic arms (e.g., 6D control) using traditional control interfaces, such as the 3-axis joystick, require the user to switch between one of several control modes (mode switching) that are subsets of the full control space. Mode switches can become extremely challenging in the assistive domain, wherein the interfaces available to individuals with motor impairments are even more limited (e.g., 1D Sip-N-Puff). Our intuition is that the intent inference will indirectly affect the number of mode switches required to complete tasks, if earlier assistance results in fewer mode switches.

6.2 Human Inference

Figure 5 shows the performance of human intent inference on the demonstration trajectories. Percentage of time the predictions were (1) correct with confidence ($C(g) > 30\%$), (2) uncertain ($C(g) < 30\%$), and (3) incorrect with confidence ($C(g) > 30\%$) are analyzed.

The results show that inferring the intended goal of robotic arm motion is a challenging task, even for humans. The human subjects made fewer incorrect predictions as compared to correct predictions but also were inclined to indicate more uncertainty. In addition, the percentage of correct predictions were comparatively lower in the case of the change-of-intent trajectories in

Fig. 6. Performance comparison across four autonomy inference methods and two interfaces, during novice teleoperation for the subjects without motor impairments (top) and end user subjects (bottom). The percentage of time the prediction is correct and confident as the trajectory progresses is shown (errors are cumulative).

which the robot changed the intended goal during motion. Interestingly, the percentage of time inference was uncertain was unaffected by whether there was a change of intent.

6.3 Autonomy Inference

6.3.1 Novice Teleoperation. We compared the predictive accuracy of our method against the amnesic and memory-based autonomy intent inference techniques. We first analyzed the inference performance with partially observed trajectories.

Figure 6 shows the average percentage time correct and confident predictions with partially observed trajectories. The amnesic inference performed poorly and failed to get the prediction correct for a higher percentage of time as compared to the other methods. All other methods performed comparatively better, with RBII-2 performing best for the change-of-intent trajectories, the

Fig. 7. Performance comparison across the probabilistic inference methods and two interfaces during novice teleoperation for subjects without motor impairments (top) and the end user subjects (bottom). The average log-loss as the trajectory progresses is shown (errors are cumulative).

memory-based approach performing best for trajectories without a change of intent, and RBII-2 consistently outperforming RBII-1. The memory-based technique utilizes the information about the past trajectory history and struggled to recover from a change of intent toward goals. The superior performance of RBII-2 on the change-of-intent trajectories is attributed to the incorporation of our more sophisticated probabilistic model of the goal-directed user actions in the RBII-2 likelihood model, which enabled the prediction to quickly recover to the correct intended goal.

The statistically significant trends were that all techniques significantly outperformed the amnesic inference technique in all scenarios ($p < .001$), except in the case of end users with the 3-axis joystick and change-of-intent trajectories. For this case, the RBII-2 technique still performed significantly better ($p < .05$ for RBII-1, $p < .01$ for RBII-2); however, the memory-based technique lost its significance. Between the RBII-1, RBII-2, and memory-based methods, there were no statistically significant differences.

Fig. 8. Average log-loss for the autonomy inference methods and two interfaces during novice teleoperation for subjects without motor impairments (top) and the end user subjects (bottom). The average log-loss is shown, with mean and standard error.

An interesting observation is that although the 3-axis joystick provides more DoF control as compared to a Sip-N-Puff interface, it could also be more difficult to operate for people with motor impairments (e.g., because of limited finger and hand motion or hand tremor) and thus be more prone to noisy signals. Thus, there can be scenarios when performing intent inference from a *higher-dimensional* interface is comparatively *more* challenging, as seen here in our end user data (e.g., under the change of intent trajectories, the memory-based method struggled to get higher percentage time correct predictions much earlier in the trajectory in the case of end users operating the 3-axis joystick).

Figure 7 shows the performance of the autonomy intent inference methods during novice teleoperation (without autonomy assistance), measured in terms of average log-loss with partially observed trajectories. (Note that a perfect inference model would have a log-loss of 0.) The memory-based technique performed comparatively better in the case of no change-of-intent trajectories, in all scenarios. It was penalized for wrong predictions by the log-loss performance metric in two scenarios: when the user expressed a change of intention during task execution and also when the tasks involved more than two goals. In the case of change-of-intent trajectories, the RBII-2 technique outperformed the other approaches. We saw one case where the RBII-2 technique was unable to outperform RBII-1: the end user group operating the 3-axis joystick and without a change of intent. This further indicates that performing intent inference for the (arguably more informative) 3-axis joystick interface could be more challenging for some user groups.

Figure 8 shows the mean log-loss across complete trajectories. The statistically significant trends were that RBII-2 significantly outperformed the memory-based approach for nearly all change-of-intent scenarios: control subjects operating both interfaces ($p < .01$) and end user subjects operating the Sip-N-Puff interface ($p < .05$). (For end users operating the joystick, the trend was observable but not significant.) For all approaches, the log-loss was significantly ($p < .001$) higher with a change of intent (unsurprisingly) versus no change of intent.

6.3.2 Benefit of Adjustable Rationality Model of Human Actions. We examined how the rationality index parameter affects performance. To test the benefit of optimizing the rationality index

Fig. 9. Benefit of user-customized adjustable rationality in our model of user actions. Note that $\neg\beta^*$ refers to using the not-individualized rationality index values of other subjects.

for each subject, we computed the average log-loss on each subject's trajectory set using first that subjects optimized value of the rationality index β^* and then the optimized values of each other subject $\neg\beta^*$ (i.e., β is not optimized for the current subject). Figure 9 shows that the optimized values of the rationality index β^* resulted in significantly ($p < .001$) lower log-loss as compared to using the not-individualized rationality index values of other subjects $\neg\beta^*$. This indicates that customizing the rationality index values for the human actions likelihood model to each user results in better intent inference performance. Last, we also note that the optimized rationality index values β^* were comparatively higher for subjects using 3-axis joystick as compared to the subjects who used the Sip-N-Puff interface, indicating that the operators with the 3-axis joystick were able to provide more rational actions.

6.3.3 Effect of Intent Inference on Shared Autonomy Performance. We analyzed the implications of the underlying intent inference approach on shared autonomy performance (with assistance). For the shared autonomy trials, we focused only on our RBII-1 and RBII-2 techniques.

Figure 10 shows the task completion times with and without assistance. The with-assistance trials involved shared autonomy operation under the RBII-1 and RBII-2 inference techniques. The without-assistance trials imply direct teleoperation (Teleop). The task completion times were significantly higher ($p < .05$) with the Sip-N-Puff, in all cases as compared to the 3-axis joystick (not marked in the figure to reduce visual clutter). Importantly, assistance under both variants of our approach, RBII-1 and RBII-2, had significantly lower ($p < .001$) task completion times than Teleop in all scenarios. Overall, for both interfaces and subject groups, better and faster inference during shared autonomy operation with RBII-2 resulted in lower task completion times than RBII-1, and this result was significant ($p < .05$) for the subjects without motor impairments. Likewise, assistance under both inference variants resulted in significantly fewer mode switches ($p < .001$) for both interfaces and subject groups. Again, RBII-2 outperformed RBII-1 on this metric (but not significantly), indicating better assistance informed by the RBII-2 intent inference.

6.3.4 Subjective Evaluation. The subjective evaluation was performed using an experiment survey after each trial to understand how subjects felt about the intent inference performance and the assistance under shared autonomy. It also gave us insight as to whether users experienced performance differences during the robot operation under different intent inference schemes. Results are summarized in Figure 11 and Table 1. Scores of 1 and 7 indicate that the responder strongly

Fig. 10. Comparison of task completion time and number of mode switches across two variants of our autonomy inference method under shared-control operation (RBII-1, RBII-2) and without autonomy assistance (Teleop) using two interfaces. Mean and standard error over all executions are shown.

Fig. 11. Subjective evaluation shows the percentage distribution of Likert scores (indicated in color codes).

agrees or disagrees, respectively. There were noticeable difference in subjective ratings for both the intent inference and shared autonomy performance across the two techniques, with RBII-2 consistently rated better than RBII-1. Results further indicate that the subjects, especially the end user group, were satisfied with the intent inference and shared autonomy performance.

Table 1. Subjective Evaluation: Autonomy Inference and Assistance

Question	Interface	Inference	Mean Score	
			Uninjured Group	End User Group
The robot was able to correctly infer the intended goal.	3-axis	RBII-1	5.2 ± 1.3	6.1 ± 1.0
		RBII-2	5.7 ± 1.1	6.4 ± 0.6
	Sip-N-Puff	RBII-1	5.0 ± 1.1	5.5 ± 0.8
		RBII-2	5.7 ± 0.7	6.3 ± 0.6
The assistance from the robot was useful to accomplish the task.	3-axis	RBII-1	5.3 ± 1.2	6.7 ± 0.5
		RBII-2	5.7 ± 1.0	6.9 ± 0.2
	Sip-N-Puff	RBII-1	5.0 ± 1.1	5.9 ± 0.9
		RBII-2	5.8 ± 0.8	6.6 ± 0.5

7 DISCUSSION

The human inference study results indicated that inferring the intended goal of a robot is a challenging task, even for humans. One important takeaway from our study is that humans tend to make fewer incorrect predictions but also indicate more uncertainty. This has important implications for the assistive domain, since most often providing the wrong assistance is worse than providing no assistance. Thus, it is worth knowing when the autonomy inference is uncertain and to what extent. The human subjects were also quickly able to switch their prediction in the case of change of intent, although with comparatively more incorrect predictions.

The autonomy intent inference methods were evaluated on novice teleoperation with and without shared autonomy assistance and using two different control interfaces. The amnesic inference failed to generate the correct prediction with sufficiently high confidence as compared to other techniques. This demonstrates the benefit of using a Bayesian approach for inference, under which taking proper account of the prior information to compute the posterior results in better performance. Memory-based inference that utilizes the information about the past trajectory history resulted in highly confident predictions comparatively earlier in the execution. However, it also more often was wrong in its predictions with high probabilities and unable to recover when the user expressed a change of intent, particularly when the tasks involved more than two objects. Some limitations of the memory-based method are recognized in an exploratory experiment in Dragan and Srinivasa [9]. RBII-1 has the same observation source for its likelihood model and performed better than the memory-based technique in the case of change-of-intent scenarios but worse when there was no change of intent. Overall, RBII-2 outperformed other approaches in terms of faster correct predictions with higher probabilities and responded well to changing user goals. The demonstrated strengths of our approach were the ability to express uncertainty and to quickly respond in the case of change of intentions from the user, which allowed it to quickly correct and recover the robot's belief from incorrect inferences. In addition to contextual observations, probabilistic modeling and incorporating the human agent's behavior as goal-directed actions with user-customized optimization of the adjustable rationality improved the overall performance. Furthermore, even under limited interface operation, our proposed technique was able to perform well to recognize human intent.

We have shown that with the probabilistic modeling of human actions as goal-directed observations, the robotics autonomy can take advantage of indirect signals that the user implicitly provides during shared autonomy, even with lower-dimensional interfaces controlling higher-DoF robot systems. The inclusion of adjustable rationality in our model accounts for suboptimal behavior in user actions, resulting in better inference performance. Notably, suboptimal user behavior

is inherent to the assistive domain. We also demonstrated the benefit of using an optimized adjustable rationality model for users, which produced better results. One weakness of our approach is that the data likelihood function under the Boltzmann rationality model is not convex. The non-convexity makes finding optimal model parameters difficult and not efficient computationally. However, in practice, finding optimal model parameters was only required once per subject.

Under direct teleoperation, the average task completion time and the average number of mode switches were significantly higher in the case of the Sip-N-Puff interface as compared to the 3-axis joystick. The inference predictions utilized to provide assistance in shared autonomy resulted in a significant reduction in the task completion time and the number of mode switches across both interfaces. Our results further verified that the underlying intent inference approach directly affects the assistance and the overall shared autonomy performance. The overall superior performance of RBII-2 as compared to RBII-1 under shared autonomy substantiates the benefit of the probabilistic modeling of human actions in addition to the contextual observations. Better and faster inference under shared autonomy with the RBII-2 technique was able to reduce the average task completion time and the average number of mode switches under shared autonomy. The subjective evaluation results further reinforce our findings. The subjects were satisfied with the intent inference and shared autonomy performance, favoring the RBII-2 technique—especially the end user group, who are the target population in assistive robotics domain.

We emphasize the importance of thoroughly evaluating the intent inference itself when used for shared autonomy operation. There are implications of the control interface on intent inference, and intent inference affects shared autonomy performance. There is a particular benefit from probabilistic formulation of intent inference within the domain of assistive robotics—having an estimate of prediction uncertainty can be leveraged during shared-control operation, to determine whether and when to provide assistance, and by how much. For superior assistance under shared autonomy, the inference approach should provide correct predictions with high probabilities earlier in task executions.

In future work, the probabilistic modeling of other observations such as spatial goal orientations, visual cues, and distance metrics, as well as the handling of continuous goal regions, could further be explored. The effect of priors (initial goal probability distribution) and the goal transition probability are also interesting directions for future research.

8 CONCLUSION

In this work, we presented an intent inference formulation that models the uncertainty over the user's goal in an RBII algorithm to probabilistically reason about the intended goal of the user without explicit communication. Our approach is able to fuse multiple observations to reason about the intended goal of the user. We introduced a probabilistic behavior model of the user's actions that incorporates an adjustable rationality and, importantly, an individualized optimization that adapts the rationality index value for user and interface personalization. We performed probabilistic modeling of implicit observations, including our user behavior model and a distance-based likelihood observation. In user studies, we first examined human inference of robot motion and then evaluated and compared the performance of our algorithm to existing intent inference approaches using multiple control interfaces that are typically available to users in the assistive domain. We furthermore examined the impact of different variants of our inference method on shared-control operation. The findings of our study showed that in addition to contextual observations, probabilistic modeling and the incorporation of human agent behavior as goal-directed actions improves intent recognition and also demonstrate the benefit of optimizing adjustable rationality to individual users. Furthermore, even under limited control interfaces, our proposed technique was able to perform well in recognizing human intent. The results also showed that the

underlying intent inference approach directly affects shared autonomy performance, as do control interface limitations.

ACKNOWLEDGMENTS

The authors would like to thank Jessica Pedersen OTR/L, ATP/SMS for helping with the recruitment of motor-impaired subjects.

REFERENCES

- [1] Jake K. Aggarwal and Michael S. Ryoo. 2011. Human activity analysis: A review. *ACM Computing Surveys* 43, 3 (2011), 16.
- [2] Brenna D. Argall. 2016. Modular and adaptive wheelchair automation. In *Experimental Robotics*. Springer Tracks in Advanced Robotics, Vol. 109. Springer, 835–848.
- [3] Reuben M. Aronson, Santini Thiago, Thomas C. Kubler, Enkelejda Kasneci, Siddhartha Srinivasa, and Henny Admoni. 2018. Eye-hand behavior in human-robot shared manipulation. In *Proceedings of 13th ACM/IEEE International Conference on Human-Robot Interaction (HRI'18)*.
- [4] Chris L. Baker, Rebecca Saxe, and Joshua B. Tenenbaum. 2009. Action understanding as inverse planning. *Cognition* 113, 3 (2009), 329–349.
- [5] Cheng-Shiu Chung, Hongwu Wang, and Rory A. Cooper. 2013. Functional assessment and performance evaluation for assistive robotic manipulators: Literature review. *Journal of Spinal Cord Medicine* 36, 4 (2013), 273–289.
- [6] Arridhana Ciptadi, Tucker Hermans, and James M. Rehg. 2013. An in depth view of saliency. In *Proceedings of the British Machine Vision Conference (BMVC'13)*.
- [7] Gergely Csibra and György Gergely. 2007. 'Obsessed with goals': Functions and mechanisms of teleological interpretation of actions in humans. *Acta Psychologica* 124, 1 (2007), 60–78.
- [8] Anca Dragan and Siddhartha Srinivasa. 2014. Integrating human observer inferences into robot motion planning. *Autonomous Robots* 37, 4 (2014), 351–368.
- [9] Anca D. Dragan and Siddhartha S. Srinivasa. 2013. A policy-blending formalism for shared control. *International Journal of Robotics Research* 32, 7 (2013), 790–805.
- [10] Chinemelu Ezech, Pete Trautman, Louise Devigne, Valentin Bureau, Marie Babel, and Tom Carlson. 2017. Probabilistic vs linear blending approaches to shared control for wheelchair driving. In *Proceedings of the IEEE International Conference on Rehabilitation Robotics (ICORR'17)*.
- [11] Alvin I. Goldman. 1992. In defense of the simulation theory. *Mind & Language* 7, 1–2 (1992), 104–119.
- [12] Michael A. Goodrich and Dan R. Olsen. 2003. Seven principles of efficient human robot interaction. In *Proceedings of the IEEE International Conference on Systems, Man, and Cybernetics*.
- [13] Deepak Gopinath, Siddarth Jain, and Brenna D. Argall. 2017. Human-in-the-loop optimization of shared autonomy in assistive robotics. *IEEE Robotics and Automation Letters* 2, 1 (2017), 247–254.
- [14] Chien-Ming Huang and Bilge Mutlu. 2016. Anticipatory robot control for efficient human-robot collaboration. In *Proceedings of 11th ACM/IEEE International Conference on Human-Robot Interaction (HRI'16)*.
- [15] Andreas Hula, P. Read Montague, and Peter Dayan. 2015. Monte Carlo planning method estimates planning horizons during interactive social exchange. *PLoS Computational Biology* 11, 6 (2015), e1004254.
- [16] Siddarth Jain and Brenna Argall. 2014. Automated perception of safe docking locations with alignment information for assistive wheelchairs. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS'14)*.
- [17] Siddarth Jain and Brenna Argall. 2016. Grasp detection for assistive robotic manipulation. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA'16)*.
- [18] Siddarth Jain and Brenna Argall. 2018. Recursive Bayesian human intent recognition in shared-control robotics. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS'18)*.
- [19] Siddarth Jain, Ali Farshchiansadegh, Alexander Broad, Farnaz Abdollahi, Ferdinando Mussa-Ivaldi, and Brenna Argall. 2015. Assistive robotic manipulation through shared autonomy and a body-machine interface. In *Proceedings of the IEEE International Conference on Rehabilitation Robotics (ICORR'15)*.
- [20] Shervin Javdani, Henny Admoni, Stefania Pellegrinelli, Siddhartha S. Srinivasa, and J. Andrew Bagnell. 2018. Shared autonomy via hindsight optimization for teleoperation and teaming. *International Journal of Robotics Research* 37, 7 (June 2018), 717–742.
- [21] Shervin Javdani, Siddhartha Srinivasa, and Andrew Bagnell. 2015. Shared autonomy via hindsight optimization. In *Proceedings of Robotics: Science and Systems*.
- [22] Oussama Khatib. 1986. Real-time obstacle avoidance for manipulators and mobile robots. In *Autonomous Robot Vehicles*. Springer, 396–404.

- [23] Dae-Jin Kim, Rebekah Hazlett-Knudsen, Heather Culver-Godfrey, Greta Rucks, Tara Cunningham, David Portee, John Bricout, Zhao Wang, and Aman Behal. 2012. 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* 42, 1 (2012), 2–14.
- [24] Scott Kirkpatrick, C. Daniel Gelatt, and Mario P. Vecchi. 1983. Optimization by simulated annealing. *Science* 220, 4598 (1983), 671–680.
- [25] Danica Kragic, Panadda Marayong, Ming Li, Allison M. Okamura, and Gregory D. Hager. 2005. Human-machine collaborative systems for microsurgical applications. *International Journal of Robotics Research* 24, 9 (2005), 731–741.
- [26] Dylan P. Losey, Craig G. McDonald, Edoardo Battaglia, and Marcia K. O'Malley. 2018. A review of intent detection, arbitration, and communication aspects of shared control for physical human–robot interaction. *Applied Mechanics Reviews* 70, 1 (2018), 010804.
- [27] Mathew Monfort, Anqi Liu, and Brian D. Ziebart. 2015. Intent prediction and trajectory forecasting via predictive inverse linear-quadratic regulation. In *Proceedings of the 29th AAAI Conference on Artificial Intelligence (AAAI'15)*.
- [28] Chris Moore and Valerie Corkum. 1994. Social understanding at the end of the first year of life. *Developmental Review* 14, 4 (1994), 349–372.
- [29] Katharina Muelling, Arun Venkatraman, Jean-Sebastien Valois, John Downey, Jeffrey Weiss, Shervin Javdani, Martial Hebert, Andrew Schwartz, Jennifer Collinger, and Andrew Bagnell. 2015. Autonomy infused teleoperation with application to BCI manipulation. In *Proceedings of Robotics: Science and Systems*.
- [30] Katharina Muelling, Arun Venkatraman, Jean-Sebastien Valois, John E. Downey, Jeffrey Weiss, Shervin Javdani, Martial Hebert, Andrew B. Schwartz, Jennifer L. Collinger, and J. Andrew Bagnell. 2017. Autonomy infused teleoperation with application to brain computer interface controlled manipulation. *Autonomous Robots* 41, 6 (2017), 1401–1422.
- [31] Claudia Pérez-D'Arpino and Julie A. Shah. 2015. Fast target prediction of human reaching motion for cooperative human-robot manipulation tasks using time series classification. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA'15)*.
- [32] Radu Bogdan Rusu, Wim Meeussen, Sachin Chitta, and Michael Beetz. 2009. Laser-based perception for door and handle identification. In *Proceedings of the IEEE International Conference on Advanced Robotics*.
- [33] Pete Trautman. 2015. Assistive planning in complex, dynamic environments: A probabilistic approach. In *Proceedings of the IEEE International Conference on Systems, Man, and Cybernetics (SMC'15)*.
- [34] Katherine Tsui, Holly Yanco, David Kontak, and Linda Beliveau. 2008. Development and evaluation of a flexible interface for a wheelchair mounted robotic arm. In *Proceedings of the 3rd ACM/IEEE International Conference on Human Robot Interaction (HRI'08)*.
- [35] Zhikun Wang, Katharina Mülling, Marc Peter Deisenroth, Heni Ben Amor, David Vogt, Bernhard Schölkopf, and Jan Peters. 2013. Probabilistic movement modeling for intention inference in human–robot interaction. *International Journal of Robotics Research* 32, 7 (2013), 841–858.
- [36] David Whitney, Miles Eldon, John Oberlin, and Stefanie Tellex. 2016. Interpreting multimodal referring expressions in real time. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA'16)*.
- [37] Erkang You and Kris Hauser. 2011. Assisted teleoperation strategies for aggressively controlling a robot arm with 2D input. In *Proceedings of Robotics: Science and Systems*.
- [38] Brian D. Ziebart, Andrew L. Maas, J. Andrew Bagnell, and Anind K. Dey. 2008. Maximum entropy inverse reinforcement learning. In *Proceedings of the 23rd AAAI Conference on Artificial Intelligence (AAAI'08)*. 1433–1438.

Received May 2018; revised May 2019; accepted August 2019