Published in final edited form as:

Rep U.S. 2018 October ; 2018: 3905–3912. doi:10.1109/IROS.2018.8593766.

Recursive Bayesian Human Intent Recognition in Shared-Control Robotics

Siddarth Jain, Brenna Argall

Siddarth Jain and Brenna Argall are both jointly affiliated with the Department of Electrical Engineering and Computer Science, Northwestern University, Evanston, IL 60208, USA and Shirley Ryan Ability-Lab, Chicago, IL 60611, USA

Abstract

Effective human-robot collaboration in shared control requires reasoning about the intentions of the human user. In this work, we present a mathematical formulation for human intent recognition 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. In addition to contextual observations, we model and incorporate the human agent's behavior as goal-directed actions with adjustable rationality to inform the underlying intent. We examine human inference on robot motion and furthermore validate our approach with a human subjects study that evaluates autonomy intent inference performance under a variety of goal scenarios and tasks, by novice subjects. Results show that our approach outperforms existing solutions and demonstrates that the probabilistic fusion of multiple observations improves intent inference and performance for shared-control operation.

I. INTRODUCTION

A large variety of application areas involve robots operating alongside people. For example, robots can help people in domestic service, search and rescue, surgery and in driving vehicles. An important application area is assistive robotics, wherein the augmentation of human control of the robot with robotics autonomy (*shared control*) [1]–[5] can alleviate some of the control burden. One fundamental requirement for effective human-robot collaboration in shared control is *human intent recognition*. In order to meaningfully assist the human collaborator the robot has to correctly reason about the intended goal of the user from a number of potential task-relevant goals—known as the intent inference problem.

One approach to infer the user's intent could be to have the user communicate the intended goal explicitly, for example via verbal commands. However requiring explicit communication from the user could lead to ineffective collaboration and increased cognitive load [6]. Humans are very good at anticipating the intentions of others from observations, demonstrating that intentions can be inferred from non-verbal communication [7]. In this work, we investigate how the robotics autonomy can take advantage of non-verbal cues and indirect signals that the user implicitly provide when performing tasks in shared-control

operation, to facilitate faster and natural interaction. We consider the user control commands as the representative user actions and we model these actions as observations in a probabilistic behavior model to inform human intent recognition.

Non-verbal cues have been utilized for intent inference in recent research within the shared-control literature [2], [8], [9], but the problem of intent inference remains understudied, with limited evaluation. The evaluation of the inferred intent is seldom the central aim of a study, and any evaluation typically is only in simple scenarios. Incorrect predictions can result in unfortunate circumstances including collisions, while imprecise predictions can impact the robot's ability to correctly assist the user. Our target domain is shared autonomy in assistive robotics, for which intent inference is a key aspect that directly impacts performance.

In this work, we present a mathematical formulation for intent inference by modeling the uncertainty over the user's goal in a recursive Bayesian filtering framework. Our framework allows for the continuous update of the probability of each goal hypothesis. Furthermore, it enables the seamless fusion of any number of observation sources, allowing the inference to leverage rich sources of contextual information. Finally, the framework can explicitly express uncertainty in the resulting inference, which is critical to know in shared control operation as assistance towards the wrong goal could be worse than providing no assistance. We perform user studies to ground the complexity of the intent inference problem, and furthermore to evaluate and compare our approach to existing approaches of inferring intent under a variety of goal scenarios and tasks. Results indicate that our approach performs well for human intent recognition and also responds well to changing user goals, thus enabling to dynamically adjust assistance to new predictions.

The remainder of the paper is organized as follows. Related work is discussed in Section II. Section III present our framework for intent inference, with implementation details in Section IV. Sections V and IV present the experiments and results, which further are discussed in Section VII. In Section VIII, we conclude with directions for future research.

II. Related Work

Intent inference—also referred to as inference of the desired goal, target, action, or behavior—has been investigated under various settings [10]. Examples include activity recognition in the area of computer vision [11] using spatiotemporal representations, and task executions in surgical systems using Hidden Markov Models [12]. A number of human-robot interaction approaches investigate the use of gestures, expressions and gaze. For example, gaze patterns are studied to understand shared manipulation [13] and anticipatory actions [14]. Gestures along with language expressions have been utilized for object inference [15]. A step further in this direction is the area of intent-driven behaviors and intent-expressive motions, such as legible motions [16].

We are interested in investigating how shared-control systems can take advantage of the indirect signals people implicitly provide when operating the robot through a direct interface, such as a joystick. Many systems simplify the intent inference problem by assuming that the robot has access to the user's intent *a priori* [3], [17]. There exist works

that rely on explicit commands from the user to communicate intent via direct interfaces, such as GUIs [18], that focus on high-level goals. Some approaches utilize instantaneous observations (e.g. distance to goal, user command) to compute an inference confidence [2], [8], [9]. A memory-based inference approach utilized in previous works involving shared autonomy [1], [4], [16] considers the history of trajectory inputs and applies Laplace's method to approximate the distribution over goals [8]. In our approach, we consider probabilistic fusion of multiple observation sources, and also model user inputs as observations using Boltzmann rationality, which has shown to explain human behavior on various data sets [19]. In the majority of the literature on shared control, the focus of the experimental work is on the control sharing and robot assistance [1], [2], [4], [9], with the intent inference being assessed only implicitly. In this paper, we present a more extensive evaluation of intent inference than is typically seen within the shared autonomy literature, because the intent inference performance directly affects different aspects of shared autonomy.

III. Framework for Human Intent Recognition

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 assists the user towards his/her intended goal.

Problem Formulation:

Assuming the environment has a discrete set of accessible goals $\mathbf{g}$, known at runtime to both user and robot (e.g. via perception algorithms [20]), the intent inference problem is that the robot has to infer (predict) the most likely goal $g^* \in \mathbf{g}$ that the user is trying to reach. With knowledge of g^* the robotics autonomy can meaningfully assist the user in shared autonomy. Note that the user might change his/her intended goal during the execution, and thus the intent inference must update in real time so that the assistance provided can dynamically adjust.

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 have shown to be effective for inference in cognitive science [21] and human-robot interaction research [15]. 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 $\mathbf{g}$, given a set of observations (features).

We represent the goal g_t as the query variable 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,

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

which becomes,

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

upon applying the definition of conditional probability to $P(g_t, g_{t-1} | \theta_{0:t-1})$ and the Markov assumption, ($g_t \perp \theta_{0:t-1} | g_{t-1}$). The computation of $b_t(g)$ thus is a recursive update, and so encodes memory of prior goal distributions. Furthermore, $P(g_t | g_{t-1})$ 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 , ($\theta^i \perp \theta^j | g$), $\forall i \neq j$. Thus, Equation 3 becomes

$$b_t(g) = P(g_t | \theta_{0:t}) \propto \prod_{\theta_t \in \Theta_t} P(\theta_t | g_t) \sum_{g_{t-1} \in g} P(g_t | g_{t-1}) b_{t-1}(g_{t-1}). \quad (4)$$

Algorithm 1: Recursive Bayesian Intent Inference

```

1 Given Goals  $g$ 
2 Initialize  $P(g_{t=0}) \forall g \in g$ 
3 Initialize  $b_{t=0}(g) \leftarrow P(\Theta_{t=0} | g_{t=0}) P(g_{t=0}) \forall g \in g$ 
4 Normalize  $b_{t=0}$ 
5 while executing do
6   Observe  $\Theta_t$ 
7   foreach  $g \in g$  do
8      $b_t(g) \leftarrow$ 
9        $\prod_{\theta_t \in \Theta_t} P(\theta_t | g_t) \sum_{g_{t-1} \in g} P(g_t | g_{t-1}) b_{t-1}(g_{t-1})$ 
9   end
10  Normalize  $b_t$ 
11  Update intent  $g_t^* \leftarrow \arg \max_{g \in g} b_t(g)$ 
12 end

```

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 $P(g_{t=0})$, $\forall g \in \mathcal{G}$, initially represents the robot's belief over the goals. The beliefs then are continuously updated, by computing the posteriors $P(g_t | \Theta_{0:t})$, $\forall g \in \mathcal{G}$, as more observations become available.

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

$$g_t^* = \arg \max_{g \in \mathcal{G}} P(g_t | \Theta_t). \quad (5)$$

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 towards 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,

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

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

IV. 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.

A. Autonomy Inference

Our approach, Recursive Bayesian Intent Inference (RBII), allows for the seamless fusion of any number of observations to perform human intent recognition. In order to examine how the incorporation of multiple 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 [1], [2], [4], [8]. We compute proximity θ^d as the Euclidean distance between the current position of the robot x_t and the goal 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, in addition to proximity to the goal, we model the actions of the human agent. Following a model of human action from cognitive science [21], we model the user as Boltzmann-rational in their actions to reach a goal g (discussed in Section IV-B).

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

$$P(g_t = g^i | g_{t-1} = g^j) = \begin{cases} 1 - \Delta & \text{if } i = j \\ \Delta/(n-1) & \text{otherwise} \end{cases} \quad (7)$$

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 initializes the probability distribution over goals to be uniform, and sets $\Delta = 0.1$.

B. Human Action Model as an Observation

We are interested in investigating how we can utilize for intent recognition the indirect signals people implicitly provide to operate the robot. Our RBII-2 implementation 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 data sets [19].

We incorporate adjustable rationality in a probabilistic behavior model such that at 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,

$$P(u_g | s, g) \propto \exp(\beta \cdot Q_g(s, u_g)) \quad (8)$$

where $Q_g(s, u_g)$ denotes the Q-value when the intended goal is g . β is a rationality index (discussed further below) that controls how diffuse are the probabilities. 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 vector norm. A moving average filter with a two second time window is applied to the observations, in order to consider a brief history of observations and reduce any undesired oscillations due to noisy or corrective control signals.

A critical detail is to include adjustable rationality in the model, as in reality a number of factors might induce sub-optimality—for example, limitations and challenges imposed by the interfaces to control high-DoF robotic systems, cognitive or physical impairments, and environment factors such as obstacles or distractions. Adjustable rationality is represented by the rationality index β and a larger value of β implies more rationality. We fit the

rationality index parameter by performing simulated annealing optimization to minimize the prediction log-loss for intent recognition on the data gathered in a pilot experiment ($\beta = 0.045$).

C. Assistance under Shared Autonomy

In this section, we discuss implementation details of the shared autonomy that utilizes the inferred intent. We implement a blending-based paradigm to provide assistance,

$$\mathbf{u}_{blend} = \mathbf{u}_h \cdot (1 - \alpha) + \mathbf{u}_r \cdot \alpha, \quad (9)$$

where $\mathbf{u}_h$ denotes the user control command, $\mathbf{u}_r$ the autonomy control command generated under a potential field policy [22] and $\mathbf{u}_{blend}$ is 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 $g \in \mathbf{g}$ such that g is treated as an attractor and all the other goals $\mathbf{g} \setminus g$ as repellers. $\alpha \in [0, 1]$ is a blending factor which 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,

$$\alpha = \begin{cases} 0 & C(\mathbf{g}) \leq \delta_1 \\ \frac{\delta_3}{(\delta_2 - \delta_1)} \cdot C(\mathbf{g}) & \delta_1 < C(\mathbf{g}) \leq \delta_2 \\ \delta_3 & C(\mathbf{g}) > \delta_2 \end{cases} \quad (10)$$

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

V. Experiments on Intent Inference

Our experimental work aimed to evaluate the performance of our intent inference algorithm, as well as impact of intent inference on shared autonomy. We performed two human subject studies, that aim (i) to characterize the complexity and variability of the intent inference problem, (ii) to compare the inference performance of our approach to existing approaches used in shared autonomy and (iii) to evaluate the impact of inference on shared autonomy assistance.

Our research platform for the designed experiments was the MICO assistive robotic arm (Kinova Robotics, Canada), a 6-DoF manipulator with a 2 finger gripper. The control interface used in the study was a 3-axis joystick that is typically been utilized for operating robotic arms.

Subject Allocation:

We recruited 12 able-bodied subjects from the local community (5 male, 7 female, aged 19–35). 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, approved by Northwestern University Institutional Review Board.

In addition to the two variants of our algorithm detailed in Section IV-A (RBII-1 and RBII-2), for comparative purposes we also implemented two approaches utilized in previous shared autonomy works—Amnesic Inference [2], [8], [9] and Memory-based Inference [1], [4], [8], [16].

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,

$$c(g) = \max\left(0, 1 - \frac{d}{D}\right) \quad (11)$$

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* [8], because it ignores all information except the instantaneous observations. In our implementation, d is the Euclidean distance $\|x_g - x_r\|$ between the current position of the robot end-effector x_r and the goal x_g . D is set to 1.0 m (maximum reach of the MICO robotic arm).

Memory-based Inference:

The memory-based prediction [8] 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, the probability of a trajectory reaching towards a specific goal g is given as $P(\xi | g) \propto \exp(-c_g(\xi))$; that is, the probability of the trajectory decreases exponentially with cost. It is assumed that the cost is additive along the trajectory. Such a solution becomes too expensive to compute in high-dimensional spaces (e.g. for robotic manipulation), and so [8] estimates the most likely goal by approximating the integral over trajectories using Laplace's method and first order approximation,

$$g^* = \arg \max_{g \in \mathcal{G}} \frac{\exp(-c_g(\xi_s \rightarrow x) - c_g(\xi_x^* \rightarrow g))}{\exp(-c_g(\xi_s^* \rightarrow g))} P(g) \quad (12)$$

where ξ^* is the optimal trajectory, s is the starting pose of robot, x is the current pose, and g is the pose of the goal. In practice the cost c_g is often the Euclidean distance and the goal probabilities are initialized with a uniform prior [1], [4], which we do in our implementation as well.

A. Human Inference Study

Humans are very good at anticipating the intentions of others with non-verbal communication and by observation of goal-directed actions [7]. 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. Related work has addressed the topic of human interpretation of robot trajectories generated by autonomy policies [16] for intent inexpressive motions. We however are interested in robot motion generated by *human* control commands, given our target domain of shared autonomy. In our experiments, the robot trajectories were pre-recorded demonstrations by a human expert operating a 3-axis joystick. We furthermore are interested in how a *change-of-intent* can affect the intent inference. We therefore consider two types of robot motion: (i) No-change-of-intent, where the robot motion maintained a single goal from start to end, and (ii) change-of-intent, where the robot motion switched during the execution the goal it was reaching towards.

Design: The study involved a variety of goal scenarios with varied complexity: in total, 8 different scenarios involving 2 to 4 potential goals (Figure 2, left). For each goal scenario, one change-of-intent demonstration was recorded (8 in total) as well as a no-change-of-intent trajectory 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.

Protocol: The subjects were given a training period in which they observed one pre-recorded teleoperation trajectory towards a goal, to get familiarized with the robot motion capabilities. The 30 teleoperated trajectories for the 8 goal scenarios then were executed (replayed) in counterbalanced order on the robot. For each trajectory, we tasked the subjects to observe the motion and predict the intended goal (Figure 2, right) by verbally mentioning the “Object Name” (e.g., “cup”). The subjects could 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”. The inference was registered as uncertain at the start of the trajectory motion, and the subject responses were recorded by the experimenter via a button press.

B. Autonomy Inference Study

Our second study aimed to evaluate how well the autonomy could infer the intent of a user under *novice teleoperation* and how the inferred intent affected shared autonomy.

Design: We adopted a within-subjects experimental design to evaluate how well the autonomy could infer the intent of the user under *novice teleoperation with and without assistance*. Four intent inference approaches were evaluated: (i) Amnesic (ii) Memory-based (iii) RBII-1 and (iv) RBII-2. Four tasks of varied complexity were used to evaluate the performance (Figure 3). Three tasks involved object retrieval where the goal was the approach grasp pose $\in \mathbb{R}^6$ on the objects. The fourth task involved pouring and placing operations where the goals were the pose of the initiation of the pour over the bowls and placing in the dish rack.

Protocol: The subjects were given a training period in which they got familiar using a 3-axis joystick to operate the robot. First, the subjects teleoperated the robot *without assistance*. They performed the tasks shown in Figure 3 and were instructed to (i) complete each goal in every task setup and(ii) for each task perform one additional trial in which they *change goal* during the task execution. The change was recorded with a time stamp via a button press. All four approaches for intent inference computed the intent online as the user teleoperated the robot. In the next stage, all the trials were performed again but now *with assistance* under shared autonomy, once with RBII-1 and then using RBII-2.

VI. Analysis and Results

We first discuss performance measures and then present the analysis with the results. For each performance measure, one factor repeated measure ANOVA (Analysis of Variance) was performed to determine significant differences ($p < 0.05$) between the intent inference approaches. Once the significance was established, multiple post-hoc pairwise comparisons were performed by using Bonferroni Confidence interval adjustments. For all figures, the notation * implies $p < 0.05$ * * implies $p < 0.01$, * * * implies $p < 0.001$.

A. Performance Measures

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

Cross-entropy Loss (log-loss): Assessing the uncertainty of a prediction is an important indicator of performance which is not captured by the percent 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 probability estimates P , we calculate the average log-loss across a trajectory as,

$$\mathcal{L}_{\log}(Y, P) = -\frac{1}{M} \sum_{i=1}^M \sum_{j=1}^N y_{i,j} \log p_{i,j} \quad (13)$$

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 at time step 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. Log-loss is unable to be computed for amnesic inference, as it is not a probabilistic method.

Assistance Onset Time: In shared autonomy it is important to consider how early the autonomy can assist the user towards the intended goal. The initial onset of assistance depends on how quickly the confidence (Equation 10) in the intended goal rises and exceeds the lower bound threshold δ_1 on the confidence.

Task Completion Time: Our intuition is that the intent inference affects the shared autonomy and thus will indirectly affect the task completion time. That is, better inference will result in correct, earlier, and stronger assistance.

Number of Control Mode Switches: Teleoperation of robotic arms using lower-dimensional control interfaces (e.g. 3-axis joystick to operate a 6-DoF arm), requires the user to switch between one of several control modes (mode switching) that are subsets of the full control space. Our intuition is that the intent inference will affect the number of mode switches required to complete tasks, if earlier assistance results in fewer mode switches.

B. Human Inference

Figure 4 shows the performance of human intent inference and comparison with the autonomy intent inference. The percentage of time predictions were (i) correct with confidence ($C(g) > 30\%$), (ii) uncertain ($C(g) < 30\%$) and (iii) 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. The percentage of correct predictions were comparatively lower in the case of the change-of-intent trajectories. Interestingly, the percentage of time uncertain was unaffected by whether there was a change-of-intent—both for the human subjects and all inference algorithms.

In the case of the autonomy inference, overall RBII-2 performed better than all other methods with higher percentage of time correct and more confident predictions—in the case of trajectories both with and without a change-of-intent. Interestingly, RBII-2 was more often correct than were the human subjects. This finding could be attributed to the result that the human subjects tend to indicate more uncertainty rather than be incorrect. The memory-based inference expressed uncertainty a smaller percentage of the time, with the result of both higher percent correct but also higher percent incorrect. Amnesic inference by far expressed the highest amount of uncertainty (typically until the very end of an execution), resulting in the lowest percent correct but, somewhat surprisingly, also the lowest percent incorrect.

C. Autonomy Inference

Novice Teleoperation: Figure 5 shows the performance of the autonomy intent inference methods, during novice teleoperation without autonomy assistance. For percentage of time correct, the amnesic inference failed to get the prediction correct for a higher percentage of time as compared to other methods ($p < 0.001$). Memory-based performed better in the case of no-change-of-intent trajectories and RBII-2 in the case of the change-of-intent trajectories, however, the results were not significant. For log-loss, the memory-based method resulted in wrong predictions with high probabilities in the case of change-of-intent trajectories and thus was penalized by the log-loss performance metric. Overall, RBII2 was better at identifying the correct goal with higher probabilities, and was significantly better than the memory-based method in the case of change-of-intent trajectories ($p < 0.01$). Figure

5 (right) shows the initial onset of assistance in the case of correct predictions. The amnesic inference was significantly slower than other methods ($p < 0.001$) for the onset of assistance in all cases. Overall, the memory-based method was able to infer the intended goal earlier than others methods. Interestingly, all methods except memory-based experienced a delay in providing assistance when there was a change-of-intent. Our further analysis of RBII-1 and RBII-2 under shared control however confirmed that even with this delay, both methods were able to recover from a change-of-intent early enough to provide assistance that improved task performance (Figure 6).

Effect of Inference on Shared Autonomy Performance: Figure 6 shows the overall comparative performance of the RBII-1 and RBII-2 intent inference approaches with shared autonomy assistance. For the task completion time, both RBII-1 and RBII-2 were significantly better than teleoperation ($p < 0.001$). Furthermore, RBII-2 resulted in significantly lower task completion time as compared to RBII-1 approach ($p < 0.05$), indicating better assistance in shared autonomy operation. For the average number of control mode switches, teleoperation without assistance had significantly higher numbers of mode switches ($p < 0.001$) than RBII-1 and RBII-2, showing that assistance resulted in fewer mode switches. RBII-2 resulted in fewer average number of mode switches than RBII-1, but this result was not significant.

VII. Discussion

The 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 both to make fewer incorrect predictions but also indicate more uncertainty. This has important implications for assistive domains, since most often providing the wrong assistance is worse than providing no assistance. Thus, there is worth in 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, though with comparatively more incorrect predictions.

The autonomy intent inference methods were evaluated both on demonstrated trajectories and on novice teleoperation (with and without shared autonomy assistance). The amnesic inference failed to get the prediction correct for a higher percentage of time, as it tended to switch to the correct prediction at the end of trajectories. Memory-based inference resulted in comparatively faster onset of assistance. However interestingly, it was often wrong in its predictions with high probabilities when the user expressed a change-of-intent. Some other limitations of the memory-based method are recognized in an exploratory experiment [8]. RBII-1, although comparatively slower, made correct predictions more often. RBII-2 outperformed other approaches in terms of faster correct predictions with higher probabilities, indicating that a fusion of observations improved performance. RBII-2 also responded well to changing user goals, thus enabling the system to dynamically adjust its assistance to new predictions.

We have shown that with the probabilistic modeling of human actions as observations the robotics autonomy can take advantage of indirect signals that the user implicitly provides during shared-control operation, even with lower dimensional interfaces controlling higher-

DoF robot systems. The inclusion of adjustable rationality in our model can account for sub-optimal behavior in user actions due to any number of reasons inherent to assistive domains. In future work, we intend to tune a customized value of rationality index for each user and also will explore performance in the case of more limited interfaces (e.g. Sip-N-Puff) that are available to people with motor-impairments. We note that other observations, such as spatial goal orientations and L_1 distance could further be explored. Eye gaze could also be utilized to reveal intent, however the data quality depend on the calibration and projection method errors of eye trackers [13], [14]. The effect of priors (initial goal probability distribution) and the goal transition probability are also interesting directions for future.

Our results further verified that the underlying intent inference approach directly affects the assistance provided and the overall shared-autonomy performance. RBII-2 resulted in significantly faster task completion times and a reduced number of control mode switches, which is of critical importance in assistive domains. For better shared-autonomy performance, the inference approach should reason about the uncertainty in its predictions and should provide correct predictions with high probabilities earlier in task executions. To this end, we emphasize the importance of thoroughly evaluating any intent inference scheme used in shared autonomy operation.

VIII. Conclusions

In this work we presented a formalism for human intent recognition that models the uncertainty over the user's goal in a recursive Bayesian algorithm. In user studies, we examined human inference on robot motion and furthermore compared the performance of our algorithm to existing intent inference approaches (both with and without shared autonomy assistance). We furthermore demonstrated the effect of intent inference on assistance in shared autonomy. Results validated that our algorithm was able to provide more correct and confident predictions in comparison to existing approaches, under different task and goal scenarios and it also responds well to changing user goals. Furthermore, the fusion of multiple observations and probabilistic modeling of human actions improved prediction accuracy, resulting in faster assistance. Our future work will evaluate our approach within a larger study involving more limited control interfaces and motor-impaired subjects.

ACKNOWLEDGMENT

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 number R01EB019335. 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 partially supported by grant from U.S. Office of Naval Research under the Award Number N00014-16-1-2247.

References

- [1]. Muelling K, Venkatraman A, Valois J-S, Downey JE, Weiss J, Javdani S, Hebert M, Schwartz AB, Collinger JL, and Bagnell JA, "Autonomy infused teleoperation with application to brain computer interface controlled manipulation," *Autonomous Robots*, pp. 1–22, 2017.

- [2]. Gopinath D, Jain S, and Argall BD, "Human-in-the-loop optimization of shared autonomy in assistive robotics," *IEEE Robotics and Automation Letters*, vol. 2, no. 1, pp. 247–254, 2017. [PubMed: 30662953]
- [3]. Kim D-J, Hazlett-Knudsen R, Culver-Godfrey H, Rucks G, Cunningham T, Portee D, Bricout J, Wang Z, and Behal A, "How autonomy impacts performance and satisfaction: Results from a study with spinal cord injured subjects using an assistive robot," *IEEE Transactions on Systems, Man, and Cybernetics-Part A: Systems and Humans*, vol. 42, no. 1, pp. 2–14, 2012.
- [4]. Javdani S, Srinivasa S, and Bagnell A, "Shared autonomy via hindsight optimization," in *Proceedings of Robotics: Science and Systems*, 2015.
- [5]. Trautman P, "Assistive planning in complex, dynamic environments: a probabilistic approach," in *Proceedings of IEEE International Conference on Systems, Man, and Cybernetics (SMC)*, 2015.
- [6]. Goodrich MA and Olsen DR, "Seven principles of efficient human robot interaction," in *Proceedings of IEEE International Conference on Systems, Man and Cybernetics*, 2003.
- [7]. Goldman AI, "In defense of the simulation theory," *Mind & language*, vol. 7, no. 1–2, pp. 104–119, 1992.
- [8]. Dragan AD and Srinivasa SS, "A policy-blending formalism for shared control," *The International Journal of Robotics Research*, vol. 32, no. 7, pp. 790–805, 2013.
- [9]. Argall BD, "Modular and adaptive wheelchair automation," in *Experimental Robotics*, 2016, pp. 835–848.
- [10]. Losey DP, McDonald CG, Battaglia E, and O'Malley MK, "A review of intent detection, arbitration, and communication aspects of shared control for physical human–robot interaction," *Applied Mechanics Reviews*, vol. 70, no. 1, p. 010804, 2018.
- [11]. Aggarwal JK and Ryoo MS, "Human activity analysis: A review," *ACM Computing Surveys (CSUR)*, vol. 43, no. 3, p. 16, 2011.
- [12]. Tao L, Zappella L, Hager GD, and Vidal R, "Surgical gesture segmentation and recognition," in *International Conference on Medical Image Computing and Computer-Assisted Intervention*, 2013.
- [13]. Aronson RM, Santini T, Kübler TC, Kasneci E, Srinivasa S, and Admoni H, "Eye-hand behavior in human-robot shared manipulation," in *Proceedings of the 2018 ACM/IEEE International Conference on Human-Robot Interaction*, 2018.
- [14]. Huang C-M and Mutlu B, "Anticipatory robot control for efficient human-robot collaboration," in *Proceedings of 11th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, 2016.
- [15]. Whitney D, Eldon M, Oberlin J, and Tellex S, "Interpreting multimodal referring expressions in real time," in *Proceedings of IEEE International Conference on Robotics and Automation (ICRA)*, 2016.
- [16]. Dragan A and Srinivasa S, "Integrating human observer inferences into robot motion planning," *Autonomous Robots*, vol. 37, no. 4, pp. 351–368, 2014.
- [17]. Jain S, Farshchiansadegh A, Broad A, Abdollahi F, MussaIvaldi F, and Argall B, "Assistive robotic manipulation through shared autonomy and a body-machine interface," in *Proceedings of the IEEE International Conference on Rehabilitation Robotics (ICORR)*, 2015.
- [18]. Tsui K, Yanco H, Kontak D, and Beliveau L, "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)*, 2008.
- [19]. Hula A, Montague PR, and Dayan P, "Monte carlo planning method estimates planning horizons during interactive social exchange," *PLoS computational biology*, vol. 11, no. 6, p. e1004254, 2015. [PubMed: 26053429]
- [20]. Rusu RB, Holzbach A, Beetz M, and Bradski G, "Detecting and segmenting objects for mobile manipulation," in *Proceedings of IEEE International Conference on Computer Vision Workshops (ICCV Workshops)*, 2009.
- [21]. Baker CL, Saxe R, and Tenenbaum JB, "Action understanding as inverse planning," *Cognition*, vol. 113, no. 3, pp. 329–349, 2009. [PubMed: 19729154]
- [22]. Khatib O, "Real-time obstacle avoidance for manipulators and mobile robots," in *Autonomous robot vehicles*, 1986, pp. 396–404.

Fig. 1.
Human intent recognition in shared autonomy.

Fig. 2.

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 the robot motion.

Fig. 3.
Task scenarios and associated grasp pose on goals used to evaluate intent inference under novice teleoperation.

Fig. 4.

Human and autonomy inference performance on the demonstration trajectories. Plots show mean and standard error. *Left:* Percentage of time inference is correct and confident. *Middle:* Percentage of time inference is uncertain. *Right:* Percentage of time inference is incorrect and confident.

Fig. 5.

Performance of different autonomy inferences on novice teleoperation (without assistance).

Plots show mean and standard errors. *Left:* Percentage of time autonomy inferences are correct. *Middle:* Average Log Loss for goal inferences. *Right:* Normalized initial assistance onset time, for correct predictions.

Fig. 6.

Performance of intent inference and effect on shared autonomy assistance. Plots show mean and standard errors. *Left*: Average task completion time. *Right*: Total average number of control mode switches.